

Multidisciplinary Problem Solving Environments for Computational Science

Elias N. Houstis, John R. Rice, Naren Ramakrishnan
Department of Computer Sciences, Purdue University
West Lafayette, IN 47907-1398
email: {enh,jrr,naren}@cs.purdue.edu

Tzvetan Drashansky
Juno Online Services, L.P.
120 West 45th Street
New York, NY 10036
email: ttd@staff.juno.com,ttd@deshaw.com

Sanjiva Weerawarana
IBM TJ Watson Research Center
Hawthorne, NY 10532
email: sanjiva@watson.ibm.com

Anupam Joshi
Department of Computer Engineering & Computer Science
University of Missouri, Columbia, MO 65211
email: joshi@trinetra.cecs.missouri.edu

C.E. Houstis
Department of Computer Science
University of Crete
email: houstis@csi.forth.gr

Abstract

The process of prototyping is part of every scientific inquiry, product design, and learning activity. Economic realities require fast, accurate prototyping using knowledge and computational models from multiple disciplines in science and engineering [1; 3; 49; 15; 47]. Thus rapid multidisciplinary problem solving or prototyping is a new grand challenge for Computational Science and Engineering (CS&E) [11; 48;

To appear in *Advances in Computers*. Work supported in part by NSF awards ASC 9404859 and CCR 9202536, ARPA ARO award DAAH04-94-G-0010DARPA, Intel Corporation and Purdue Research Foundation.

14]. In this prototyping scenario it is safe to assume that The Network (*The Net*) is the Computer [29] consisting of geographically distributed computational units, ranging from workstations to massively parallel machines and physical instruments, and software resources (i.e. libraries and problem solving environments (PSEs)). Moreover, *The Web* and its technologies can be viewed as an object-oriented operating system kernel that allows the development of an MPSE as a distributed application utilizing resources and services from many sources. The realization of this vision will require the formulation and development of new mathematical and software frameworks for PSEs [27] and multidisciplinary PSEs (MPSEs) including the tools, enabling technologies, and underlying theories needed to support physical prototyping in the classroom, laboratory, desk, and factory. MPSEs will be network based, adaptable and intelligent with respect to end-users and hardware platforms and will use collaborating software systems and agent based techniques. They will allow wholesale reuse of scientific software and provide a natural approach to parallel and distributed problem solving. In this chapter, we describe the research that is needed to realize the MPSE concept and present a software architecture of an MPSE framework based on the *agent* approach supported by domain specific knowledge bases in a networked computing setting. The viability of such a framework is demonstrated for partial differential equation applications.

1. INTRODUCTION

It is predicted that by the beginning of the next century, the available computational power will enable anyone with access to a computer to find an answer to any question that has a known or effectively computable answer. In [9] we have made several research recommendations for the development of *problem solving environment* (PSE) technologies. We believe that PSE technologies will contribute toward the realization of this prediction for physical modeling in order to provide students, scientists, and engineers with environments that allow them to spend more time doing science and engineering, and less time struggling with the details of the underlying computation.

The predicted growth of computational power and network bandwidth suggests that computational modeling and experimentation will be one of the main tools in *big* and *small* science. In this scenario, computational modeling will shift from the current single physical component design to the design of a whole physical system with a large number of components that have different shapes, obey different physical laws and manufacturing constraints, and interact with each other through geometric and physical interfaces. For example, the analysis of an engine involves the domains of *thermodynamics* (gives the heat flows throughout the engines), *reactive fluid dynamics*, (gives the behavior of the gases in the piston-cylinder assemblies), *mechanics* (gives the kinematic and dynamic behaviors of pistons, links, cranks, etc.), *structures* (gives the stresses and strains on the parts) and *geometry* (gives the shape of the components and the structural constraints). The design of the engine requires that these different domain-specific analyses interact in order to find the final solution. The different domains share common parameters and interfaces but each has its own parameters and constraints. We refer to these multi-component based physical systems as *multidisciplinary applications* (MAs) and their PSEs as *multidisciplinary problem solving environments* (MPSEs). In

the following, we will specifically concentrate on such environments for physical systems modeled by partial differential equations (PDEs) and ordinary differential equations (ODEs).

The realization of the above scenario, which is expected to have significant impact in industry, education, and training, will require (i) the development of new algorithmic strategies and software for managing the complexity and harvesting the power of the expected high performance computing and communication (HPCC) resources, (ii) PSE technology to support programming-in-the-large and reduce the overhead of HPCC computing, (iii) an *active programming* paradigm capable of realizing the interactions between the physical and software components in a reusable mode, (iv) the selection of computational/hardware resources and the determination of their parameters for the specified simulations. Some of the objectives of this chapter include review of the research issues involved in the development of MPSEs, identification of a framework for the numerical simulation of multidisciplinary applications, and the specification of some enabling theories and technologies needed to support and realize this framework in targeted applications. The MPSE is the software implementation of this framework. It is assumed that its elements are discipline-specific problem solving environments (PSEs) and libraries. The MPSE design objective is to allow the natural specification of multidisciplinary applications and their simulation with interacting PSEs through mathematical and software interfaces across networks of heterogeneous computational resources.

Another safe prediction is that the future global information infrastructure (GII) will impact many aspects of life, including the way we learn and do science, access information, work, and collaborate. It will allow computing everywhere [17]. Learning and training simulators will be part of every classroom and laboratory. The concept of the classroom, the laboratory, and the student/scientist/engineer desk environments will evolve to some virtual form based on an array of multimedia devices. These virtual environments, sometimes called *collaboratories*, can be implemented naturally using the MPSE technology. The discipline-specific PSEs will be used to build learning and training simulators in some areas of computational science and engineering.

For the scope of this discussion, the network (*the Net*) is assumed to be the host for multidisciplinary problem solving. We assume existing network software infrastructure to support distributed applications, PSEs, libraries of solvers, and distance learning and collaboratory services over the Net. We envisage that all these resources and services will take the form of special servers in the network which are constantly updated by their creators. The problem solving power of math/engineering libraries will be encapsulated into PSEs that users will access over the Net as a routine part of science and engineering. We have already developed such servers to experiment with web based dissemination and use of PDE software and PSEs [27; 51].

This chapter is organized as follows. We start by defining the terms of PSEs and MPSEs and discuss the research issues involved in Sections 2 and 3. Then in Sections 3 and 4 we present an agent-based approach for designing and building MPSEs together with the problems of resource and solution methodology selection. We then introduce, in Section 6, the *SciAgents* system, which provides the solver and mediator agents for PDE based application MPSEs. In Section 7, we describe

PYTHIA, a multiagent advisory system that uses a distributed knowledge corpus. We show how these various agents can interact with each other to automate the process of solving multi-physics problems. Finally, in Section 8 we consider in detail two case studies using current prototypes that show the applicability and the potential of the MPSE concept and that demonstrate our approach for its implementation.

2. DOMAIN SPECIFIC PSES

Even in the early 1960s, scientists had begun to envision problem-solving computing environments not only powerful enough to solve complex problems, but also able to interact with users on human terms. The rationale of the PSE research is that the dream of the 1960s will be the reality of the 21st: High performance computers combined with better algorithms and better understanding of computational science have put PSEs well within our reach.

What are PSEs? A PSE is a computer system that provides all the computational facilities needed to solve a target class of problems. These facilities include advanced solution methods, automatic selection of appropriate methods, use of the application's language, selection of appropriate hardware and powerful graphics, symbolic and geometry based code generation for parallel machines, and programming-in-the-large. The scope of a PSE is the extent of the problem set it addresses. This scope can be very narrow, making the PSE construction very simple, but even what appears to be a modest scope can be a serious scientific challenge. For example, we have created a PSE for bioseparation analysis [13] which has a narrow scope, but is still a complex challenge as we incorporate both a computational model and an experimental process supported by physical laboratory instruments. We are also creating a PSE called PDELab [52] for partial differential equations (PDEs). This is a far more difficult area than bioseparation and the resulting PSE will be less powerful (less able to solve all the problems posed to it), less reliable (less able to guarantee the correctness of results), but more generic (more able to attempt to solve a broad class of PDE problems). Nevertheless, PDELab will provide a quantum jump in the PDE solving power delivered into the hands of the working scientist and engineer. What are the PSE related research issues to be addressed? A substantive research effort is needed to lay the foundations for building PSEs. This effort should be directed towards i) a PSE kernel for building scientific PSEs, ii) a knowledge based framework to address computational intelligence issues for PDE based PSEs, iii) infrastructure for solving PDEs, and iv) parallel PDE methodologies and virtual computational environments.

3. MPSES FOR PROTOTYPING OF PHYSICAL SYSTEMS

If PSEs are so powerful, what then is an MPSE? In simple terms, an MPSE is a framework and software kernel for combining PSEs for tailored, flexible multidisciplinary applications. A physical system in the real world normally consists of a large number of components where the physical behavior of each component is modeled by a PDE or ODE system with various formulations for the geometry, PDE, ODE, interface/boundary/linkage and constraint conditions in many different geometric regions. One needs a mathematical/software framework which, first, is applicable to a wide variety of practical problems, second, allows for software reuse in order

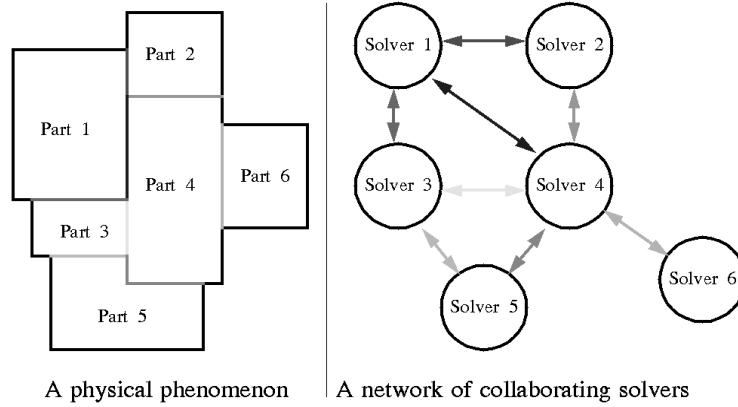


Fig. 1: The representation of a physical phenomenon (left) by a mathematical network (right) of solvers (circles) and interface conditions (arrows).

to achieve lower costs and high quality, and, finally, is suitable for some reasonably fast numerical methods. Most physical systems can be modeled as a mathematical network whose nodes represent the physical components in a system or artifact. Each node has a mathematical model of the physics of the component it represents and a solver agent for its analysis. The relationship between a physical phenomenon and the mathematical network is illustrated in Figure 1. Individual components are chosen so that each node corresponds to a simple PDE or ODE problem defined on a regular geometry.

