A. Summary

This proposal brings together several important research areas including parallel compilers, data parallel SPMD libraries and object oriented programming models. We research combinations of these ideas that achieve high performance, in an approach that implies more work for the programmer than envisaged in systems such as High Performance Fortran (HPF), but which can more clearly be implemented in a robust fashion over a range of languages.

Building on earlier research in the Parallel Compiler Runtime Consortium (PCRC) project, we will be investigating a language model that combines characteristic data-parallel features from the HPF standard with an explicitly SPMD programming style. This language model, which we call the *HPspmd* model, is designed to facilitate direct calls to many of the established libraries for parallel programming with distributed data. By relaxing the HPF requirement of semantic equivalence between a data parallel program and a sequential Fortran program, and asking the programmer to explicitly manage non-local references, the task of the compiler is hugely simplified. On the other hand, with good bindings (eg class-library bindings) to suitable libraries, a level of expressiveness comparable to HPF can often be achieved. Because of the emphasis on direct library management of communication, the model is expected to fare better than HPF on *irregular* applications.

The requested funding will be used to further research and optimization in a Java instantiation of the language model (HPJava), to design and implement HPspmd bindings to a number of established SPMD libraries for data parallel programming, and to implement large applications within the framework. The libraries that will be targeted for interfacing from our the framework include some subset of Adlib, CHAOS, Global Arrays, MPI, DAGH and ScaLAPACK. Early work on HPJava relied on the runtime library Adlib, delivered in the PCRC project. Development of other library bindings and applications is expected to drive further refinement of the language model, especially in support of irregular problems.

Collaborators at the University of Peking will be working on Fortran and C++ bindings of the same model. Some support for related work on Java for large scale computing, led by Syracuse, will be provided by Sun Microsystems.

B. Table of Contents

Replace with NSF Form 1359.

C. Project Description

C.1 Motivation of the proposed work

It is generally accepted that data parallel programming has a vital role in high-performance scientific computing. The basic implementation issues related to this paradigm are well understood. But the choice of high-level programming environment remains uncertain. Five years ago the High Performance Fortran Forum published the first standardized definition of a language for data parallel programming [19, 30]. In the intervening period considerable progress has been made in HPF compiler technology, and the HPF language definition has been extended and revised in response to demands of compiler-writers and end-users [20]. Yet it seems to be the case that most programmers developing parallel applications—or environments for parallel application development—*do not* code in HPF. The slow uptake of HPF can be attributed in part to immaturity in the current generation of compilers. But there is the suspicion that many programmers are actually more comfortable with the lower-level Single Program Multiple Data (SPMD) programming style, perhaps because the effect of executing an SPMD program is more controllable, and the process of tuning for efficiency is more intuitive. (Partially, no doubt, this reflects a status quo where *expert* programmers build parallel programs and less experienced programmers merely use them.)

SPMD programming has been very successful. There are countless applications written in the most basic SPMD style, using direct message-passing through MPI [21] or similar low-level packages. Many higher-level parallel programming environments and libraries assume the SPMD style as their basic model. Examples include ScaLAPACK [4], PetSc [2], DAGH [36], Kelp [18, 31], the Global Array Toolkit [34] and NWChem [3, 28]. While there remains a prejudice that HPF is best suited for problems with very regular data structures and regular data access patterns, SPMD frameworks like DAGH and Kelp have been designed to deal directly with irregularly distributed data, and other libraries like CHAOS/PARTI [16, 37] and Global Arrays support unstructured access to distributed arrays. These successes aside, the library-based SPMD approach to data-parallel programming certainly lacks the uniformity and elegance of HPF. All the environments referred to above have some idea of a distributed array, but they all describe those arrays differently. Compared with HPF, creating distributed arrays and accessing their local and remote elements is clumsy and error-prone. Because the arrays are managed entirely in libraries, the compiler offers little support and no safety net of compile-time checking.

The proposed work will investigate a class of programming languages that borrow certain ideas, various run-time technologies, and some compilation techniques from HPF, but relinquish some of its basic tenets, in particular: that the programmer should write in a language with (logically) a single global thread of control, that the compiler should determine automatically which processor executes individual computations in a program, and that the compiler should automatically insert communications if an individual computation involves accesses is to array element held outside the local processor.

