SciNapse: A Problem-Solving Environment for Partial Differential Equations

ROBERT L. AKERS, ELAINE KANT, CURTIS J. RANDALL, STANLY STEINBERG, AND ROBERT L. YOUNG

SciComp Inc.

♦   ♦   ♦

The SciNapse code generation system transforms high-level descriptions of partial differential equation problems into customized, efficient, and documented C or Fortran code. Modelers can specify mathematical problems, solution techniques, and I/O formats with a concise blend of mathematical expressions and keywords. An algorithm template language supports convenient extension of the system's built-in knowledge base.

♦

artial differential equations can represent the essence of a broad range of problems in engineering, science, and other technical fields. Those who need to solve systems of these equations numerically, however, often do not have the right combination of knowledge—expertise in a technical discipline, in numerical analysis, and in computer science or programming—to do an efficient job of it. Therefore many researchers have been attracted to the vision of a *problem-solving environment*, or PSE, that could provide comprehensive help in solving systems of PDEs.[1]

We have worked on such PSEs for nearly 10 years, using applications ranging from wave propagation[2,3] and computational fluid dynamics to financial modeling.[4] Our current PSE, SciNapse, focuses on code generation; it is a system for solving scientific computing problems without actually programming by hand, and could function as part of a larger PSE system. SciNapse has generated codes that solve

- ♦ the unsteady Maxwell's equations in 3D dispersive, anisotropic media;
- ♦ the Black–Scholes equation for valuation of multiple-asset derivative securities in computational finance, including the effects of stochastic asset price volatility

and interest rates, and discrete sampling of spot prices;
- ♦ nonlinear, multidimensional, multispecies reaction diffusion equations for chemical and nuclear applications; and
- ♦ time-domain solution of viscoelastodynamic equations in 3D anisotropic media.

SciNapse currently can generate codes that solve a wide range of initial boundary value problems for systems of PDEs, as well as many steady-state problems. The system can apply finite-difference methods to any region that can be mapped to a rectangle in any number of dimensions, though codes using very high dimensions may require excessive computational power to execute.

The codes SciNapse generates for these applications can include features such as general coordinate transformations and grid generators, various linear solvers and preconditioners, higher-order differencing techniques, automatic interpolation of equation parameters from multidimensional tabular input data, jump conditions in both space and time dimensions, free boundaries, and imposition of solution constraints such as positivity. SciNapse currently does not have the data structures to represent finite-element solutions on unstructured grids. It also currently lacks the knowledge to solve PDE problems with boundary element, boundary integral, or

Monte Carlo techniques. We are in the process of adding algorithms for solving nonlinear PDEs and knowledge about nonlinear optimization algorithms. The goal is to generate codes in which the finite-difference PDE solution becomes the forward engine for solving multiparameter inverse problems via nonlinear optimization.

Crucial aspects of our PSE technology that make it useful in practice include the specification language, refinement of specification to code, reporting and help systems, customization of the environment, and the use of templates for representing algorithms. Problem specifications in SciNapse typically range from several lines to a half page, and the synthesized code is usually about 25 times the length of the specification. On a 200-MHz personal computer, the system generates code at about 50 lines per minute.

History and plans

The SciNapse project has evolved over about eight years with an average of two or three computer scientists as implementers and one or two mathematicians and physicists as advisors, testers, and users. Three of the project members have been involved continuously. SciNapse currently comprises approximately 75,000 lines of Mathematica code. About half (the most rapidly growing section) represents knowledge of mathematics and PDE solutions, about one quarter is general computer science knowledge, and about one quarter is problem-solving system and interface support.

We originally developed the system at Schlumberger, focusing on oil-field applications such as the modeling of seismic and acoustic logging tools, with Connection Machine Fortran as a target language. Several of the generated codes, after some hand tuning, were used in company design projects. After founding SciComp and transferring the project to this new company, we gradually shifted much of the focus to financial applications such as the pricing and hedging of derivative securities and the generation of C code. The generated codes have been extensively tested, with some in use at major investment banks.

At this point the system development is well advanced in the financial modeling area, and SciComp plans to release a commercial finance product within the next year. We plan to produce PDE packages for students and professional engineers in the future.

An overview of SciNapse

A critical feature of any PSE for PDEs is hav-

ing an easily understood, high-level *problem specification language*, or PSL, that supports natural descriptions of the problem's geometry, mathematics, and desired interfaces. SciNapse's PSL allows problems and solution strategies to be stated concisely and abstractly, much as one might describe the problem to a colleague. The language allows easy specification and modification of input and output formats, and (optionally) of solution algorithms and interfaces to numerical libraries.

Because the specifications are at such a high level, the synthesis system must bridge the gap from coordinate-free mathematics to target-language code. SciNapse has an extensive knowledge base of coordinate-free constructs (such as Laplacian), equations (such as Navier-Stokes), discretization rules (such as Crank Nicholson), solvers (such as preconditioned conjugate gradient and SOR), optimization rules, and so on that it applies to produce code. Users can extend and customize the knowledge base. Using this knowledge base, the system converts equations into discretized, scalar components, and weaves

◆

SciNapse allows users to focus on physics and mathematics rather than on the mechanics of combining library components.