What are the mathematical network methodologies required? What are the research issues? There exist many standard, reliable PDE/ODE solvers that can be applied to these local node problems. In addition there are nodes that correspond to interfaces (e.g. ODEs, objective functions, relations, common parameters and their constraints) that model the collaborating parts in the global model. To solve the global problem, we let these local solvers collaborate with each other to relax (i.e., resolve) the interface conditions. An interface controller or mediator agent collects boundary values, dynamic/shape coordinates, and parameters/constraints from neighboring subdomains and adjusts boundary values and dynamic/shape coordinates to better satisfy the interface conditions. Therefore, the network abstraction of a physical system or artifact allows us to build a software system which is a network of collaborating well defined numerical objects through a set of interfaces. Some of the theoretical issues of this methodology are addressed in [39; 5] for the case of collaborating PDE models. The results obtained so far verify the feasibility and potential of network-based prototyping.

What are the software methodologies for implementing the mathematical network? What are the research issues? A successful architecture for PSEs requires heavy reuse of existing software within a modular, object oriented framework consisting of layers of objects. The kernel layer integrates those components common to most PSEs or MPSEs for physical systems. We observe that this architecture can be combined with an agent oriented paradigm and collaborating solvers [5] to create an MPSE as a powerful prototyping tool using a variety of infrastructure tools. The

designs for MPSEs must be application and user driven. We should not restrict our design just to use the current technology of high performance computers, powerful graphics, modular software engineering, and advanced algorithms. We should see an MPSE as delivering problem solving services over the Net. This viewpoint leads naturally to collaborating, agent based methodologies. This, in turn, leads to very substantial advantages in both software development and quality of service as follows. We envision that a user of MPSE will receive at his location only the user interface. Thus, the MPSE server will export to the user's machine an agent that provides an interactive user interface while the bulk of the software and computing is done at the server's site using software tailored to a known and controlled environment. The server site can, in turn, request services from specialized resources it knows, e.g., a commercial PDE solver, a proprietary optimization package, a 1000 node supercomputer, an ad hoc collection of 122 workstations, a database of physical properties of materials. Each of these resources is contacted by an agent from the MPSE and all of this can be managed without involving the user, without moving software to arbitrary platforms, and without revealing source codes.

What are the design objectives of an MPSE for physical system design? What are the research issues? These mathematical networks can be very big for major applications. For a realistic vehicle simulation, there are perhaps 100 million variables and many different time scales. This problem has very complex geometry and is very non-homogeneous. The answer is 20 gigabytes in size and requires about 10 teraflops to compute. An answer is a data set that allows one to display an accurate approximate solution at any point. This data set is much smaller than the computed numerical solution. The mathematical network has about 10,000 subdomains and 35,000 interfaces. A software network for this simulation is a natural mapping of a physical system and simulates how the real world evolves. This allows software reuse for easy software update and evolution, things that are extremely important in practice. The real world is so complicated and diverse that we believe it is impractical to build monolithic, universal solvers for such problems. In this application each physical component can be viewed both as a physical object and as a software object. In addition, this mathematical network approach is naturally suitable for parallel computing as it exploits the parallelism in physical systems. One can handle issues like data partition, assignment, and load balancing on the physics level using the structure of a given physical system. Synchronization and communication are controlled by the mathematical network specifications and are restricted to interfaces of subdomains, which results in a coarse-grained computational problem.

4. AGENT BASED COMPUTING PARADIGM FOR MPSES

We envisage an MPSE as a mathematical network whose nodes represent the physical components in a system or artifact. Each node has a mathematical model of the physics of the component it represents and a solver agent for its analysis. We propose to use the *multi-agent computing framework* to provide run-time support for MPSEs where we replace the multiphysics problem by a set of simple(r) simulation problems on simple geometries which must be solved simultaneously while satisfying a set of interface conditions. These simpler problems may reflect the underlying structure/geometry/physics of the system to be simulated, or may be artificially

created by techniques such as domain decomposition. Given a collection of *solver agents* for these smaller problems on simple geometries, we create a network of collaborating solvers by introducing *mediator agents* between them. Each solver deals with one of the subproblems defined earlier. The original multiphysics problem is solved when one has all the equations satisfied on the individual components and these solutions “match properly” on the interfaces between the components. This latter part is the responsibility of the mediator agents which facilitate the collaboration between solver agent pairs. The term “match properly” is defined by the physics if the interface is where the physics changes. For heat flow, for example, this means that temperature is the same on both sides of the interface and that the amount of heat flowing into one component is the same as the amount flowing out of the other. If the interface is artificial (introduced to make the geometry simple or the work smaller) then “match properly” is defined mathematically and means that the solutions join smoothly (have continuous values and derivatives).

Many agent-based systems have been developed[12; 42; 43; 46; 55] which demonstrate the power of the agent-oriented paradigm. It provides modularity and flexibility, so it is easy to dynamically add or remove agents, to move agents around the computing network, and to organize the user interface. An agent based architecture provides a natural method of decomposing large tasks into self-contained modules, or conversely, of building a system to solve complex problems by a collection of agents, each of which is responsible for small part of the task. Agent-based systems can minimize centralized control.

The agent-based paradigm is useful in scientific computing to handle complex mathematical models in a natural and direct way. It allows *distributed problem solving* [32] which is distinct from merely using distributed computing. The expected behavior of the simple problem solvers, computing locally and interacting with the neighboring solvers, naturally take on the behavior of a *local problem solver* agent. The task of mediating interface conditions between adjacent subproblems is given to *mediator* agents and their ability to autonomously pursue their goals can resolve the problems during the solution process without user intervention and converge to the global solution.

Several researchers have addressed the issue of coordinating multi-agent systems. For instance Smith and Davis [44] propose two forms of multi-agent cooperation, task sharing and result sharing. Task sharing essentially involves creating subtasks, and then farming them off to other agents. Result sharing is more data directed. Different agents are solving different tasks, and keep on exchanging partial results to cooperate. They also proposed using “contract nets” to distribute tasks. Wesson et al., showed[54] how many intelligent sensor devices could pool their knowledge to obtain an accurate overall assessment of a situation. The specific task presented in their work involves detecting moving entities, where each “sensor agent” sees only a part of the environment. They reported results using both an hierarchical organization, as well as an “anarchic committee” organization, and found that the latter was as good as, and sometimes better than the former. Cammarata et al. [4] present strategies for cooperation by groups of agents involved in distributed problem solving, and infer a set of requirements on information distribution and organizational policies. They point out that different agents may have different capabilities, limited knowledge and resources, and thus differing appropriateness in

solving the problem at hand. Lesser et al. [26] describes the FA/C (functionally accurate, cooperative) architecture in which agents exchange partial and tentative results in order to converge to a solution. Joshi [16] presents a learning technique which enhances the effectiveness of such coordination. It combines neuro-fuzzy techniques [45] with an epistemic utility criterion.

5. THE RESOURCE SELECTION PARADIGM FOR MPSES

In this paradigm for a networked MPSE environment, the solver and mediator agents form a potentially large pool of computational objects spread across the Net. Moreover, there are many possible choices for their instantiation. For example the past few decades has seen a huge amount of sophisticated code being developed to solve specific, homogeneous problems. Mediators today are almost nonexistent and a large number will have to be created to allow disparate solvers to interact. Clearly, expecting the user to be aware of all the potentially useful solvers on the Net is not realistic. Nor is a user likely to know all the hardware choices available to solve the problem. This problem is an obvious generalization of the *algorithm selection* problem formulated by Rice [40], we call it the *resource selection* problem in the context of MPSEs. We propose the use of *advisory agents* that accept a problem definition and some performance/success criteria from the user, and that then suggest software components and hardware resources that can be deployed to solve this problem. This is very similar to the idea of *recommender systems* that is being proposed for harnessing distributed information resources. While the recommender problem has been identified for networked information resources and initial research done [38], the resource selection problem remains largely ignored for harnessing networked computational resources. Note that the problem is different from invoking a known method remotely on some object, a problem where many distributed object oriented techniques are being developed and proposed. To appreciate the need for advisory agents, consider the present day approximation to “networked” scientific computing. Several software libraries for scientific computing are available, such as Netlib, Lapack/ScaLapack, etc. There are even some attempts to make such systems accessible over the web, such as Web //ELLPACK [from Purdue, <http://pellpack.cs.purdue.edu/>] and NetSolve [from UTK/ORNL, <http://www.cs.utk.edu/netsolve/>]. The GAMS [2] system helps users to identify and locate the right class of software for their problem. However, the user has to select the specific routine most appropriate for the given problem, download the software along with its installation and use instructions, install the software, compile (and possibly port) it, and then learn how to invoke it appropriately. Clearly this is a non-trivial task even for a single piece of software, and it can be enormously complex when multiple software components need to be used. Using networked resources today can be viewed as the modern day equivalent of programming ENIAC, which required direct manipulation of connecting wires. Systems are needed to abstract away the detail of the underlying networked system from the user and allow interaction with this system in the application domain. This is where MPSEs with inherent “intelligence” come in. We posit that multiagent systems, consisting of broad class of solver, mediator, and advisory agents can be used to create MPSEs with the desired characteristics.

As mentioned earlier, our prototype software to validate these ideas is being cre-

ated for PDE based systems. The numerical solution of PDEs depends on many factors including the nature of the operator, the mathematical behavior of its coefficients and its exact solution, the type of boundary and initial conditions, and the geometry of the space domains of definition. Most numerical solvers for PDEs normally require a number of parameters (mesh sizes, iteration parameters, sparse matrix representations) from the user in order to obtain a solution within a specified error level while satisfying certain resource (e.g., memory and time) constraints. The problem of selecting a solver and its parameters for a given PDE problem to satisfy the user's computational objectives is difficult and of great importance. The user must also select a machine from among the many available on the network, including parallel machines. Depending on the mathematical characteristics of the PDEs, there are "thousands" of numerical methods to apply, since very often there are several choices of parameters or methods at each of the several phases of the solution. It is unrealistic to expect that engineers and scientists will or should have the deep expertise to make "intelligent" combinations of selections of methods, their parameters, and computational resources that will satisfy their objectives.