If these foundational assumptions are removed from the HPF model, does anything useful remain? In fact, yes. What will be retained is *an explicitly SPMD programming model*

complemented by syntax for representing distributed arrays, syntax for expressing that certain computations are localized to certain processors, and syntax for expressing concisely a *distributed form* of the parallel loop. The claim is that these are just the features needed to make calls to various data-parallel libraries, including application-oriented libraries and high-level libraries for communication, about as convenient as, say, making a call to an array transformational intrinsic function in Fortran 90. We hope to illustrate that, besides their advantages as a framework for library usage, the resulting programming languages can conveniently express various practical data-parallel algorithms. The resulting framework may also have better prospects for dealing effectively with *irregular* problems than is the case for HPF.

This proposal brings together several important research areas including parallel compilers, data parallel SPMD libraries and object oriented programming models. We research combinations of these ideas which achieve high performance with an approach that implies more work for the programmer than envisaged in systems such as HPF, but can more clearly be implemented in a robust fashion on a range of languages. Explicitly we are combining our research on the use of Java and Web technologies with the high performance SPMD libraries and some of the compiler techniques developed as part of HPF research. Java has many features that suggest it could be a very attractive language for scientific and engineering or what we now term "Grande" applications. Clearly Java needs many improvements both to the language and the support environment to achieve the required linkage of high performance with expressivity. This cannot be guaranteed but we have helped set in motion a community activity involving academia, government and industry (including IBM, Intel, Microsoft, Oracle, Sun and perhaps most importantly James Gosling from Javasoft) which is designed to both address language changes and the establishment of standards for numerical libraries and distributed scientific objects. The Java environment is still malleable and we are optimistic that this effort will be succesful and Java will emerge as a premier language for large scale computation. Our research will be aimed at multi-language programming paradigms but our new implementations will focus on Java exploiting existing high performance C++ and Fortran libraries. Our collaborater Professor Xiaoming Li from Peking University will be developing the Fortran and C++ aspects of this general high level SPMD environment. We can consider our work from either of two points of view; bringing the power of Java to a data parallel SPMD environment or alternatively researching the expression of data parallelism within Java. Note that we are adopting a more modest approach than a full scale data parallel compiler like HPF; we believe this is an appropriate approach to Java where the situation is changing rapidly and one needs to be very flexible.

We should stress what we are not doing! Many of the discussions of Java at the recent "Grande" workshops [23–25] have focussed on its use in distributed object and mobile or Web client based computing. In fact our group also is looking into this for composing large scale distributed systems. However in this proposal, we are addressing "hard-core" science and engineering computations where data parallelism and the highest performance are viewed as critical.

The work proposed in this project continues research conducted in the the Parallel Compiler Runtime Consortium (PCRC) project [14]. PCRC was a DARPA-supported project involving Rice, Maryland, Austin, Indiana, CSC, Rochester and Florida, with NPAC as

prime contractor. Achievements included construction of an experimental HPF compilation system [42], delivery of the NPAC PCRC runtime kernel (Adlib) [11] and early work on the design and implementation of HPJava [10].

C.2 Objective and expected significance

Our system aims to support a programming model that is a flexible hybrid of the data-parallel, language-oriented, HPF style, and the established and popular, library-oriented, SPMD style. We refer to this model as *HPspmd*.

Primary goals of the current project include

1. Providing a small set of syntax extensions to various base languages (including Java, Fortran, and C++). These syntax extensions add distributed arrays as language primitives, and introduce a few new control constructs, such as the distributed loop.

2. Providing bindings from the extended languages to various communication and arithmetic libraries. These may include libraries modelled on, or simply new interfaces to, some subset of Adlib, CHAOS, Global Arrays, MPI, DAGH, ScaLAPACK, etc. Supporting the libraries for irregular communication will be an important goal.