◆

them into algorithm templates that can be specified by name. The system chooses appropriate data structures and generates a pseudocode solution that is then translated into the desired target language. Along the way, it optimizes the mathematical problem, the abstract algorithms and data structures, and the resulting code.

The user interface in SciNapse consists of the PSL, "level summaries" that track the progress of the code generation process, an information system or help system that contains a hyperlinked glossary and specification language descriptions, and a simple graphics display system.

SciNapse is built on top of a general-purpose knowledge-based system that includes objects, rules, and the use of computer algebra (the entire system is implemented in Mathematica). Much of the algorithmic information is stored

```
(* Geometry *)
    Region[0 <= x <= 2 && 0 <= y <= 3 && 0 <= t <= 1,
      Cartesian[{x,y},t]];
(* The PDE *)
    When[Interior, der[f,t] == a laplacian[f]; a == 2; SolveFor[f]];
    Variable[a, Scalar, DiffusionConstant, "diffusion constant"];
(* Boundary Condition (for full boundary) *)
    When[Boundary, f == 0];
(* Boundary Condition (override on one edge) *)
    When[y == 3, f == 100;];
(* Initial Condition *)
    When[InitialValue, f == 50 Sin[Pi x/2] Sin[Pi y/3]];
(* Parameters *)
    Inline[{a}];
(* Evolution Algorithm *)
    Movie[frames == 5];
(* Discretization *)
    CrankNicholson;
    RelativeErrorTolerance[.01];
    Double;
(* Solver Algorithm* )
    ConjugateGradient[SSORPre[omega == 1.0, maxit == 1]];
(* Runtime Interfaces *)
    Output[f, OneFile, Labelled, "PCG-SSOR.dat"];
```

Figure 1. Specification for a diffusion problem.

as transformation rules or in extensible collections of very abstract templates.

Important features of code synthesis

Code *synthesis* systems such as SciNapse are much more powerful for numerically solving PDEs than are numerical prototyping environments such as Matlab, the classical assembly of components from a numerical library, or even a C++ numerical library with an intelligent front end. SciNapse can, in fact, do all the things these systems can do, but it is far more flexible.

A system that can automatically synthesize code makes it possible to have specification languages of an even higher level than prototyping languages. For example, simple keywords can stand for known equations, algorithms, and discretization methods. Synthesis allows the generation of efficient implementations, whereas interpreted code executes slowly and may not scale up to large data sets. Because the knowledge-based kernel of SciNapse is a general intelligent agent, the system could be extended in the areas of equation and algorithm selection, model identification, and results analysis.

The combination of highly abstract SciNapse templates with code synthesis allows users to focus on the physics and mathematics of problem solutions rather than on the mechanics of combining library components. This methodology allows the use of arbitrary higher-order methods without explicit target-language representation, validates whether specific compositions satisfy assumptions

made by individual components, and performs *global* optimization of programs. Components can be created to have the target language, data structures, and interface needed for a particular environment, yet be easily modified for reuse.

Although some of SciNapse's templates generally correspond to the modules or objects found in conventional or C++ libraries (solvers, time-steppers, evolution algorithms), the combination of highly abstract templates with code synthesis has several important differences. With most libraries, you have to write the stencil setups or discretization routines yourself, make data-structure decisions, and explicitly call many low-level routines to initialize, allocate, free, and finalize codes. SciNapse's templates are not code subroutines that are called. Rather, they are very abstract skeletons of algorithms (such as solvers) that are independent of spatial dimensionality, data structures, and target language considerations. The system automatically generates optimized data structures based on the automatically generated stencil arrays or discretized equations and uses those data structures in custom-generated solvers. It automatically optimizes control structures according to dataflow requirements. Further global optimization is used to generate code appropriate to the specified output architecture (such as serial, parallel, or distributed) and language (the choice of which may allow array-level operations or may affect the order of nested loops).

The ability to generate a variety of program interfaces makes synthesis especially useful for generating target codes suitable for PSEs having paradigms such as plug and play or dynamic assembly. Flexible generation of interfaces also facilitates incremental replacement of components of legacy codes with automatically synthesized codes, retaining the high-level specifications for modification and reuse. Easily understandable specifications also promote cooperative use of PSEs. Although we are quite some distance from this full vision, high-level problem specification with synthesis of target codes provides many opportunities for realizing the full power of PSEs.

Specifying PDE problems in SciNapse

To illustrate the problem specification language, we will describe a simple diffusion equation and summarize how it is specified in SciNapse.