The PYTHIA [53] project at Purdue has focussed on creating a knowledge based system that selects scientific algorithms to achieve desired tasks in computing. It determines a near-optimal strategy (i.e., a solution method and its parameters) for solving a given problem within user specified resource (i.e., limits on execution time and memory usage) and accuracy requirements (i.e., level of error). While the ideas behind PYTHIA are quite general, our current implementations operate in conjunction with systems that solve (elliptic) partial differential equations (PDEs), such as the ELLPACK and //ELLPACK PSEs developed at Purdue. The methodology of PYTHIA is to gather performance information about PDE solvers on standardized test problems and use this data plus feature information about PDE problems to determine good algorithms to solve the PDEs. The efficacy of this approach is dependent on the breadth and diversity of the method and problem sets used to create the performance evaluation information.

We now briefly describe some attempts at developing intelligent systems for assisting in various aspects of the PDE solution process. In [41], Rice describes an abstract model for the algorithm selection problem, which is the problem of determining a selection (or mapping) from the problem feature space to the algorithm space. Using this abstract model Rice describes an experimental methodology for applying this abstract model in the performance evaluation of numerical software. In [30], Moore et al. describe a strategy for the automatic solution of PDEs at a different level. They are concerned with the problem of determining (automatically) a geometry discretization that leads to a solution guaranteed to be within a prescribed accuracy. In [6; 7], Dyksen and Gritter describe a rule based expert system for selecting solution methods for elliptic PDE problems based on problem characteristics. This work differs significantly from our approach, which uses only performance data as the basis of the algorithm selection methodology. While these rules help some, we argue that using problem characteristics solely is not sufficient because the performance of a solver depends on quantities which cannot be measured symbolically and a priori. Further, software performance depends not only on the algorithms used, but on their implementations as well. In [23], Kamel et al. describe the expert system ODEXPERT for selecting numerical solvers for ini-

tial value ordinary differential equation (ODE) systems. ODEXPERT uses textual parsing to determine some properties of the ODEs and performs some automatic tests (e.g., a stiffness test) to determine others. Once all the properties are known, it uses its knowledge base information about available ODE solution methods (represented as a set of rules) to recommend a certain method. After a method has been determined, it selects a particular implementation of that method based on other criteria and then generates source code (Fortran) for the user. If necessary, symbolic differentiation is used to generate code for the Jacobian as well. Leake has recently begun some work in the area of using traditional case based reasoning systems to select appropriate methods for solving sparse linear systems [24]. Our group has also been actively involved in using several techniques, such as neural nets, neuro-fuzzy systems and Bayesian nets ([53; 22; 35; 21; 19; 20]) to address related issues of classifying PDE problems based on their performance characteristics, and then using this classification to predict an appropriate solution method for new problems. We have also formulated the algorithm selection problem as conducting knowledge discovery in domains of computational science [36; 37]. This work shows that such data mining approaches can be used to form relational descriptions of PDE objects which lead to more powerful schemas for resource selection (in terms of both representation and prediction).

6. SCIAGENTS SYSTEM

In this section, we describe in detail the *SciAgents* software architecture, and explain how to use it for complex PDE-based models from MPSEs. As an application of our MPSE approach, SciAgents employs two major types of computing agents – *solvers* and *mediators*. It interacts with the *recommender* agents as described later. Each solver agent is considered a ‘black box’ by the other agents and it interacts with them using an interagent language for the specific problem. This feature allows all computational decisions for solving one individual subproblem to be taken independently from the decisions in any other subproblem – a major difference from the traditional approaches to multidisciplinary simulations. Each mediator agent is responsible for adjusting an interface between two neighboring subproblems. Since the interface between any two subproblems might be complex in itself, there may be more than one mediator assigned to adjust it, each of them operating on separate piece of the whole interface. Thus the mediators control the data exchange between the solvers working on neighboring subproblems by applying mediating formulas and algorithms to the data coming from and going to the solvers. Different mediators may apply different mediating formulas and algorithms depending on the physical nature of their interfaces. The mediators are also responsible for enforcing global solution strategies and for recognizing (locally) that some goal (like “end of computations”) has been achieved.

The solvers and mediators form a network of agents to solve the given global problem. A schematic view of the functional architecture of a SciAgents MPSE containing an example network is given in Figure 2. The computations (and the major data exchange) are concentrated in the network of solver (PSE) and mediator agents. The solver agents communicate with the recommender agents (as consultants) through queries to obtain “advice” on computation parameters. The user interacts with the system through the global and local user interfaces which

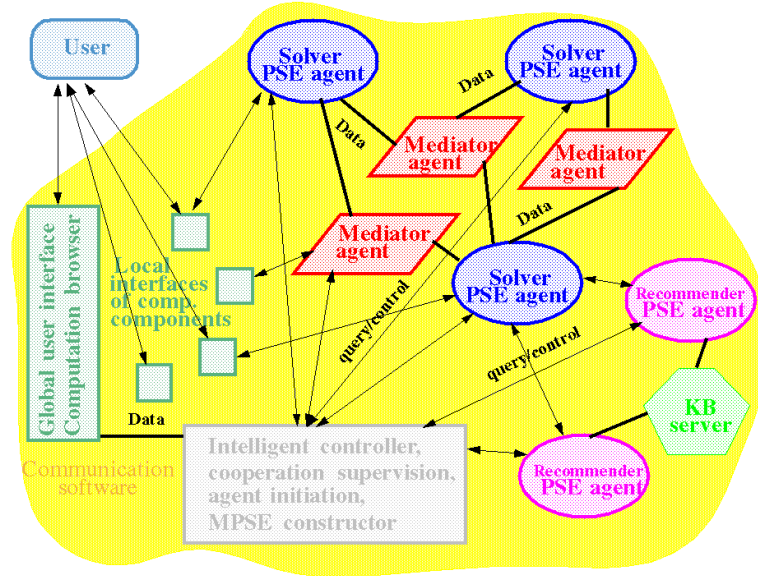


Fig. 2: Functional architecture of a SciAgents solver for an MPSE. The computations (and the major data exchange) are concentrated in the network of solver (PSE) and mediator agents. The solver agents communicate with the recommender ones through queries to obtain “advice” on computation parameters. The user interacts with the system through the global and local user interfaces which send queries and receive replies from the various agents.

send queries and receive replies from the various agents. The intelligent controller and the MPSE constructor can be integrated into a single “agent” which controls the global state of the computations and instantiates, queries, and manages (if necessary) the other agents.

We now describe how the user builds (“programs”) this network. The agent framework provides a natural abstraction to the user in the problem domain and hides the details of the actual algorithms and software involved in the problem solving. The user firsts breaks down the geometry of the composite domain into simple subdomains with simple models to define the subproblems for each subdomain. Then the physical conditions along each interface between the subdomains are identified. All this is done in the terms of the user’s problem domain. The user is provided with an *MPSE constructor (agent instantiator)* — a process which displays information about the templates and creates active agents of both kinds, capable of computing. Initially, only *templates of agents* — structures that contain information about solver and mediator agents and how to *instantiate* them, are available. Then the user constructs the proper network of computing agents

by simply *instantiating* various agents. The user selects solvers that are capable of solving the corresponding subproblems and mediators that are capable of mediating the physical conditions along the specific interfaces, and assigns subproblems and interfaces, respectively, to each of them. The user interacts with the system using a visual programming approach which has proved useful in allowing the non-experts to “program” by manipulating images and objects from their problem domain. In our case, a visual environment is useful for the MPSE constructor, or when the user wants to request some action or data.

Once an agent is instantiated, it takes over the communication with the user and with its environment (the other agents) and tries to acquire all necessary information for its task. Each PSE (solver agent) retains its own interface and can interact with the user. It is convenient to think of the user as another agent in these interactions. The user defines each subproblem independently, interacting with the corresponding solver agent through its user interface and similarly interacting with the mediators to specify the physical conditions holding along the various interfaces.

The agents *actively* exchange partial solutions and data with other agents without outside control and management. In other words, each solver agent can request the necessary domain and problem related data from the user and decide what to do with it (should it, for instance, start the computations or should it wait for other agents to contact it?). After each mediator agent has been supplied with the connectivity and mediating data by the user, it contacts the corresponding solver agents and requests the information it needs. This information includes the geometry of the interface, the functional capabilities of the solvers with respect to providing the necessary data for adjusting the interface, visualization capabilities, etc. All this is done without user involvement. By instantiating the individual agents (concentrating on the individual subdomains and interfaces) the user builds the highly interconnected and interoperable network that is tailored to solve the particular multiphysics problem, by *cooperation* between individual agents.

The user’s high-level view of the MPSE architecture is shown in Figure 3. The global communication medium used by all entities in the MPSE is called a *software bus* [50]. The MPSE constructor communicates with the user through the user interface builder and uses the software bus to communicate with the templates in order to instantiate various agents. Agents communicate with each other through the software bus and have their own local user interfaces to interact with the user. The order of instantiating the agents is not important. If a solver agent is instantiated and it does not have all the data it needs to compute a local solution (i.e., a mediator agent is missing), then it suspends the computations and waits for some relaxer agent to contact it and to provide the missing values (this is also a way to “naturally” control the solution process). If a mediator agent is instantiated and a solver agent on either side of its interface is missing, then it suspends its computations and waits for the solver agents with the necessary characteristics (the right subdomain assigned) to appear. This built in synchronization is, we believe, an important advantage of the SciAgents architecture. It results from each agent adapting to its environment. We go into more detail about inter agent communication later.