3. Testing and evaluating HPJava and the HPspmd model in general on large scale applications.

A major thrust of the proposed work will be on researching compiler (or preprocessor) support for our extended languages, and development of exemplar interfaces from the new languages to a subset of the libraries mentioned above. The research aspects of the proposed work involve investigation of compiler optimizations and safety checks peculiar to the new languages, extensions to the basic language model to improve support of irregular problems, and design of attractive class-library bindings for the various SPMD environments involved in the project.

The next four subsections overview the language extensions we are investigating, the libraries we will study, issues concerning low-level MPI programming in the proposed environment, and the parallel machine model.

C.2.1 HPspmd language extensions

We aim to provide a flexible hybrid of the data parallel and low-level SPMD approaches. To this end HPF-like distributed arrays appear as language primitives. A design decision is made that all access to *non-local* array elements should go through library functions—for example, calls to a collective communication library, or simply `get` and `put` functions for access to remote blocks of a distributed array. This puts an extra onus on the programmer; but making communication explicit encourages the programmer to write algorithms that exploit locality, and simplifies the task of the compiler writer.

For the newcomer to HPF, one of its advantages lies in the fact that the effect of a particular operation is logically identical to its effect in the corresponding sequential program. This means that, assuming the programmer understands conventional Fortran, it is very

easy for him or her to understand the behaviour of a program at the level of what values are held in program variables, and the final results of procedures and programs. Unfortunately, the ease of understanding this "value semantics" of a program is counterbalanced by the difficulty in knowing exactly how the compiler translated the program. Understanding the *performance* of an HPF program may require the programmer to have very detailed knowledge of how arrays are distributed over processor memories, and what strategy the compiler adopts for distributing computations across processors.

The language model we discuss has various similarities to the HPF model, but the HPF-style semantic equivalence between the data-parallel program and a sequential program is abandoned in favour of a literal equivalence between the data-parallel program and an SPMD program. Because understanding an SPMD program is presumably more difficult than understanding a sequential program, our language may be slightly harder to learn and use than HPF. But understanding performance of programs should be much easier.

The distributed arrays of the new languages will be kept strictly separate from ordinary arrays. They are a different kind of object, not type-compatible with ordinary arrays. An important property of the languages we describe is that if a section of program text *looks like* program text from the unenhanced base language (Java or Fortran 90 for example), it is translated exactly as for the base language—as local sequential code. Only statements involving the extended syntax behave specially. This makes preprocessor-based implementation of the new languages very straightforward, allows sequential library code to be called directly, and gives the programmer good control over the generated code—he or she can be confident no unexpected overhead have been introduced in code that looks like ordinary Fortran (for example).

In the baseline language we adopt a distributed array model semantically equivalent to to the HPF data model in terms of how elements are stored, the options for distribution and alignment, and facilities for describing *regular sections* of arrays. Distributed arrays may be subscripted with global subscripts, as in HPF. *But* a subscripting operation must not imply access to an element on a different processor. We will sometimes be refer to this restriction as the *SPMD constraint*. To simplify the task of the programmer, who must ensure an accessed element is held locally, the languages will typically add *distributed control* constructs. These play a role something like the `ON HOME` directives of HPF 2.0 and earlier data parallel languages [29]. A further special control construct will facilitate access to all elements in the locally held section of a particular array (or group of aligned arrays). This is the *distributed loop* or *overall construct*.

Java, Fortran and C++ versions. A Java instantiation (HPJava) of the HPspmd language model outlined above has been described in [9, 10]. A brief review is given in section C.4.1. HPJava is a superset of the Java language that adds predefined classes and some additional syntax for dealing with distributed arrays. It also adds three new control constructs, including the *overall* distributed loop, which is used to traverse local elements of distributed arrays.

In [7] we have outlined possible syntax extensions to Fortran to provide similar semantics to HPJava. As emphasized previously, a distinguishing property of the proposed system, compared to HPF, is that it includes ordinary Fortran as a strict subset, *and ordinary Fortran constructs are unchanged by the translator*. The proposed system would *not*

attempt to exploit parallelism even in constructs such as the array syntax of Fortran 90 or the FORALL statement of Fortran 95, because those constructs operate on the standard sequential arrays of the language. This policy drastically simplifies the translator, and gives the programmer much finer control over the generated code.