Specifying the equation

Diffusion equations have many applications, such as modeling the cooling of electrical and mechanical components, the diffusion of chemicals, and the diffusion of populations in biology. Despite their simplicity, solving even the most elementary diffusion equations numerically is nontrivial. The diffusion equation specification illustrated in the PSL (Figure 1) is often given in textbooks in coordinate-free notation as

$$\frac{\partial f}{\partial t} = a\nabla^2 f, \tag{1}$$

where ∇^2 is the Laplacian (written `laplacian[]` in the PSL) and a is the diffusion coefficient. When described in Cartesian coordinates with x and y as spatial variables and t as the time variable, Equation 1 becomes

$$\frac{\partial f}{\partial t} = a\left(\frac{\partial^2 f}{\partial x^2} + \frac{\partial^2 f}{\partial y^2}\right), \tag{2}$$

where the solution $f = f(x, y, t)$ is a function of the space and time variables.

The physical region in x-y-space where the PDE is to be solved is the simple rectangle $0 \leq y \leq 3$ and $0 \leq x \leq 2$. A Dirichlet boundary condition is chosen for the full boundary; that is, the solution is set to zero on each piece of the boundary, except where overridden by $f = 100$ on the $y = 3$ edge. Time is the infinite interval starting at $t = 0$, and the initial condition is

$$f(x, y, 0) = 50 \sin\left(\frac{\pi x}{2}\right) \sin\left(\frac{\pi y}{3}\right). \tag{3}$$

Because the solution of this problem decays rapidly to zero, we set the final time t_f to a rather low value, say 1. The goal is to create a movie of the solution: a sequence of 3D plots of the solution $f(x, y, t)$ at equally spaced times.

Figure 1 gives a problem specification that will generate a Fortran code to solve this problem. Comments, delineated as *(* comment *)*, are optional. The specification language mirrors a typical mathematical specification to make it as natural as possible. So the statements in Figure 1 are direct translations of the mathematical statements into specifications in the PSL. We discuss such specifications in detail elsewhere.[5]

For straightforward problems, this is the most that needs to be specified. SciNapse makes any remaining choices required to translate the problem specification into a numerical code. For example, SciNapse will select appropriate discretization methods and solvers if the specification does not indicate what to do.

How specifications are organized

A specification typically has a mathematical part and an interface part.

The mathematical specification is organized in an outline, with the description of the geometric region's parts providing the structure. Equations can be associated with the interior, boundary and boundary parts, and initial time. This outline also provides a structure for expressing the discretization rules to be used on particular parts of the regions and for algorithms used to solve the problem. Equations have a similar parts structure including systems of equations, individual equations, and terms of equations. Again, discretizations or other directives can be applied to systems, equations, or terms.

The interface part includes the specification of the target code language, how variables are to be initialized (read interactively, from files, inlined, computed from functions, and so on), and output formats. SciNapse can generate data appropriate for use with visualization packages, it can read and interpolate data sets, and it has a simple interface to Mathematica's graphics utilities. The system also can generate some types of external subroutines and calls to them. We will extend it so that generated code can include calls to sophisticated graphics packages, extensive external libraries, and complex external data sets.

A specification can occur at multiple levels, as shown in Figure 2. Specifying at the highest possible level of abstraction maximizes clarity of understanding and ease of revision and optimization. Most specification is done at the top three levels in the figure. For example, the user can simply specify keywords like `Convection Diffusion` to obtain the convection-diffusion equation or `CrankNicholson` to get the Crank-Nicholson time discretization. The coordinate-free level allows the specification of differential equations in terms of the divergence, gradient, curl, and Laplacian.

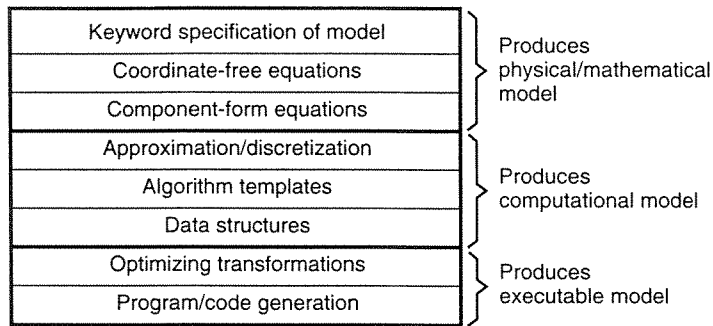Many problems require only keywords and co-

Keyword specification of model	Produces physical/mathematical model
Coordinate-free equations	
Component-form equations	
Approximation/discretization	Produces computational model
Algorithm templates	
Data structures	
Optimizing transformations	Produces executable model
Program/code generation	

Figure 2. Levels of specification and abstraction in the SciNapse code generation system.

IEEELevelSummary.nb

File   Edit   Cell   Format   Input   Kernel   Find   Window   Help

Summary for CoordinateFree level:
TotalRegion1:
   Global:
      Discretization:
         {CrankNicholson}
      When[Interior,
         Eq:
            $\frac{\partial f}{\partial t} == a \triangle[f]$
         Eq1:
            $a == 2$
         Discretization:
            {CrankNicholson}
      ]
      When[Boundary,
         BCEq:
            $f == 0$
            When [y == 3 && x >= 0 && x <= 2,
               BCEq1:
                  $f == 100$
            ]
      ]
      When[InitialValue,
         IVEq:
            $f == 50 Sin[\frac{\pi x}{2}] Sin[\frac{\pi y}{3}]$
      ]
      Outputs:
         ResultEq:
            f
         Discretization:
            {CrankNicholson}
   Discretization:
         {CrankNicholson}

Figure 3. Summary at the coordinate-free level.

ordinate-free specifications. At the next step down, the component level, differential equations can be specified using a particular coordinate system for the variables and vector components. The user rarely specifies directly at the discrete level. SciNapse refines specifications into successively lower levels, with the final stages involving a numerical algorithm specified in the internal programming language that as a last step is translated into a standard compilable programming language.