Since agent instantiation happens one agent at a time, the data which the user has to provide (domain, interface, problem definition, etc.) is strictly local, and

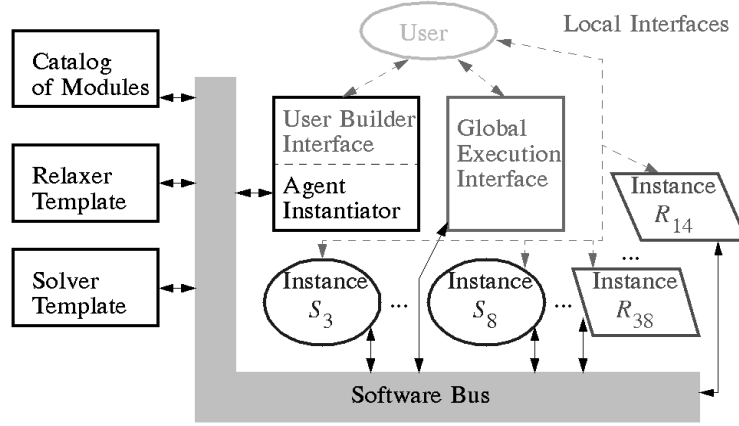


Fig. 3: Software architecture of an MPSE: the user's abstraction. The user initially interacts with the User Interface Builder to define the global composite problem. Later the interaction is with the Global Execution Interface to monitor and control the solution of the problem. Direct interaction with individual solvers and mediators is also possible. The agents communicate with each other using the *software bus*.

the agents collaborate in building the computing network. The user actually does not even need to know the global model. We can easily imagine a situation when the global problem is very large. Different specialists may only model parts of it. In such a situation, a user may instantiate a few agents and leave the instantiating of the rest of the cooperating agents to colleagues. Naturally, some care has to be taken in order to instantiate all necessary agents for the global solution and not to define contradictory interface conditions or mediation schemes along the “borders” between different users.

The collection of agent interfaces that a user interacts with is the only software the user actually needs to run locally in order to solve a problem. Therefore, this architecture abstracts successfully from the user the location of the main computations (the location of the solvers and the mediators) and allows for great flexibility in this direction, including running the MPSE over the Internet and distributing the agents over the Net.

This user view of the SciAgents architecture is too abstract for an actual implementation where one has to design the internal architecture of each agent and the detailed communication among the agents. We refer the reader to [5] for these important details. We only mention here that the agent architecture utilizes the locality of the communication patterns described before and the fact that whenever a mediator is active (computing), the corresponding solvers are idle and vice versa. Also, the asynchronicity of the communication and the need of implementing the “pro-active” feature of the agents prompt us to employ many active threads in a single agent (multithreading).

Coordination of the Solution Process We discuss now some important aspects of the cooperation between the agents during the solution process. There are well-defined global mathematical conditions for terminating the computations, for example, reaching a specified accuracy, or impossibility to achieve convergence. In

most cases, these global conditions can be “localized” either explicitly or implicitly. For instance, the user may require different accuracy for different subdomains and the computations may be suspended locally if local convergence is achieved. Note that local convergence can be achieved and then later lost due to changes from other agents.

The local computations are governed by the mediators (the solvers simply solve the PDE problems). The mediator agents collect the errors after each iteration and, when the desired accuracy is obtained, *locally* suspend the computations and report the fact to the intelligent controller. The suspension is done by issuing an instruction to the solvers on both sides of this interface to use the boundary conditions for the interface from the previous iteration in any successive iterations they may perform (the other interfaces of the two subdomains might still not have converged). The solvers continue to report the required data to the submediators and the submediators continue to check whether the local interface conditions are satisfied with the required accuracy. If a solver receives instructions to use the old iteration boundary conditions for all its interfaces, then it stops the iterations. The iterations may be restarted if the interface conditions handled by a given mediator agent are no longer accurately satisfied (even though they once were). In this case, the mediator issues instructions to the two solvers on both sides of its interface to resume solving with new boundary conditions. If the maximum number of iterations is reached, the mediator reports failure to the intelligent controller and suspends the computations. The only global control exercised by the intelligent controller is to terminate all agents in case all mediators report local convergence or one of them reports a failure. The messages used in the interagent communication are given in full detail in [18], we provide a small example in the next section.

The above scheme provides a robust mechanism for cooperation among the computing agents. Using *only* local knowledge, they perform only local computations and communicate only with “neighboring” agents. They *cooperate* in solving a global, complex problem, and none of them exercises centralized control over the computations. The global solution “emerges” in a well-defined mathematical way from the local computations as a result of intelligent decision making done locally and independently by the mediator agents. The agents may change their goals dynamically according to the local status of the solution process – switching between observing results and computing new data.

Other global control policies can be imposed by the user if desired – the system architecture allows this to be done easily by distributing the control policy to all agents involved. Such global policies include continuing the iterations until the all interface conditions are satisfied, and recomputing the solutions for all subdomains if the user changes something (conditions, method, etc.) for any domain.

Software Reuse and Evolution One of the major goals of this MPSE approach is to design a system that allows for low-cost and less time-consuming methods of building the software to simulate a complex mathematical model of physical processes. This goal cannot be accomplished if the existing rich variety of problem solving software for scientific computing is not used. More precisely, there are many well-tested, powerful, and popular PSEs for solving problems very similar or identical to the subproblems that appear when breaking the global model into “simple” subproblems defined on a single subdomain. These PSEs could easily and

accurately solve such a “simple” subproblem. It is, therefore, natural to reuse such PSEs as solver agents. However, our architecture requires the solvers to behave like agents (e.g., understand agent languages, use them to communicate data to other agents), something the existing PSEs in scientific computing do not do.

Our solution to this problem is to provide an *agent wrapper* for PSEs and other software modules, which takes care of the interaction with the other agents and with the other aspects of emulating agent behavior. The wrapper encapsulates the original PSE and is responsible for running it and for the necessary interpretation of parameters and results. This is not simply a “preprocessor” that prepares the PSE’s input and a “postprocessor” that interprets the results, since the mediation between subproblems may require communicating intermediate results to the mediators and/or accepting some additional data from them. Designing the wrapper is sometimes complicated by the “closed” nature of extant PSEs — their original design is not flexible or “open” enough to allow access to various parts of the code and the processed data. However, it is our opinion that the PSE developers can design and build such a wrapper for a very small fraction of the time and the cost of designing and building entire new PSE or custom software for every new problem. The wrapper, once written, will enable the reuse of this PSE as a solver agent in different MPSEs, thus amortizing the cost further. As part of the specifications of the wrapper the developers have to consider the mediation schemes involving sub-models within the power of the PSE. An additional task is to evaluate the PSE’s user interface — since the user defines the local submodel through it, it is important that the interface facilitates the problem definition in user’s terms well enough. Our experience with //ELLPACK was that building a wrapper for a substantial (more than a million lines of code), diverse, and non-homogeneous PDE solver could be done efficiently, it required about a thousand lines of code.

7. PYTHIA SYSTEM

We see that the role played by the recommender agents is paramount for the effectiveness of *SciAgents*. When queried by the solver agents, they provide consulting advice on a suitable scheme (and associated computation parameters) to solve a given problem so as to achieve desired performance criteria. An example PDE problem is given in Figure 4. A prescribed solution strategy could be “*Use the 5-point star algorithm with a 200×200 grid on an $nCube/2$ with 16 processors. Confidence: 0.90*” (Notice that a recommender agent provides a level of confidence in the selected strategy). In essence, the recommender agents serve as knowledge engines that provide domain-specific inference for PDE problems. If any particular recommender agent lacks the expertise to provide this recommendation, it will collaborate with other recommender agents and select the best answer. These agents can also be made to interact directly with the user, via the agent instantiator. Thus PYTHIA is a collaborative, multi-agent system [33] that uses collective knowledge to prescribe a strategy to solve a given problem in scientific computation. The agents themselves are referred to as PYTHIA agents and are implemented by a combination of C language routines, shell scripts and systems such as CLIPS (the C Language Integrated Production System) [10]. The agents communicate using the Knowledge Query and Manipulation Language (KQML) [8], using protocol defined performatives. All PYTHIA agents understand and utilize a private language

PROBLEM #28	$(w u_x)_x + (w u_y)_y = 1,$ where $w = \begin{cases} \alpha, & \text{if } 0 \leq x, y \leq 1 \\ 1, & \text{otherwise.} \end{cases}$
DOMAIN	$[-1, 1] \times [-1, 1]$
BC	$u = 0$
TRUE	unknown
OPERATOR	Self-adjoint, discontinuous coefficients
RIGHT SIDE	Constant
BOUNDARY CONDITIONS	Dirichlet, homogeneous
SOLUTION	Approximate solutions given for $\alpha = 1, 10, 100$. Strong wave fronts for $\alpha \gg 1$.
PARAMETER	α adjusts size of discontinuity in operator coefficients which introduces large, sharp jumps in solution.

Fig. 4. A problem from the PDE population.

(PYTHIA-Talk) that describes the meaning (content) of the KQML performatives. This design allows the seamless integration of the recommender agents into the MPSE architecture.

A PYTHIA agent relies heavily on the problem set used in its performance evaluation knowledge base so the effectiveness of a recommender agent depends on its ‘experience’. For example, one agent’s expertise might come from its test base of computational fluid dynamics PDE solvers and problems while a second agent’s expertise might be based on heat conduction problems. Our multi-agent methodology recognizes that there are many, many different kinds of PDE problems and any single recommender agent is likely to be limited by its knowledge base. Thus, the approach taken is to create several different PYTHIA agents, each of which has information about some class(es) of PDE problems and can predict an appropriate solver for a given PDE of those classes. If a PYTHIA agent discovers that it does not have enough confidence in the prediction it is making, it could query all other PYTHIA agents, obtain answers from all of them and use this information to decide which one is “most reasonable”. This could entail a huge amount of network traffic and inordinate delays. A better approach is to use the information obtained by the initial broadcast type of queries to infer the most experienced PYTHIA agent for the problem at hand. This naturally raises the following issues:

- (1) Given more than one applicable agent, how does one determine the best agent(s) for a given PDE problem? In other words, what is the mapping from a given problem to the best PYTHIA agent?
- (2) Can the notion of best agent be inferred automatically or does it require user input?
- (3) How does one learn and adapt to the changing dynamics of the scenario? Agents may come into existence, some may go extinct, their knowledge corpus may change dynamically etc. How do we learn the mapping in this case and update it suitably?