So far as C++ is concerned, a working prototype of our language model exists in the form of the ad++ interface to Adlib [5, 12]. This extends C++ only by class libraries and macros. In C++ we can use features like operator-overloading, templates, reference-valued functions, and macros to effectively prototype new language constructs. *But* the current ad++ is very inefficient (and the concrete syntax is quite clumsy) compared with what could be achieved with a purpose-built compiler or preprocessor.

In the proposed work, research into optimizing compilers and preprocessor for HPspmd versions of Fortran and C++ will be led by our collaborater Professor Xiaoming Li from Peking University.

General translation issues. The language extensions described earlier were devised partly to provide a convenient interface to a distributed-array library developed in the Parallel Compiler Runtime Consortium (PCRC) project [14].

Compared with HPF, translation of the HPspmd languages is very straightforward. The HPJava compiler, for example, is being implemented initially as a translator to ordinary Java, through a compiler construction framework developed in the PCRC project. The distributed arrays of the extended language appear in the emitted code as a pair—an ordinary Java array of local elements and a Distributed Array Descriptor object (DAD). In the initial implementation, details of the distribution format, including non-trivial details of global-to-local translation of the subscripts, are managed in the runtime library. Even with these overheads, acceptable performance is achievable, because in useful parallel algorithms most work on distributed arrays occurs inside *overall* constructs with large ranges. In normal usage, the formulae for address translation can be linearized inside these constructs, and the cost of runtime calls handling non-trivial aspects of address translation (including array bounds checking) can be amortized in the startup overheads of the loop. These compiler optimizations will be important in the base level translator. If array accesses are genuinely irregular, the necessary subscripting cannot usually be *directly* expressed in our language; subscripts cannot be computed randomly in parallel loops without violating the SPMD restriction that accesses be local. This is not necessarily a shortcoming: it forces explicit use of an appropriate library package for handling irregular accesses (such as CHAOS, see section C.2.2).

The basic HPJava translator will be available by the start date of the proposed work. In figure C.1 we give benchmark results for HPJava examples manually converted to Java, following the translation scheme outlined above. The examples are essentially the ones described in section C.4.1. The parallel programs are executed on 4 sparc-sun-solaris2.5.1 using MPICH and the Java JIT compiler in JDK 1.2Beta2, through a JNI interface to Adlib for collective communications. In both cases arrays are 1024 by 1024. For Jacobi iteration, the timing is for about 90 iterations. Timings are compared with *sequential* Java and C++ versions of the code (horizontal lines). Note that poor scaling in the Cholesky case is attributable to the poor performance of MPICH on this platform *not* overheads of HPJava. Scaling will be much improved by using SunHPC MPI.

The single-processor HPJava performance is better than sequential Java, because the pure Java version was coded in the natural way, using two-dimensional arrays—quite inefficient in Java. The HPJava translation scheme linearizes arrays. (We remark that in recent workshops James Gosling has stated that this is his preferred approach to adding generalized array-like structure in Java.) Although absolute performance is still somewhat lower than C++, Java performance has improved dramatically over the last year, and we expect to see further gains. Parity between Java and C or Fortran no longer seems an unrealistic expectation. In fact, even if the performance of Java does not rapidly approach that of C and Fortran, Java remains an excellent research platform for the general language model we espouse. It combines strong support for dynamic and object-oriented programming in a *relatively simple* language, for which preprocessors for extended versions of the language ("little languages") are a straightforward proposition.