Program refinement

Paralleling the steps an expert might take when programming by hand, SciNapse automatically refines a specification in a stepwise fashion from the most abstract level through several more concrete levels, finally creating a numerical code. The system constructs a complete program description at a given level by combining constructs from that level in the specification with refinements of constructs from earlier levels. It expands keyword specifications at the appropriate process level. The mathematical and programming constructs used during refinement are represented as objects in the knowledge base; transformation rules expand the keywords and more abstract constructs into elaborated constructs. After each stage, SciNapse checks the problem state for consistency appropriate to that level of abstraction.

First, if the PDEs, boundary conditions, or initial conditions are given in coordinate-free form, the system refines them into component equations using a specific coordinate system. The resulting scalar equations must be transformed from continuous to discrete form, so SciNapse analyzes the equations and selects appropriate discretization methods (unless overridden by recommendations from the user that are consistent with the rest of the specification).

Next comes the algorithm selection level. User design decisions, if given, are validated; if they are not specified or not valid, SciNapse makes the decisions based on the nature of the problem. Next come data-structure decisions, which the system normally makes without the user's input, based on the equations, discretizations, and algorithm choices. SciNapse then makes generic programming choices at the program level, constructing a pseudocode from the templates as elaborated by the data-structure choices. Control structures and operations are optimized at the pseudocode level. Finally, SciNapse generates target-language code from the optimized pseudocode.

"Level summaries" report on progress

After each stage of refinement, SciNapse generates a "level summary," an output report that enables users to see directly how their specifications were understood and integrated with the system's inferences. The summary structures reflect the geometric organization of the problem specification, and like the specifications, use a notation as close to that of traditional math as possible. We want SciNapse to be a mathemat-

ical partner, solving the problem with the user.

Figure 3 shows a small part of the Mathematica notebook that SciNapse produces after the coordinate-free stage in processing the specification in Figure 1. Here the function f in the equations appears with no arguments in the Laplacian expression. Users manipulate the standard Mathematica notebook interface for viewing and interacting with level-summary reports. The menu buttons at the top show possible interactions. The nested brackets at the right break the report contents into hierarchically grouped cells that can be closed and opened to show different levels of outlining. Clicking on an active element (such as a variable or equation name) produces a display elaborating that element. Combining this sort of hyperlinking with outlining provides flexibility without requiring that the user remember too much context.

Figure 4 shows the component-level summary, with equations expanded into the variables determined by the choice of a Cartesian coordinate system. The Laplacian is expanded into derivatives of the spatial coordinates. Cells marked with downward-pointing arrowheads, such as the line "Summary for Discrete level" near the bottom, indicate groups that have been closed; the visible line is the first cell of the group. To open or close the display of a group, a user double-clicks on its bracket.

Figure 5 shows the problem after discretization. The derivatives have now been replaced by difference expressions, and f is now a function of the space and time coordinates. The values for the solver and the evolution algorithm are now present.

The general-purpose knowledge-based system

SciNapse's high-level specification language supports and encourages users in the iterative design and experimentation so critical in problem solving. Automatically translating such specifications into efficient code requires code synthesis.