We use a quantitative measure of reasonableness [16; 34], to automatically generate exemplars to learn the mapping from PDE problems to PYTHIA agents. This is needed because the computational scientist cannot be expected to have such information in this dynamic scenario. For example, in response to a query from

the user about a particular PDE problem, each PYTHIA agent might suggest a different method with varying levels of confidence in the recommended strategy. Moreover, each of these agents might have different levels of expertise (such as the kind of PDEs it knows about) and different ‘training’ history. The user, thus, cannot be expected to know which one of them is most suitable for the problem if all these responses are supplied. Our measure of reasonableness allows the automatic ‘ranking’ of the PYTHIA agents for a particular problem (class). This measure combines two factors, the probability of an agent’s prediction q being true, and the predictor’s utility. Specifically, the reasonableness of a proposition is defined as follows [25]:

$$r(q) = p(q)U_t(q) + p(\sim q)U_f(q),$$

where $U_t(q)$ denotes the positive utility of accepting q if it is true, $U_f(q)$ denotes the negative utility of accepting q if it is false and $p(q)$ be the probability that q is true.

In the case of PYTHIA, each agent produces a number denoting confidence in its recommendation being correct, so $p(q)$ is trivially available, and $p(\sim q)$ is simply $1 - p(q)$. For the utility, we use the following definition:

$$U_t(q) = -U_f(q) = f(N_e),$$

where f is some squashing function mapping the domain of $(0, \infty)$ to a range of $(0, 1]$, and N_e is the number of exemplars of a given type (that of the problem being considered) that the agent has seen. We chose $f(x) = \frac{2}{1+e^{-x}} - 1$ because it reflects the number ($x = N_e$) of problems of the present type that it has seen.

Having defined our notion of reasonableness, we still need a way to learn a mapping from a PDE problem to the most reasonable PYTHIA agent. We have evaluated standard statistical methods, gradient descent methods, machine learning techniques and other classes of algorithms [22], but it has been our experience that specialized techniques developed for this domain perform better than conventional off-the-shelf approaches [20]. In particular, we have designed a neuro-fuzzy technique that infers efficient mappings, caters to mutually non-exclusive classes (as the PDE problem classes naturally are) and learns the classifications in an on-line manner [22]. For the purposes of this paper, it is sufficient to understand that this scheme provides a mapping from a PDE problem to the best available recommender agent and that the mappings can be learnt in an incremental fashion using this reasonableness measure. While this mapping could be done by any of the PYTHIA agents by ‘housing’ a copy of the learned classification in each of them, we chose to create a central agent, PYTHIA-C whose main task is to perform this mapping. This just serves to demonstrate the learning aspect of the agents as distinct from their other capabilities.

8. CASE STUDIES

8.1 Solving Composite PDE Problems

The main issue is what mediation schemes can be applied to a composite PDE problem — in other words, how to obtain a global solution out of the local solutions produced by the single-domain solvers. To do this, SciAgents uses interface relaxation [5; 28]. Important mathematical questions of the convergence of the method,

the behavior of the solution in special cases, etc., are addressed in [31]. Typically, for second order PDEs, there are two physical or mathematical interface conditions involving values and normal derivatives of the solutions on the neighboring subdomains. The interface relaxation technique is as follows.

- . *Step 1.* Make initial guesses as boundary conditions to determine the subproblem solutions.
- . *Step 2.* Solve the subproblem in each subdomain and obtain a local solution.
- . *Step 3.* Use the solution values on the interfaces to evaluate how well the interface conditions are satisfied. Use a *relaxation formula* to compute new values of the boundary conditions.
- . *Step 4.* Iterate steps 2 and 3 until convergence.

We now describe the solution of a composite PDE problem using four solvers and five mediators. It models the heat distribution in the walls of a chemical or a nuclear reactor and in the surrounding isolating and cooling structures, see Figure 5. The subdomains are shown, with the solver agents S_i , $i = 0, \dots, 3$ simulating the local process in each subdomain and the mediators M_j , $j = 0, \dots, 4$ mediating the interface piece they are written on. The unknown function is T and the exterior boundary conditions are shown next to the corresponding boundary pieces. The reactor keeps the inside temperature of its wall at 1000 degrees and the outside walls of the cooling structures are kept at, more or less, room temperature. The boundary conditions along the x and y axes reflect the symmetry of the construction. We denote by Γ_{ik} the k -th boundary piece of the i -th subdomain. The differential operators L_i , $i = 1, 2, 3$ are

$$\begin{aligned} L_1 : T_{xx} + T_{yy} + \alpha_1 T &= \beta_2(x^2 + y^2 - 2) \\ L_2 : T_{xx} + T_{yy} + \alpha_2 T &= 0 \\ L_3 : T_{xx} + T_{yy} - \gamma_3(T_x + T_y) + \alpha_3 T &= 0 \end{aligned} \tag{1}$$

The parameters are: $\alpha_1 = 0.2$, $\alpha_2 = 0.4$, $\alpha_3 = 0.3$, $\beta_2 = -60$, $\gamma_3 = 10$. We denote by Ω_i the subdomain associated with S_i , $i = 0, \dots, 3$. We use as interface conditions the continuity of temperature and heat flow across the subdomain interfaces. Note that even though the interface between Ω_0 and Ω_1 , Ω_2 , and Ω_3 looks like a single curve from the point of view of Ω_0 , it is divided into three pieces Γ_{02} , Γ_{03} and Γ_{04} , so that the mediators M_0 , M_1 , and M_2 can each be assigned a single piece to mediate. The time we spent from writing down the problem on paper to getting a contour plot of the solution on the screen was 5 hours (this includes some manual calculations and adjusting the relaxation formulas for better convergence).

A user begins solving this problem by drawing Figure 5. The sketch identifies the subdomains (the solvers), the mediators, each boundary piece in every subdomain, and the endpoints of the interfaces. The sketch is necessary since the currently implemented version of *SciAgents* requires input as a script file. However, we believe that (with the possible exception of the boundary piece identifiers) such a sketch will be necessary even with the best imaginable graphical user interface. We only expect the user to annotate this initial sketch.




Fig. 5: A sketch of a composite PDE problem modeling the heat distribution in the walls of a chemical or a nuclear reactor and in the surrounding isolating and cooling structures. The subdomains are shown, with the solver agents S_i , $i = 0, \dots, 3$ simulating the local process in each subdomain and the mediators M_j , $j = 0, \dots, 4$ mediating the interface piece they are written on. The unknown function is T and the exterior boundary conditions are shown next to the corresponding boundaries. We denote by Γ_{ik} the k -th boundary piece of the i -th subdomain.

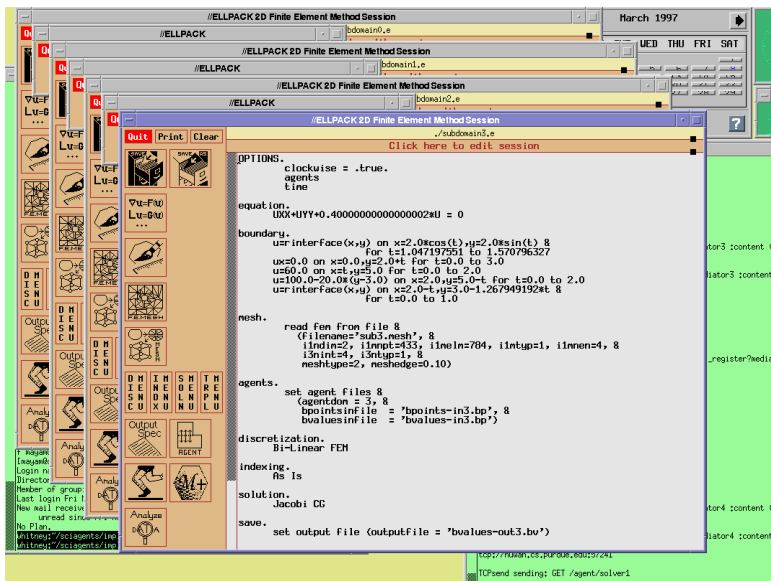


Fig. 6: Four copies of the `//ELLPACK` interface are presented to the user for defining the four PDE subproblems.

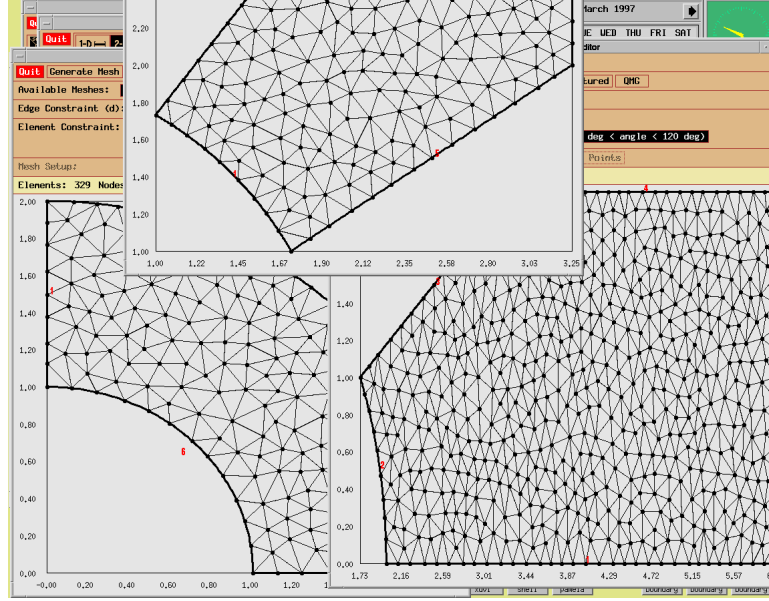


Fig. 7: A snapshot of the display during the subproblem definition process. Parts of three //ELLPACK domain tools containing three of the subdomain geometries and finite element meshes are visible. The user can discretize each subdomain completely independently from the others. For example, the densities of the above meshes are different.

After making the sketch the user constructs the *SciAgents* input file and starts *SciAgents*. This starts the global controller (containing the agent instantiator) and it instantiates the agents on the appropriate machines and builds the network of four solvers and five mediators that is to solve the problem. After that, the “computing” thread of the global controller starts a shell-like interface with two major commands: **pause** and **tolerance** for control and steering the computations. The **pause** prompts the controller to issue messages to all agents to save their current state and to exit. The **tolerance** command changes dynamically the tolerance of a given mediator or of all mediators.