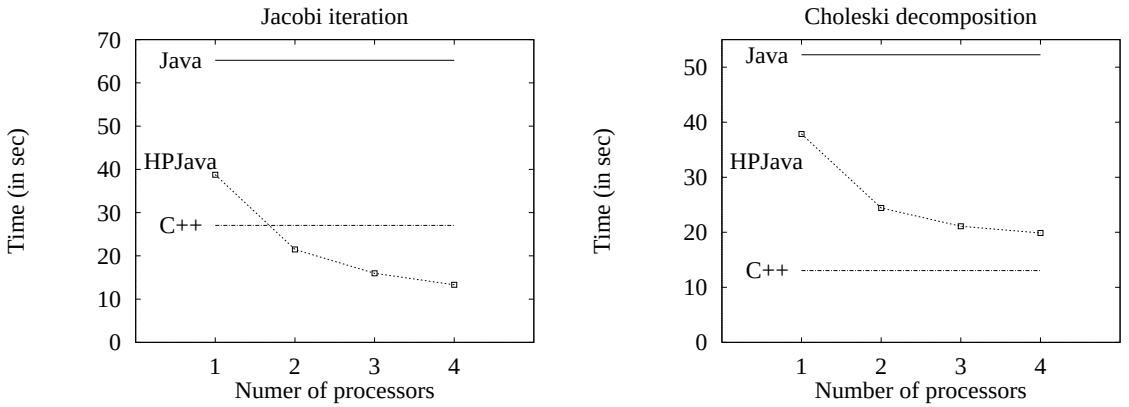

Figure C.1: Preliminary HPJava performance

C.2.2 Integration of high-level libraries, regular and irregular

Libraries are at the heart of the HPspmd model. From one point of view, the language extensions are simply a framework for invoking libraries that operate on distributed arrays. The base language model was originally motivated by work on HPF runtime libraries carried out in the Parallel Compiler Runtime Consortium (PCRC) project [14] led by Syracuse (and earlier related work by one of us [12]).

Hence an essential component of the proposed work is to define a series of bindings from our languages to established SPMD libraries and environments. Because our language model is explicitly SPMD, such bindings are a more straightforward proposition than in HPF, where one typically has to pass some extrinsic interface barrier before invoking SPMD-style functions.

Various issues must be addressed in interfacing to multiple libaries. For example, low-level communication or scheduling mechanisms used by the different libraries may be incompatible. As a practical matter these incompatibilities must be addressed, but the main thrust of the proposed research is at the level of *designing* compatible interfaces, rather than solving interference problems in specific implementations.

We will group the existing SPMD libraries for data parallel programming into three

classes, loosely based on the complexity of design issues involved in integrating them into our language framework.

In the first class we have libraries like ScaLAPACK [4] and PetSc [2] where the primary focus is similar to conventional numerical libraries—providing implementations of standard matrix algorithms, say, but operating on elements in regularly distributed arrays. We believe that designing HPspmd interfaces to this kind of package will be relatively straightforward

ScaLAPACK for example, provides linear algebra routines for distributed-memory computers. These routines operate on distributed arrays—in particular, distributed matrices. The distribution formats supported are restricted to two-dimensional block-cyclic distribution for dense matrices and one-dimensional block distribution for narrow-band matrices. Since both these distribution formats are supported by HPspmd (it supports all HPF-compatible distribution formats), using ScaLAPACK routines from the HPspmd framework should present no fundamental difficulties. Problems can only arise if the caller attempts to pass in matrix with a distribution format unsupported by the ScaLAPACK routines. The interface code between HPspmd and ScaLAPACK (which converts between array descriptors) must either flag a runtime error in this case, or remap the argument array (using, for example, the `remap` primitive of Adlib [11]).

In the second class we place libraries conceived primarily as underlying support for general parallel programs with regular distributed arrays. They emphasize high-level communication primitives for particular styles of programming, rather than specific numerical algorithms. These libraries include rutimes libraries for HPF-like languages, such as Adlib and Multiblock Parti [1], and the Global Array toolkit [34].

Adlib is a runtime library was initially designed to support HPF translation. It provides communication primitives similar to Multiblock PARTI, plus all Fortran 90 transformational intrinsics for arithmetic on distributed arrays. It also provides some gather/scatter operations for irregular access.

The array descriptor of Adlib supports the full HPF 1.0 distributed array model—including all standard distribution formats, all alignment options including replicated alignment, and a facility to map an array to an arbitrary subgroup of the set of active processors. The runtime array descriptor of the HPspmd languages will be an enhanced version of the Adlib descriptor (with a few extra features, such as support for the `GENBLOCK` distribution format of HPF 2.0 [20]). The Adlib collective communication library will provide initial library support for regular applications in HPspmd.