Our code synthesis system is built on top of a general-purpose knowledge-based system, which we have written in Mathematica. The system includes an integrated object system, rule system, and planning system. It supports a mixed decision style in which the system can make any design decisions not specified by the user. For instance, the system will make any discretization or algorithm choices left unspecified. The algorithmic information is stored as trans-
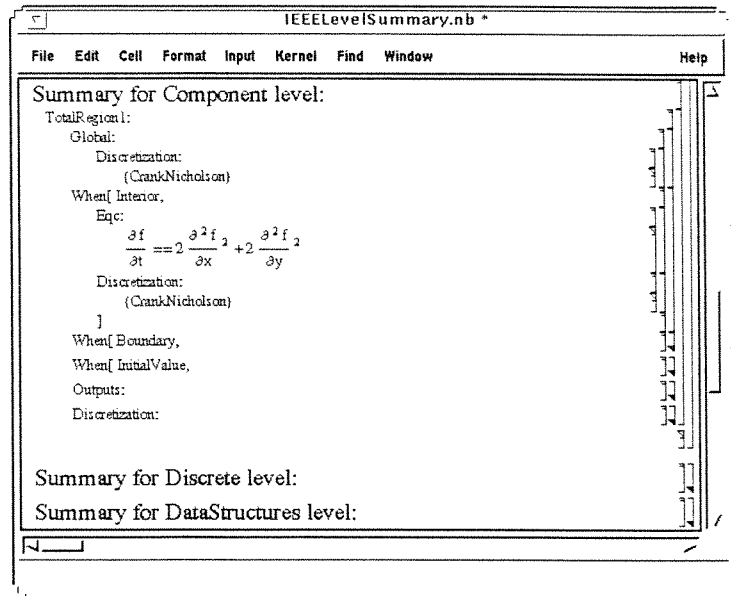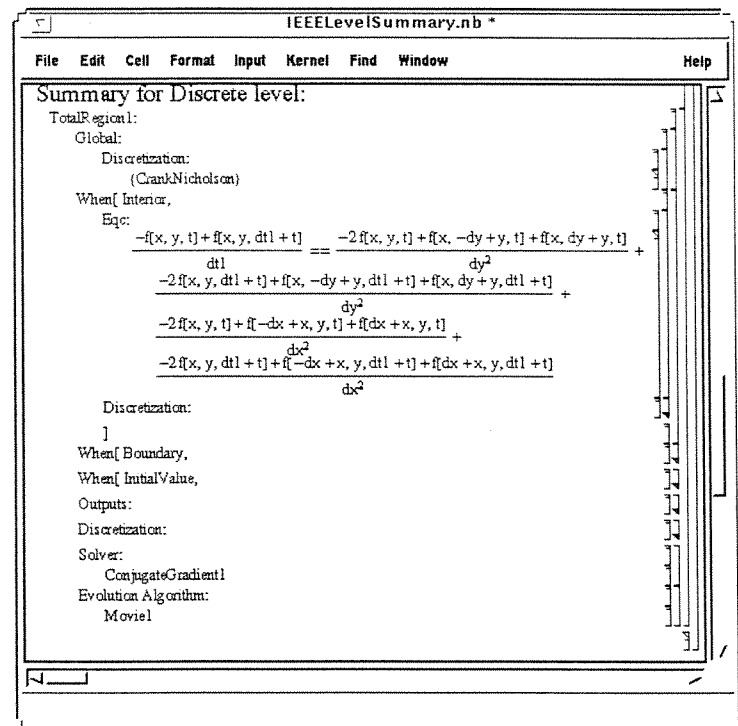
Figure 4. Summary at the component level.

Figure 5. Summary at the discrete level.

formation rules or in an extensible collection of templates. SciNapse contains its own internal numerical programming language that is easily translated to commonly used languages such as Fortran or C.

Representation via objects and rules

SciNapse objects explicitly represent common mathematical constructs such as a geometric region or part, a system of equations, an individual equation, a term, or a problem variable. Objects also represent programming constructs such as a linear solver, a subroutine, or a program variable. Objects carry attributes describing methods for elaborating the object and its properties, attributes describing associated design decisions, and explanatory help information.

Refinement rules detail the methods referred to in the object attributes. For example, rules may expand keywords and abstract constructs, make discretizations, or check specifications for internal consistency of physical units and tensor order.

SciNapse creates instances of objects during the refinement process that persist throughout the entire synthesis. Some attributes on object instances record relationships between objects, such as linking an equation to all the variables it uses and linking a variable to all its equations. Users may inspect object instances to examine intermediate results such as an equation's discretization.

Explicitly represented design choices

SciNapse's knowledge base explicitly represents design choices to be made by the user or the system. Considering only the world of initial boundary value problems, hundreds of design choices exist, leading to thousands of possible programs. To maximize flexibility, we want users to be able to specify any design choices they desire, but *not* be required to specify those for which they have no preferences or those that should be obvious to the system. Some choices are simple selection among algorithms, but the algorithms are represented as highly abstract templates that are customized to the problem. Other choices involve the method of input (such as reading from a file, entering interactively, or computing from a subroutine), file formats, and data structures (for instance, an array can be full, diagonal, time-independent, stencil, and so on). These choices help interface the generated code to the rest of the PSE. Typically the system will optimize programs without advice from the specification.

To make design choices easier for users to find and understand and easier for developers to extend, design choices are represented by attributes on objects. For example, the DiscretizationMethod attribute on an equation, system, or term can have values such as CrankNicholson or CentralDifference[4] (the "4" indicates fourth order). If a design choice is identified by a unique keyword, simply mentioning that keyword in the appropriate "When[region, ...]" context is sufficient (see Figure 1).

Associated with each choice can be constraint rules, heuristic rules, and defaults. Constraints are definite requirements on the values allowed. They can be based on previous choices or on the nature of the mathematical problem. For example, some preconditioner choices can be ruled out because of solver choices and matrix properties. Constraints cannot be ignored: violating them will lead to incorrect codes. Heuristics are system suggestions and can be overridden by the specification. Defaults specify what to do in the absence of decisions made by the heuristics or the specification file.

The planning system

SciNapse's planning system sets goals to instantiate objects and then refine them by filling in attributes. This includes an agenda mechanism that schedules refinements and design decisions for making selections such as algorithms and data structures. The planning system uses the method descriptions to automatically determine an ordering of refinements that ensures that all data to make choices is in place before a decision is considered.