After the initial exchange of data to check that all agents are ready, the user sees four copies of the //ELLPACK user interface (see Figure 6). All four subproblems are defined (see Figure 7 for a snapshot during this process) and selecting a discretizer, linear solver, etc., in one subdomain does not lead to any requirement or necessity about selections in the neighboring subdomains. If a subdomain is huge, one may choose to use a 32-node Intel Paragon for it, while the neighboring tiny subdomain may be simulated on the same host where the wrapper is running. There are only two requirements for global synchronization of the local definitions: each subdomain geometry has to be input in terms of the global coordinate system (hence the need of the coordinates of the boundary pieces in the sketch), and for each interface piece, the right-hand side of the boundary conditions has to be the function **rinterface(x,y)**. It is the user’s responsibility to make sure that the relaxation formulas used for each interface piece correspond to the left-hand sides

of the boundary conditions entered in the two solver's user interfaces. For the example, the boundary condition used at all interfaces is $\mathbf{T} = \mathbf{rinterface}(\mathbf{x}, \mathbf{y})$ and the relaxation formula is (U is the solution on the "left" side, V is the solution on the "right" side; U_n is the normal derivative; f is a factor given below; the formula is always applied pointwise for each point from any solver's grid/mesh on the interface):

$$U^{new} = V^{new} = \frac{1}{2}(U^{old} + V^{old}) - f \times (U_n^{old} - V_n^{old}) \quad (2)$$

The form of the factor f is $f = \frac{|U_n^{old}| + |V_n^{old}|}{(|U_n^{old}| + |V_n^{old}|)f_0}$ which scales the relaxation properly (and avoids dependencies on the choice of the coordinate system) and regulates the rate of change of the boundary conditions along the interface from iteration to iteration by changing f_0 . It is sometimes hard to predict the "optimal", or even the acceptable, values of f_0 .

The user input results in writing the script for the actual future runs. The user exits the //ELLPACK interface which prompts the wrapper to collect the initial data and to send them to the mediators. They compute initial right-hand sides of the boundary conditions. After the mediators provide all necessary boundary conditions, the wrapper runs the script which, in turn, runs the executable(s). When the iteration is completed the wrapper takes over again and extracts all required data from the computed solution and sends it to the mediators, waiting for the new boundary conditions from them. Thus, at the next iteration, no new compilation and user actions are necessary, since the same script (and executable(s)) is run by the wrapper.

For this example, we had to change the factor f_0 twice before the process began to converge, especially for mediators M_3 and M_4 . This seems to be due to the natural singularity that occurs at the reentrant corners of the global domain which affects the stability of the convergence.

When a mediator observes convergence (the change of the boundary conditions for the next iteration is smaller than the tolerance), it reports this to the global controller, and after all mediators report convergence, the global controller issues a message to all agents to stop. In this case we had convergence after 53 iterations. Figure 8 shows a combined picture of all four subdomain solutions. Note that all contour lines match when crossing from one subdomain to another, there are even a few which go through three subdomains, and one going through all four subdomains. This is solid evidence that the interface relaxation technique works in this problem.

To experiment with the applicability of *SciAgents* to more difficult problems we solved several variations of the above example replacing L_1 , L_2 , and L_3 with nonlinear operators (exhibiting different nonlinearity for different L_i). Since //ELLPACK uses a Newton iterative procedure to solve a nonlinear problem, the global solution process becomes a multi-level iteration where one *SciAgents* step involves a complete Newton iteration in it. Also, while one can plausibly handle the linear example above by considering a single PDE with discontinuous coefficients on a single domain, this approach is not feasible for nonlinear problems. Using *SciAgents* we were able to solve the following sets of PDEs (the increased complexity is




Fig. 8: A combined picture of all subdomain solutions of the example problem in Equation (1). The global solution corresponds to the physical intuition about the behavior of the modeled real-world system. All contour lines match when crossing from one subdomain to another, there are even a few which go through three subdomains and one going through all four subdomains.

reflected in a two to three-fold increase of the number of iterations necessary for the convergence to the global solution).

—Set 1

$$\begin{aligned} -L_1 : TT_{xx} + (1 + T)T_{yy} + aT(1 + T) &= b(x^2 + y^2 - 2) \\ -L_2 : T_{xx}/(1 + (x - y)^2) + T_{yy}/(1 + (4x - 5y)^2) + cT/(101 + t) &= 0 \\ -L_3 : T_{xx} + T_{yy} - d(T_x + T_y) + cT &= 0 \end{aligned}$$

—Set 2

$$\begin{aligned} -L_1 : T_{xx} + T_{yy} + ae^{x+y+T/500} &= b(x^2 + y^2 - 2) \\ -L_2 : TT_{xx} + TT_{yy} + (T_x + 20)T_y + 2(T_x - 20)T_x &= 0 \\ -L_3 : T_{xx} + T_{yy} - b(T_x + T_y) + aT &= 0 \end{aligned}$$

—Set 3

$$\begin{aligned} -L_1 : T_{xx} + T_{yy} + \alpha_1 T(1 + T/1000) &= \beta_2(x^2 + y^2 - 2) \\ -L_2 : T_{xx} + T_{yy} + \alpha_2 T &= 0 \\ -L_3 : T_{xx} + (1 + T/1000)T_{yy} + (T_x/500 + 3)T_x + \alpha_3 T &= 0 \end{aligned}$$

8.2 Intelligent PDE Computing with PYTHIA

In this section, we describe how PYTHIA can be used to determine reasonable strategies for PDE problem solving. In our prototype implementation, our PYTHIA agents' expertise stems from the following classes of PDEs (we also list the number of samples in each class from our study that involves about 167 PDE problems):

- (1) SINGULAR: PDE problems whose solutions have at least one singularity (6 exemplars).
- (2) ANALYTIC: PDE problems whose solutions are analytic (35 exemplars).
- (3) OSCILLATORY: PDE problems whose solutions oscillate (34 exemplars).
- (4) BOUNDARY-LAYER: Problems with a boundary layer in their solutions (32 exemplars).
- (5) BOUNDARY-CONDITIONS-MIXED: Problems that have mixed boundary conditions in their solutions (74 exemplars).
- (6) SPECIAL: Problems that do not belong to the above classes (10 problems).

Note that these classes are not mutually-exclusive, so their total membership is 191 problems. In other words, there are different PYTHIA agents, each of which can recommend a solver for a PDE belonging to its representative class(es) of problems. Also, a problem can belong to more than one class simultaneously (a given PDE can *both be* analytic *and have* mixed boundary conditions). Detecting the presence of such mutually non-exclusive classes is critical to selecting a good solver for the PDE.

To test our ideas, we made five experiments, with 2, 3, 4, 5 and 6 PYTHIA agents respectively. In each experiment, each PYTHIA agent knows about a certain class(es) of PDE problems. For example, with 6 PYTHIA agents, each agent knows about one of the above classes of PDEs. In the '3-agent' experiment, agent 1 knows about problem classes 1 and 2, agent 2 knows about classes 3 and 4 and the third agent knows about classes 5 and 6. The population of 167 PDE problems was split into two parts: a large set of 111 problems and a smaller set of 56 problems. We conducted two sets of experiments: In each scenario, we first

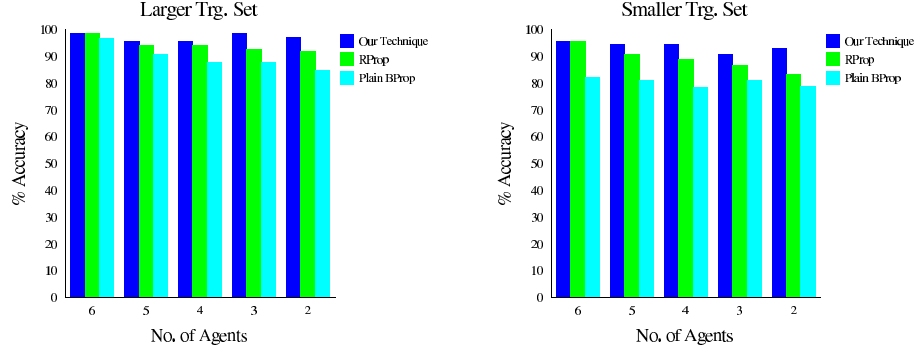


Fig. 9: Performance of learning algorithms. The graph on the left depicts the results with the larger training set and the one on the right shows the results with the smaller training set. In each case, recommendation accuracy figures for the 5 experiments (with 2, 3, 4, 5 and 6 agents) are presented for all the three learning algorithms considered in this paper.

trained our technique on the larger set of $\{\text{problem, agent}\}$ pairs (using the notion of reasonableness defined earlier) and tested our learning on the smaller set of 56 exemplars. In the second experiment, the roles of these two sets were reversed. We also compared our technique with two very popular gradient descent techniques for training feedforward neural networks, namely, Vanilla (Plain) Backpropagation (BProp) and Resilient Propagation (RProp). Figure 9 summarizes the results.

It can be easily seen that our method consistently outperforms BProp and RProp on learning the mapping from problems to agents. Also, performance on the larger training set was expectedly better than that on the smaller training set. Moreover, our algorithm operates in an on-line mode; new data do not require retraining on the old. Our technique was also tested for this ability; for the larger training set, we incrementally trained our algorithm on the 111 PDEs and the accuracy figures on the test set were found to rise steadily to the figures shown in Fig. 9. In the collaborative networked scenario of an MPSE, where the resources change dynamically, this feature of our neuro-fuzzy system enables us to automatically infer the capabilities of multiple PYTHIA agents. If the capabilities of agent 1 were to change, for example, in the 6-agent scenario, then our network could infer the new mappings without losing the information already learnt. This feature is absent in most other methods of classification such as BProp and RProp in which the dimensionality of the network is fixed and it is imperative that the old data be kept around if these networks are to update their learning with new data.

The PYTHIA project web pages at <http://www.cs.purdue.edu/research/cse/pythia> provide information about this collaborative PYTHIA methodology and facilities to invoke it remotely. At the outset, there is a facility to provide feature information about a PDE problem. In particular, there are forms that guide the user in providing information about the operator, function, domain geometry and boundary conditions. Once these details are given, the information is submitted to the central PYTHIA agent, PYTHIA-C, that performs further processing. As mentioned before, it first classifies the given PDE problem into categories of problems as described above. Having classified the problem into one or more of these classes,

the PDE is taken to an appropriate PYTHIA agent for this class of problems, which in turn predicts an optimal strategy and reports back to the user.