The Global Array (GA) toolkit, developed at Pacific Northwest National Lab, provides an efficient and portable "shared-memory" programming interface for distributed-memory computers. Each process in a MIMD parallel program can asynchronously access logical blocks of distributed arrays, without need for explicit cooperation by other processes ("one-sided communication"). This model has been popular and successful. GA is a foundation of the NWChem [3, 28] computational chemistry package.

The existing interface to Global Arrays only supports two-dimensional arrays with general block distribution format. Distributed arrays are created by calls to Fortran functions which return integer handles to an array descriptor. The authors of the package are currently investigating generalization to support multi-dimensional arrays, with more general distribution format. They have already expressed interest in making their library accessible

through the kind of language extensions for distributed arrays described in this proposal.

Besides providing a much more tractable interface for creation of multidimensional distributed arrays, our syntax extensions will provide a more convenient interface to primitives such as `ga_get`, which copies a patch of a global array to a local array. Advantages over the existing API include the fact is that the interface can be made uniform for all ranks of arrays, and various sorts of checking can subsumed by the general mechanisms for array section creation, leading to improved safety and compile-time analysis.

Regular problems (such as the linear algebra examples in section C.4.1) previous section) are an important subset of parallel applications, but of course they are far from exclusive. Many important problems involve data structures too irregular to express purely through HPF-style distributed arrays.

Our third class of libraries therefore includes libraries designed to support irregular problems. These include CHAOS [16, 37] and DAGH [36].

We anticipate that irregular problems will still benefit from regular data-parallel language extensions (because, at some level they usually resort to representations involving regular arrays). But lower level SPMD programming, facilitated by specialized class libraries, are likely to take a more dominant role when dealing with irregular problems.

The CHAOS/PARTI runtime support library provides primitives for efficiently handling *irregular* problems on distributed memory computers. The complete library includes partitioners to choose optimized mapping on arrays to processors, functions to remap input arrays to meet the optimized partitioning, and functions which optimize interprocessor communications. After data is repartitioned (if necessary) CHAOS programs involve two characteristic phases. The *inspector* phase analyses data access patterns in the main loop, and generates a schedule of optimized optimized communication calls. The *executor* phase involves executing a loop essentially similar to the loop of the original sequential program.

How best to capture this complexity in a convenient HPspmd interface will be a subject of research in the proposed work. A baseline approach (in HPJava, for example) is to handle the translation tables, schedules, etc of CHAOS as ordinary Java objects, constructed and accessed in explicit library calls. Presumbly the initial values for the data and indirection arrays will be provided as normal HPspmd distributed arrays. The simplest assumption is that the CHAOS preprocessing phases yield new arrays: the indirection arrays may well be left as HPspmd distributed arrays, but the data arrays may be reduced to ordinary Java arrays holding local elements (in low-level SPMD style). Then, with no extensions to the currently proposed HPJava language, the parallel loops of the executor phase can be expressed using *overall* constructs. More advanced schemes may incorporate irregular maps into generalized array descriptor [13, 17, 20]. Extensions to the HPspmd language model may be indicated.

DAGH (Distributed Adaptive Grid Hierarchy) was developed at Texas, Austin as a computational toolkit for several projects including the Binary Black Hole NSF Grand Challenge Project. It provides the framework to solve systems of partial differential equations using adaptive mesh refinement methods. The computations can be executed sequentially or in parallel according to the specification of the user. In the parallel case DAGH takes over communication, updating ghost regions on the boundaries of component grids.

Conceivably the HPspmd distributed array descriptor could be generalized to directly represent a DAGH grid hierarchy. This is probably unrealistic. DAGH implements a

non-trivial storage scheme for its grid hierarchy, based on space-filling curves. It seems unlikely that the details of such a structure can be sensibly handled by a compiler. A more straightforward possibility is to represent the individual grid functions (on the component regular meshes of the hierachy) as essentially standard HPspmd distributed arrays. Since DAGH is supposed to maintain storage for these functions in Fortran-compatible fashion, it should be practical to create an HPspmd array descriptor for them. The hierarchy itself would be represented as a Java object from a library-defined class. This is a crude outline of a particular scenario. Devising practical and convenient HPspmd bindings for DAGH and similar application-oriented libraries is a research topic in the proposed work.