Given more knowledge, the planning mechanism could be used in several new ways. For example, it could derive and/or select equations or variants thereof based on problem properties, set up default parameters, or analyze data from program runs.

Context-dependent code optimization

SciNapse's internal representation of numerical programs is independent of target language, and records the maximum parallelism possible in algorithms and equations. In this abstract representation, context-dependent global optimizations are easy and productive. These include early elimination of problem variables not needed to compute the desired outputs, and selection of optimal data structures, such as arrays based on problem-specific stencil sizes. The system also applies standard optimizations such as introducing temporary variables, loop merging and loop unrolling, and algebraic simplification.

Although the internal representation is language-independent, the last step of optimization before code generation does provide for target-

language-specific optimizations. Currently Sci-Napse generates codes in C and Fortran 77, so, for example, loop enumerations are reversed for C and Fortran arrays, and dynamic allocation is available in C.

Because the parallelism inherent in a problem solution is retained until the very last step of code synthesis, adding architecture-specific code generators is reasonably straightforward (at one point we added a generator for Connection Machine Fortran, which is similar to Fortran 90). Here, no loop enumerations will be generated for array operations. We anticipate generating codes for distributed architectures using languages like MPI in the future.

Extensive use of computer algebra

The availability of a computer algebra makes it easy for SciNapse to estimate truncation errors, compute derived quantities (such as sensitivities), transform equations into general coordinates, and discretize terms with arbitrarily high-order differencing formulas. Computer algebra capabilities allow the system to represent templates independently from the problem details, yet still generate efficient code and globally optimize it—for instance, by determining stencils from discretized equations and by using the knowledge about stencil shape to optimize array representations.

Testing

Correctness of the codes generated by Sci-Napse is ensured in several ways, both by overall system testing and by testing of individual codes. We have collected a large number of test problems for which we compare the results of SciNapse-generated codes against analytic solutions (when possible) and also against some alternative computations. We have a large regression suite that we rerun whenever the system is modified. In testing individual problems, the equations of the specification are checked for consistency of tensor order and physical dimensions (if specified). Also, the truncation error for the approximations is automatically computed and inserted as a comment in the code. A numerical convergence rate test wrapper is available to verify that the synthesized code performs in accordance with the symbolically calculated truncation error.

Customization

A knowledge-based PSE should be easy to customize for different application areas. SciNapse's

PSL allows any chunk of specification language to be put in a separate file and be called by name. A typical use for this "macro" construct is to specify frequently used systems of equations with "equation generators." Conditional specification keywords can provide more customization. For example, a keyword parameter to the Navier-Stokes equation generator can control whether to include an internal or a total energy conservation equation.

Macro invocations can include keywords and also rules that replace any symbol with another symbol, constant value, or expression. Such rules can change variable names, set variable values, substitute expressions, or even pass in a set of problem-specific state equations.

SciNapse users and developers can customize algorithms, discretization methods, and decision rules. The template language, discussed in the next section, enables the definition of algorithms tailored to a specific application area. Language constructs exist to name collections of discretization methods and to define new ones. It is also possible to add new constraints and heuristics on design choices.

Using the preceding constructs, developers can define all the common names of equations, boundary conditions, derivatives and sensitivity terms, outputs, solvers, or any other design choice. Application users can then write specifications completely in terms of those keywords, yielding very concise specifications that are easy for practitioners in that application area to understand.

Specifying a problem at the highest possible level of abstraction maximizes clarity of understanding.

Algorithm templates

Some of the expert knowledge brought to bear in code synthesis is represented as abstract algorithm templates.[6] These templates capture the essentials of an algorithm; they are elaborated with problem-specific aspects and then translated into the chosen target language. Templates support the scalability of a PSE because they make it easy for end users as well as developers to add new algorithms. Many linear solvers and other algorithms are recorded in the literature in a form very similar to SciNapse's templates.[7,8]

The Template Declaration Language

Templates are defined in an algorithm de-

```
Template[ConjugateGradient, SolverTemplate,
  SubroutineName[CGSolver],
  LocalVars[p[SameAs[y], "search direction vector"],
    q[SameAs[y], "image of search direction vector"],
    zz[SameAs[y], "conditioned residual"],
    alpha[Real, "multiple of search direction vector"],
    beta[Real, "orthogonalization constant"],
    rho[Real, "temporary variable"],
    old[Real, "old value of rho"],
    iter[Integer, "iteration counter"]],
  Templates[Preconditioner],
  Code[r = y - svar.xx;
    old = r.r;
    p = 0;
    Until[iter, maxit, StopCode,
      comment["Apply Preconditioner."];
      zz = r;
      Preconditioner[r, svar, zz];
      rho = r.zz;
      beta = rho / old;
      p = zz + beta p;
      q = svar.p;
      alpha = p.q;
      alpha = rho / alpha;
      xx = xx + alpha p;
      r = r - alpha q;
      old = rho]]];
```

Figure 6. Template declaration for a conjugate gradient solver.

scription language that is independent of spatial dimensionality, data structures, and target language considerations. The rich language of pseudocode expressions includes various matrix operators (for example, dot products, arithmetic operators, transpose, norm, diagonal, and identity) and transcendental operators, as well as the usual scalar operators for numbers and Booleans. It also includes easily recognizable procedural control structures. In addition to conventional loop, while, and until constructs, special constructs allow looping over the range of abstractly specified subsets or supersets of array dimensions, and mapping of operation sequences over variable lists. A special dimension can be identified and loops defined over just the special or the non-special dimensions. Some parallel constructs record opportunities for distributed computation.