8.3 Learning and Adaptation in MultiAgent Systems

The above experiment can be visualized as an example where the central agent PYTHIA-C is in a *learning mode*, cycles through the training set, and learns mappings from the given PDEs to appropriate agents. From this point on PYTHIA-C is in the *stable mode*. It will only ask the best agent to answer a particular question. If PYTHIA-C finds a PYTHIA agent's recommendation unacceptable, it will ask the next best agent, until all agents are exhausted. This is facilitated by our neuro-fuzzy learning algorithm. By varying an acceptance threshold in the algorithm, we can get an enumeration of "not so good" agents for a problem type. If PYTHIA-C determines no plausible solution exists among its agents or itself, then PYTHIA-C gives the answer that "is best". When giving such an answer, the user is notified of PYTHIA-C's lack of confidence.

While this scheme serves most purposes, an issue still pending is the mode of switching between the learning and stable modes. PYTHIA-C switches from learning to stable mode after an *a priori* fixed number of problems (this was 111 in our first set of experiments, for example). The timing of the reverse switch back to learning is a more interesting problem; we report on three different methods.

Time based: This simple approach is where PYTHIA-C reverts to learning after a fixed time period. At such points, PYTHIA-C cycles through its training set, queries other agents, gets back answers, determines reasonableness values and finally learns new mappings for the PDE problems. Figure. 10 depicts the results with the six-agent case and the time based approach using the larger training set. Initially, each agent starts up with approximately 1/3 of their total knowledge base and this knowledge steadily increases with time. At periodic time intervals, PYTHIA-C switches to learning mode and cycles through the larger training set with each of the agents in the experiment. The performance is then measured with the smaller training set. As can be seen, the accuracy figure steadily improves for each of the six individual agents to the accuracy observed in the previous static experiment. PYTHIA-C's accuracy improves from 40.85% to 98.20% in this experiment.

We conducted another experiment with this method, one more realistic for multi-agent systems. We begin the experiment with no 'known' agents, i.e., PYTHIA-C initially does not know about the existence of any agents or their capabilities. Then, each agent is introduced into the experiment with a small initial knowledge base and then their knowledge base is slowly increased. For example, Agent 1 comes into the setup with a small knowledge base and announces its existence to PYTHIA-C which creates a class for Agent 1. It then reverts to learning mode (though wasteful) and learns mappings from PDE problems to agents (in this case, there is only one agent). After some time, Agent 3 comes into the experiment and this process is repeated. This is repeated until all six agents are introduced. While the addition of new agents and associated classes is taking place, the abilities of existing agents (like Agent 1) also increase simultaneously. Thus, these events happen in parallel; i.e., addition of new agents and additions to the knowledge base of existing agents. Because our neuro-fuzzy scheme has the ability to introduce new classes on the fly,

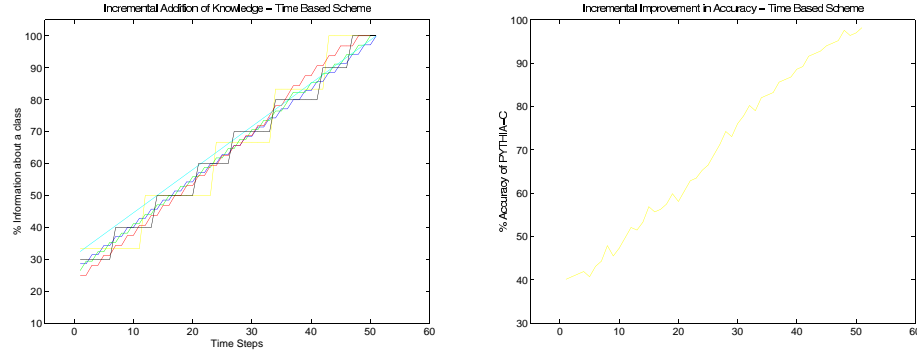


Fig. 10: Results with the time based scheme for 6 agents using the larger training set. The graphs on the left show the systematic increase in the abilities for each of the agents individually and the one on the right shows the corresponding improvement in accuracy of the central agent, PYTHIA-C

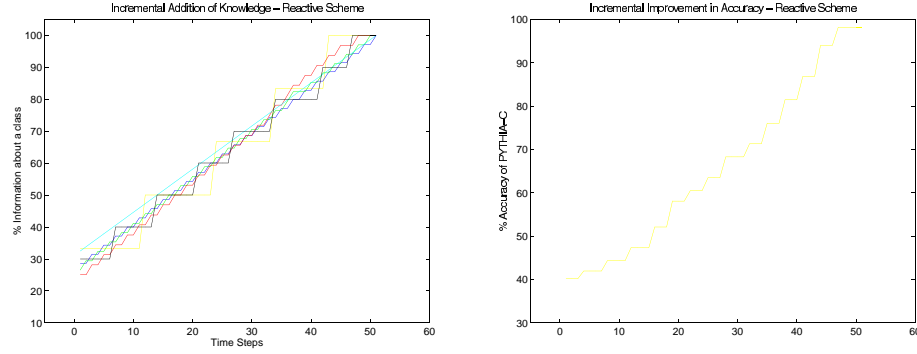


Fig. 11: Results with the reactive method for 6 agents using the larger training set. The graphs on the left show the systematic increase in the abilities for each of the agents and the one on the right shows the corresponding improvement in accuracy of PYTHIA-C

PYTHIA-C can handle this situation well. The accuracy figures converge to the values previously obtained.

Reactive: In this method, a PYTHIA agent notifies PYTHIA-C whenever its confidence for some class of problems has changed significantly. PYTHIA-C reverts to learning when it next receives a query about this class of problems. Each agent started with the same initial knowledge base as before and this is slowly increased. As the agents indicate the resulting increase in confidence to PYTHIA-C, it reverts to learning mode from time to time. The accuracy figures for PYTHIA-C approach the same values as before; they follow a monotonic pattern, but a more slowly increasing pattern, see Figure 11.

Time based reactive: This is a combination of the two methods above where PYTHIA-C sends out a “has anyone’s abilities changed significantly” message at fixed time intervals, and switches to learning if it receives a positive response. Each agent, starts with the same knowledge base and this is slowly increased. Figure 12 shows that the accuracy figures for PYTHIA-C are again a monotonic increasing

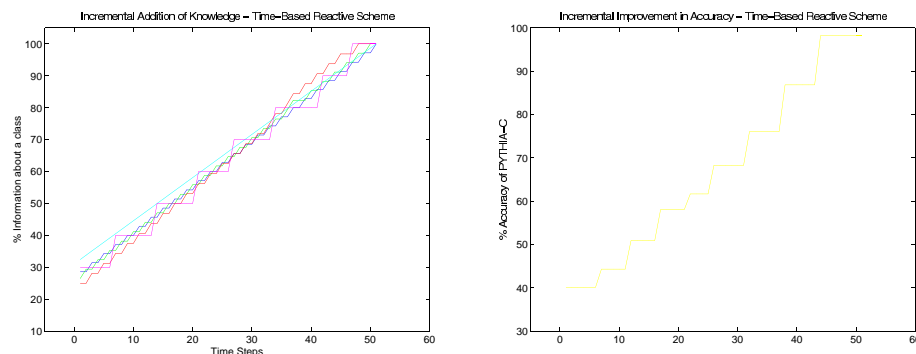


Fig. 12: Results with the time based reactive method for 6 agents using the larger training set. The graphs on the left show the systematic increase in the abilities for each of the PYTHIA agents and the one on the right shows the corresponding improvement in accuracy of PYTHIA-C.

and rising slightly faster than for the reactive method.

Our experiments with the three methods show that they enable the central agent PYTHIA-C to keep track of the dynamic capabilities (in our case, the knowledge base) of other agents. These methods also enable PYTHIA-C to handle situations where agents appear and disappear over time.

9. CONCLUSION

In this chapter, we have described an agent oriented architecture for MPSEs, consisting of solver, mediator and recommender agents. This architecture enables us to combine existing PSEs and libraries into MPSEs. The SciAgents [5] and PYTHIA [33] systems provide the solver, mediator and recommender agents required to realize multidisciplinary problem solving environments. Our ongoing research focusses on many more aspects of these problems. We are extending the functionality of the SciAgents system to address more complex problem domains and are also investigating strategies to choose interface relaxation schemes. In the PYTHIA system, we are working on enhancing the knowledge bases to provide more flexible resource selection schemes. In the PDE domain, for instance, PYTHIA can be used to select partitioning strategies (for parallel PDE solving), mesh refinement techniques and selecting solvers for linear systems that arise from the discretization of PDEs. We are also interfacing the PYTHIA system with the GAMS system for mathematical software to facilitate software delivery. Together, the systems presented here address rapid multidisciplinary prototyping — one of the most important grand challenge problems in computational science and engineering.

REFERENCES

- [1] K. C. Bernard. Ordering Chaos: Supercomputing at the edge. In D. Leebaert, editor, *Technology 2001: The Future of Computing and Communications*. MIT Press, Cambridge, MA, 1992.
- [2] R.F. Boisvert, S.E. Howe, and D.K. Kahaner. The Guide to Available Mathematical Software Problem Classification System. *Comm. Stat. - Simul. Comp.*, vol.20(4):pp.811–842, 1991.
- [3] R. E. Burkart. Reducing the R&D Cycle Time. Technical report, Research Tech. Mgmt., 1994.