C.2.3 Java MPI linkage

In HPF, with its global-thread-of-control model, a proper interface to the underlying message-passing platform is only practical through the *extrinsic procedure* mechanism. In HPspmd it is possible to access the MPI interface directly. In Fortran and C++ bindings of HPspmd probably the only major issue arising is access to the local elements of distributed arrays as standard sequential Fortran or C++ arrays, which can be passed to the standard MPI functions. Inquiry functions on distributed arrays return the sequential arrays as pointers or handles (depending on the language instantiation).

We have already implemented a Java language binding for MPI, version 1.1 [6, 8]. Our current approach is a relatively direct transcription of standard MPI bindings, but Java object serialization introduces new possibilities for passing compound objects. Similar projects on Java MPI bindings are in progress elsewhere [15, 27].

C.2.4 Integration of thread-based single Java VM and multi-VM data parallel

Our language model is primarily aimed at distributed memory computers, including networks of workstations or PCs. Clearly the Java version of HPspmd also holds special promise in the domain of metacomputing—targeting heterogeneous systems. At the other extreme, the same model can be straightforwardly implemented on symmetric multiprocessors—using threads within a single Java virtual machine. The most naive approach is to directly simulate the SPMD model in this environment with a fixed set of threads. Further possibilities arise if a few restrictions on variable usage are added to the language model. The main program can execute as a single thread, with multiple threads forked only when an *overall* construct is encountered. These issues will be investigated further.

C.3 General plan of work

Work at NPAC will initially focus on the Java binding of the HPspmd language model (HPJava). The basic HPJava translator will be available for further development and initial experiments with applications. This version of HPJava will rely heavily on run-time library functions for basic operations such as subscript translation (incorporating

only essential optimizations on distributed loops). Initially the only communication library available will be Adlib.

One thread in the proposed work will be to produce an optimized version of the initial HPJava translator. For example, static information will be exploited to inline and simplify calls to the runtime library wherever possible. Runtime checks on multi-dimensional array-bound violations and adherence to the "SPMD constraint" (requiring that accesses be local) will be eliminated where possible. Ultimately it would be desirable to produce a true compiler (rather than source-to-source translator) for HPJava. This will not be a primary goal in the proposed work, which emphasizes rapid implementation of, and experimentation with, novel language ideas, driven by application and library requirements.

A second major thread will be design and limited implementation of HPspmd interfaces to libraries described in section C.2.2 (and MPI). This work will be coupled with the development of suitable demonstrator applications that exploit the libraries. Initial examples will be taken from the HPFA kernels maintained at NPAC [38], converted to use the Adlib library. The proposal includes support for application scientists familiar with DAGH and the binary black hole problem, and GA and computational chemistry. Fast Multipole and its associated irregular MPI-based library for earthquake problems is another area of current interest at NPAC.

This application work, and in particular the requirements of the library bindings, is expected to drive the third thread: further development of the base HPspmd language model, especially in regard of supporting "irregular" problems.

The fourth major thread will involve taking the HPspmd ideas and embedding them in more conventional scientific programming languages: Fortran and C++. The main design and implementation work here will be carried out by our collaborators from the University of Peking, led by Professor Xiaoming Li. Professor Li has collaborated closely with NPAC over several years, and worked at NPAC for two years during the PCRC project, leading our HPF compiler effort.

C.3.1 Three year workplan

Year one: In the first year we will be studying and implementing optimizations in the basic HPJava translator. Our early experiments give us confidence that the basic HPJava translation scheme can give good performance, with minimal overheads, for problems involving large arrays. But there is considerable scope for improving performance on smaller problem sets, especially in reducing run-time overheads associated with subscript conversion. Also in this year an interface will be made between HPJava and the Global Arrays toolkit, at least for some set of platforms. Application efforts will concentrate on a finite difference problem derived from the theory of Black Holes. Study of requirements for irregular problems will be an important activity.