Data declarations in templates are more general than those of Fortran or C. SciNapse's type annotations, rather than providing fixed storage declarations, allow declarations in terms of configurations of other variables, including those from the problem's equations. In this way they resemble the *automatic data objects* of Fortran 90 and avoid the need for large libraries of special cases. Declarations can state that a new array has the same shape as another array, or the same shape but with some dimensions deleted or added, or with dimensionality and range based

on the ranges and shapes of other problem or template variables.

Templates can be composed similarly to subroutines in a programming language. The need for a subordinate or "called" template is indicated by including the identifier in the local "Templates" declaration. In the pseudocode, the "call" is represented as a reference to that identifier, signifying only that some class of routine, such as a solver or a preconditioner, is needed at that point. The specific choice of solver or preconditioner is not fixed in the calling template declaration, but rather is a design choice. During synthesis, the template identifier takes on the "value" of an instance of the subordinate template class selected by the specification or system's decision mechanism. For example, in Figure 6, we declare the identifier Preconditioner. Its appearance in the code indicates where the preconditioner call will occur. Arguments to the called template may be automatically inferred by rules.

SciNapse's knowledge base represents each template as a class in a hierarchy of template classes. This hierarchy allows for a separation of concern in template declarations as well as defining families of options for various kinds of algorithms. Among the major classes of templates are solvers, time evolution templates, stepping algorithms, interpolations, and inverse problem or test wrappers (for automatically generating convergence rate testing code, for instance). Using normal object inheritance, aspects like variables or template place holders defined high in the hierarchy are known in all the subclasses and need not be redeclared. For example, in Figure 6, the StopCode declaration is inherited by the ConjugateGradient solver from its superclass, SolverTemplate.

A partial class hierarchy for time evolution templates appears in Figure 7.

Names for both variables and template placeholders can also be drawn from other templates. Mechanisms are available to refer to the variables that exist in the template's "caller" or in the SciNapse problem state, and to pass information from templates to their subordinate templates.

A template declaration can also provide an indexing scheme, that is, a function that transforms the generic indexing operations in the code declaration to more concrete data representations and operations appropriate for the synthesized data structures.

A template declaration can include programming directives such as whether the system

should expand the template in line or encapsulate it as a named subroutine.

Template translation

SciNapse contains a rich collection of templates written by its developers and submitted to the system via a translator that builds them into the system's knowledge base of objects and rules. The same translator will be available to SciNapse users, so that if the templates they need are not already in the system, they may declare them and have them incorporated as first class objects in the synthesis environment. To ensure that all the pieces fit together in a sensible manner, the language supports good syntax-checking and type-checking, and a variety of semantic checks both at template declaration time and at synthesis time.

When a template is declared, the template translator automatically creates an attribute in the template class corresponding to each variable that the template can access and each subordinate template place holder; it also creates methods for filling in the values of these attributes.

Templates during synthesis

During synthesis, SciNapse creates an instance of the selected template class and fills in its attributes as the problem demands. Next, the system fills in the names of the specific variables that can be accessed by extracting them from the problem state or from calling templates, or, in the case of local variables, by defining variables of the appropriate kind and shape. For the attributes corresponding to template placeholders, the methods instantiate the new templates, and the process continues recursively.

In the next step of template processing, the system builds pseudocode for the various template instances. The work is done by methods derived by the template translator from the code part of the template declaration. The first step is to fill in the pseudocode from the declaration with the variables and code for subordinate templates previously derived. Thus, where a subordinate template is called, the template code (see later) for that template is inserted, either in-line or as a subroutine call. The pseudocode for the conjugate gradient solver in the diffusion example appears in Figure 8.

Next, the indexing function, if present, is applied to the pseudocode. This function transforms the abstract indexing given in the pseudocode to a concrete form appropriate for

```
EvolutionTemplate
    SteadyState
    TimeDependent
      TimeDependentEvolve
        Movie
        Evolve
        Motion
        DiscreteEventsGeneral
      TimeDependentFixed
        TestingTemplate
          ConvergenceRateTest
        EvolveFixed
        MotionFixed
```

Figure 7. Class hierarchy for evolution templates.

```
seq[assign[r1, g1 - SA1 . xx1],
  assign[old1, r1 . r1], assign[p1, 0],
  seq[assign[iter1, 1],
    seq[if[tol1 < 8*eps1,
      print["Warning: Linear tolerance is too small.,"
        tol1, " < 8*," eps1, " ."], seq[]],
      comment["Compute the static stopping criteria
        information"],
      commentAbout["The norm of the right hand side,"
        assign[normg1, norm[g1]]],
. . .]]]
```

Figure 8. Pseudocode for preconditioned conjugate gradient.

```
call[CGSolver, eps1, f, g1, iMax, jMax, SA1]
```

Figure 9. Template code for calling conjugate gradient.

the shape of the data structures for the particular variables in the code. No such transformation is specified for the conjugate gradient solver in our example.