- [4] S. Cammarata et al. Strategies of Cooperation in Distributed Problem Solving. In Bond and Gasser, editors, *Readings in Distributed Artificial Intelligence*, pages 102–105. Morgan Kaufmann, 1988.
- [5] T. T. Drashansky. *An Agent-Based Approach to Building Multidisciplinary Problem Solving Environments*. PhD thesis, Dept. Comp. Sci., Purdue University, December 1996.
- [6] Wayne R. Dyksen and Carl R. Gritter. Elliptic Expert: An Expert System for Elliptic Partial Differential Equations. *Mathematics and Computers in Simulation*, 31:333–343, 1989.
- [7] Wayne R. Dyksen and Carl R. Gritter. Scientific Computing and the Algorithm Selection Problem. In E. N. Houstis, J. R. Rice, and R. Vichnevetsky, editors, *Expert Systems for Scientific Computing*, pages 19–31. North-Holland, 1992.
- [8] Fritzson, R. KQML: A Language and Protocol for Knowledge and Information Exchange. In *Proc. 13th Intl. Distributed Artificial Intelligence Workshop*, pages 134–143. Springer-Verlag, 1994.
- [9] E. Gallopoulos, E. N. Houstis, and J. R. Rice. Computer as Thinker/Doer: Problem-Solving Environments for Computational Science. *IEEE Computational Science and Engineering*, 1(2):11–23, 1994.
- [10] J. C. Giarratano. *CLIPS User's Guide, Version 5.1*. NASA Lyndon B. Johnson Space Center, 1991.
- [11] S.S. Grimajl. The First ICASE/LARC Industry Roundtable: Session Proceedings. Technical Report ICASE Interim Report 26, ICASE, NASA Langley Research Center, Hampton, VA, 1995.
- [12] B. Hayes-Roth et al. Guardian. A Prototype Intelligent Agent for Intensive-care Monitoring. *Artif. Intell. Med*, 4(2):165–185, 1992.
- [13] C. M. Hoffmann, E. N. Houstis, J. R. Rice, A. C. Catlin, M. Gaitatzes, S. Weerawarana, N.-H. Wang, C. Takoudis, and D. Taylor. SoftLab - A Virtual Laboratory for Computational Science. *Math. Comp. in Simulation*, 36, 1994.
- [14] E.N. Houstis, A. Joshi, J.R. Rice, and S. Weerawarana. MPSEs: Multidisciplinary Problem Solving Environments. White Paper presented at the *America in the Age of Information: A forum*, Committee on Information and Communications, National Science and Technology Council, 1996.
- [15] Industrial Research Institute. Proceedings: Roundtable Meeting on Reducing R&D Cycle Time. Technical report, Industrial Research Inst., Washington DC, 1992.
- [16] A. Joshi. To Learn or Not to Learn ... In G. Weiss and S. Sen, editors, *Adaptation and Learning in Multiagent Systems*, volume 1042 of *Lecture Notes in Artificial Intelligence*, pages 127–139. Springer Verlag, 1996.
- [17] A. Joshi, T. Drashansky, E. Houstis, and S. Weerawarana. SciencePad: An Intelligent Electronic Notepad for Ubiquitous Scientific Computing. In *International Conference on Intelligent Information Management Systems*, June 1995.
- [18] A. Joshi, T.T. Drashansky, J.R. Rice, S. Weerawarana, and E.N. Houstis. On Learning and Adaptation in Multiagent Systems: A Scientific Computing Perspective. Technical Report TR-95-040, Dept. Comp. Sci., Purdue University, 1995.
- [19] A. Joshi, N. Ramakrishnan, J.R. Rice, and E. Houstis. A Neuro-Fuzzy Approach to Agglomerative Clustering. In *Proc. IEEE Intl. Conf. on Neural Networks*, volume 2, pages 1028–1033. IEEE Press, July 1996.
- [20] A. Joshi, N. Ramakrishnan, J.R. Rice, and E. Houstis. On Neurobiological, Neuro-Fuzzy, Machine Learning and Statistical Pattern Recognition Techniques. *IEEE Trans. Neural Networks*, 8(1):18–31, 1996.
- [21] A. Joshi, S. Weerawarana, and E.N. Houstis. The Use of Neural Networks to Support “Intelligent” Scientific Computing. In *Proceedings Intl. Conf. Neural Networks, World Congress on Computational Intelligence*, volume IV, pages 411–416, 1994. (Orlando, Florida).
- [22] A. Joshi, S. Weerawarana, N. Ramakrishnan, E.N. Houstis, and J.R. Rice. Neuro-Fuzzy Support for PSEs: A Step Toward the Automated Solution of PDEs. *Special Joint Issue of IEEE Computer & IEEE Computational Science and Engineering*, 3(1):pp.44–56, 1996.

- [23] M. S. Kamel, K. S. Ma, and W. H. Enright. ODEXPERT: An Expert System to Select Numerical Solvers for Initial Value ODE Systems. *ACM Trans. Math. Software*, 19:44–62, 1993.
- [24] D. Leake. Case-based selection of Problem Solving Methods for Scientific Computation. <http://www.cs.indiana.edu/hyplan/leake/cbmatrix.html>, 1996.
- [25] K. Lehrer. *Theory of Knowledge*. Westview Press, Boulder, CO, USA, 1990.
- [26] V. R. Lesser. A Retrospective View of FA/C Distributed Problem Solving. *IEEE Transactions on Systems, Man, and Cybernetics*, 21(6):1347–1363, 1991.
- [27] S. Markus, S. Weerawarana, E.N. Houstis, and J.R. Rice. Scientific Computing via the World Wide Web: The Net//ELLPACK PSE Server. Technical Report CSD TR-97-022, Department of Computer Sciences, Purdue University, 1997.
- [28] S. McFaddin and J. Rice. Collaborating PDE Solvers. *Appl. Num. Math.*, 10:279–295, 1992.
- [29] Sun Microsystems. The Network is the Computer. Trademark, 1996.
- [30] Peter K. Moore, Can Ozturan, and Joseph E. Flaherty. Towards the Automatic Numerical Solution of Partial Differential Equations. In E. N. Houstis, J. R. Rice, and R. Vichnevetsky, editors, *Intelligent Mathematical Software Systems*, pages 15–22. North-Holland, 1990.
- [31] Mo Mu and J. R. Rice. Modeling with Collaborating PDE Solvers — Theory and Practice. *Computing Systems in Engineering*, 6:87–95, 1995.
- [32] T. Oates et al. Cooperative Information Gathering: A Distributed Problem Solving Approach. Technical Report TR-94-66, Computer Science, University of Massachusetts, Amherst, 1994.
- [33] N. Ramakrishnan. *Recommender Systems for Problem Solving Environments*. PhD thesis, Dept. Comp. Sci., Purdue University, 1997.
- [34] N. Ramakrishnan, A. Joshi, E.N. Houstis, and J.R. Rice. Neuro-Fuzzy Approaches to Collaborative Scientific Computing. In *Proceedings of the IEEE International Conference on Neural Networks*, volume I, pages 473–478. IEEE Press, 1997.
- [35] N. Ramakrishnan, A. Joshi, S. Weerawarana, E.N. Houstis, and J.R. Rice. Neuro-Fuzzy Systems for Intelligent Scientific Computing. In *Proc. Artificial Neural Networks in Engineering ANNIE '95*, pages 279–284, 1995.
- [36] N. Ramakrishnan and J.R. Rice. GAUSS: An Automatic Algorithm Selection System for Quadrature. Technical Report TR-96-048, Dept. Computer Sciences, Purdue University, 1996.
- [37] N. Ramakrishnan, J.R. Rice, and E.N. Houstis. Knowledge Discovery in Computational Science: A Case Study in Algorithm Selection. Technical Report TR-96-081, Dept. Computer Sciences, Purdue University, 1996.
- [38] P. Resnick and H. Varian. Recommender Systems. *Comm. ACM*, 40(3):56–58, March 1997.
- [39] J. R. Rice. Processing PDE Interface Conditions. Technical Report TR-94-041, Dept. Comp. Sci., Purdue University, 1994.
- [40] J.R. Rice. The Algorithm Selection Problem. *Advances in Computers*, 15:pp.65–118, 1976.
- [41] J.R. Rice. Methodology for the algorithm selection problem. In L. Fosdick, editor, *Performance Evaluation of Numerical Software*, pages 301–307. North Holland, 1979.
- [42] J. C. Schlimmer and L. A. Hermens. Software Agents: Completing Patterns and Constructing User Interfaces. *Journal of Artificial Intelligence Research*, 1(61-89), 1993.
- [43] Y. Shoham. Agent-Oriented Programming. *Artificial Intelligence*, 60(1):51–92, 1993.
- [44] R. G. Smith and R. Davis. Frameworks for Cooperation in Distributed Problem Solving. In Bond and Gasser, editors, *Readings in Distributed Artificial Intelligence*, pages 61–70. Morgan Kaufmann, 1988.
- [45] L.H. Tsoukalas and R.E. Uhrig. *Fuzzy and Neural Approaches in Engineering*. John Wiley and Sons, Inc., Third Avenue, New York, 1997.
- [46] L. Z. Varga et al. Integrating Intelligent Systems into a Cooperating Community for Electricity Distribution Management. *Intl. J. Expert Systems with Applications*, 7(4):42–49, 1994.
- [47] J. T. Vessey. Speed-to-Market Distinguishes the New Competitors. Technical report, Research Tech. Mgmt., 1991.

- [48] R.G. Voigt. Requirements for Multidisciplinary Design of Aerospace Vehicles on High Performance Computers. Technical Report ICASE Report No. 89-70, ICASE, NASA Langley Research Center, Hampton, VA, 1989.
- [49] Jr. W. R. Johnson. Anything, Anytime, Anywhere: The Future of Networking. In D. Leebaert, editor, *Technology 2001: The Future of Computing and Communications*. MIT Press, Cambridge, MA, 1992.
- [50] S. Weerawarana. *Problem Solving Environments for Partial Differential Equation Based Systems*. PhD thesis, Dept. Comp. Sci., Purdue University, 1994.
- [51] S. Weerawarana, E. Houstis, J. Rice, M. Gaitatzes, S. Markus, and A. Joshi. Web//ELLPACK: A Networked Computing Service on the World Wide Web. Technical Report TR 96-011, Department of Computer Sciences, Purdue University, 1996.
- [52] S. Weerawarana, E. N. Houstis, J. R. Rice, A. C. Catlin, C. L. Crabill, C. C. Chui, and S. Markus. PDELab: An Object-Oriented Framework for Building Problem Solving Environments for PDE Based Applications. In *Proc. Second Annual Object-Oriented Numerics Conference*, pages 79–92, Rogue-Wave Software, Corvallis, OR, 1994.
- [53] S. Weerawarana, E.N. Houstis, J.R. Rice, A. Joshi, and Houstis C.E. PYTHIA: A knowledge based system to select scientific algorithms. *ACM Trans. Math. Software*, 22(4):447–468, 1996.
- [54] R. Wesson et al. Network Structures for Distributed Situation Assessment. In Bond and Gasser, editors, *Readings in Distributed Artificial Intelligence*, pages 71–89. Morgan Kaufmann, 1988.
- [55] M. Wooldridge and N. Jennings. Intelligent Agents: Theory and Practice. (submitted to Knowledge Engineering Review).