Finally, if the template is to be cast into a subroutine, SciNapse encapsulates its code as a subroutine object with appropriate variable declarations and replaces the pseudocode with a call to the subroutine. This call will be incorporated into the code for the calling templates. If the code is to be in line, no action is needed. The resulting code for calling the conjugate gradient solver in the diffusion example appears in Figure 9.

Because SciNapse already can handle large and complex problems, we believe that the technologies that we have in place form the basis of a powerful PSE to assist with the solving of initial boundary value problems for partial differential equations. We plan to extend the system to include knowledge about more complex algorithms (such as fractional step-

ping methods) and geometries (such as unstructured grids). Toward this end, we continue to work on improvements to the system's algorithm and geometric representations as well as adding objects and rules to the knowledge base.

SciNapse has now been tested by many users: some PDE experts, some not. A special financial package called SciFinance,[4] which is built on top of SciNapse, is being used by major financial institutions. While this preliminary testing has uncovered a number of problems, none have to do with the fundamental design of the system. In fact, all reports indicate that the system is easy to understand and use, and impressive in what it can do. ♦

Acknowledgments

References

1. "Problem-Solving Environments Home Page," WWW: http://www.cs.purdue.edu/research/cse/pses/.
2. E. Kant, "Synthesis of Mathematical Modeling Software," *IEEE Software*, Vol. 10, No. 3, 1993, pp. 30–41.
3. E. Kant et al., "Scientific Programming by Automated Synthesis," in M.R. Lowry and R.D. McCartney, eds., *Automating Software Design*, AAAI Press/MIT Press, Menlo Park, Calif., 1991, pp. 169–205.
4. C. Randall and E. Kant, "Numerical Options Models Without Programming," *Proc. IEEE/IAFE 1997 Conf. Computational Intelligence for Financial Engineering (CIFEr)*, IEEE, Piscataway, N.J., 1997, pp. 15–21.
5. R.L. Akers et al., "Automatic Synthesis of Numerical Codes for Solving Partial Differential Equations," to be published in *Mathematics and Computers in Simulation*, special issue on Nonstandard Applications of Computer Algebra, H. Hong, E. Roanes Lozano, and S. Steinberg, eds.
6. E. Kant and S. Steinberg, "Automatic Program Synthesis from Abstract PDE Specifications," to be published in *Proc. 15th IMACS World Congress on Scientific Computation, Modelling and Applied Mathematics* [Berlin, 24–29 Aug. 1997], Int'l Assoc. for Math. and Computers in Simulation, c/o Dept. of Computer Science, Rutgers Univ., New Brunswick, N.J., 1997; also on CD-ROM.
7. R. Barrett et al., *Templates for the Solution of Linear System: Building Blocks for Iterative Methods*, SIAM, Philadelphia, 1994.
8. C.T. Kelley, *Iterative Methods for Linear and Nonlinear Equations*, Frontiers in Applied Mathematics, SIAM, Philadelphia, 1995.

Robert L. Akers is a computer scientist at SciComp Inc. He earned his MA and PhD in computer sciences at the University of Texas at Austin. His interests include programming languages, program synthesis, and the formal analysis and algebraic manipulation of programs. He is a member of the ACM.

Elaine Kant is the founder and president of SciComp Inc. She received a BS in mathematics from MIT and a PhD in computer science from Stanford University. Before starting SciComp, she worked for Schlumberger and taught computer science at Carnegie Mellon University. Her research interests focus on automating scientific problem solving and program synthesis. She is a fellow of the American Association for Artificial Intelligence and a member of the ACM and the IEEE Computer Society.

Curtis J. Randall is vice president of SciComp. He received the BS and MS degrees in nuclear engineering from the University of Wisconsin, Madison, and a PhD in applied science from the University of California, Davis-Livermore. Before joining SciComp, he worked at Lawrence Livermore Laboratory and Schlumberger. His research interests include computational physics and computational finance. He is a member of IAFE, SEG, and ASA.

Stanly Steinberg is professor of mathematics at the University of New Mexico and senior research consultant at SciComp. He received his PhD from Stanford University in mathematics. He has worked with computer algebra and numerical methods for partial differential equations for over 20 years. He is interested in finite-volume methods for discretizing partial differential equations and has applied methods in computational logic to stability problems for discretized partial differential equations.

Robert L. Young is the director of technology at SciComp. He received his PhD in computer sciences from the University of Texas at Austin. His research interests center on knowledge-based interactive systems. He is a member of ACM and AAAI.

Corresponding author: Elaine Kant, SciComp Inc., 5806 Mesa Drive, Suite 250, Austin, TX 78731, USA; e-mail, kant@scicomp.com; WWW, http://www.scicomp.com/.