Year two: The requirements identified in the first year's activity will feed into implementation of class library interfaces for irregular problems. This will include CHAOS-like support for irregular access to arrays, and DAGH-like support for adaptive meshes. Fast multipoles will be one focussed example of a more complex problem tackled as an application in this phase. We will look at some representative computation chemistry problems

using the GA binding developed in year one. Any extensions to the basic HPspmd language model indicated by experiences with irregular applications and libraries will be implemented. Work on optimization and compile-time checking of the translator will continue, to produce the robust system needed by the application programmers.

Year three: Emphasis will be on integration. By this stage we will have bindings to several libraries operating on various platforms. We also expect to have compilers for Fortran and C++ versions of the HPspmd languages, developed by our collaborators in China. In so far as practical we must ensure that bindings to different libraries interoperate without interference, and document any problems. The software developed in the project will be placed in the public domain.

C.3.2 Collaborations

As explained above the project involves an important collaboration with Peking University. This will require mutual visits and continuation of ongoing electronic collaboration. NPAC already have substantial sharing of software with the Peking group, exemplified by our HPF front end [32] and the `f2j` Fortran to Java translator [22], where the software was built in China but used in NPAC activities, who provided design expertise.

Some input into this project is expected from work supported by Sun Microsystems. They are providing funding for a project led by NPAC to investigate Java for large scale computing. This work will support students at Syracuse, Indiana and Illinois. It will look at Java for NCSA Alliance Grand Challenges.

C.4 Related work

C.4.1 Applicant's related work

HPJava. HPJava [9, 10] is an instance of the HPsmpd language model. HPJava extends the base Java language by adding predefined classes and some additional syntax for dealing with distributed arrays, and three new control constructs.

As explained in the previous section, the underlying distributed array model is equivalent to the HPF array model. As a matter of detail, distributed array mapping is described in terms of a slightly different set of basic concepts. HPF describes the decomposition of an array through alignment to some *template*, which is in turn distributed over a *processor arrangement*. The analogous concepts in our parametrization of the distributed array are the *distributed range* (or simply *range*) and the *process group* (or simply *group*). A distributed range is akin a single dimension of an HPF template—it defines a map from an integer global subscript range into a particular dimension of a process group. A *process group* is equivalent to an HPF processor arrangement, or to a certain subset of such an arrangement. Switching from templates to ranges and groups is a change of parametrization only. In itself it does not change the set of allowed ways to decompose an array. The new primitives fit better with our distributed control constructs, and correspond more directly to components of our run-time array descriptor. Ranges and groups are treated as proper objects in the extended language. They are values that can be stored in variables

```
Procs1 p = new Procs1(NP) ;
on(p) {
  Range x = new CyclicRange(N, p.dim(0));

  float [[,#]] a = new float [[N, x]] ;
  float [[]]   b = new float [[N]] ;  // buffer

  Location l ;
  Index m ;

  for(int k = 0 ; k < N - 1 ; k++) {

    at(l = x [k]) {
      float d =  Math.sqrt(a [k, l]) ;

      a [k, l] = d ;
      for(int s = k + 1 ; s < N ; s++)
        a [s, l] /= d ;
    }

    Adlib.remap(b [[k + 1 : ]], a [[k + 1 : , k]]);

    over(m = x | k + 1 : )
      for(int i = x.idx(m) ; i < N ; i++)
        a [i, m] -= b [i] * b [x.idx(m)] ;
  }

  at(l = x [N - 1])
    a [N - 1, l] = Math.sqrt(a [N - 1, l]) ;
}
```

Figure C.2: Choleski decomposition.

or passed to procedures. The group and ranges describing a particular distributed array are accessible through inquiry functions.

To motivate the discussion of HPJava, we will refer to figure C.2, which gives a parallel implementation of Choleski decomposition in the extended language. In pseudocode, the sequential algorithm is

$$
\begin{aligned}
&For\ k = 1\ to\ n-1\\
&\quad l_{kk} = a_{kk}^{1/2}\\
&\quad For\ s = k+1\ to\ n\\
&\qquad l_{sk} = a_{sk}/l_{kk}\\
&\quad For\ j = k+1\ to\ n\\
&\qquad For\ i = j\ to\ n\\
&\qquad\quad a_{ij} = a_{ij} - l_{ik}l_{jk}\\
&l_{nn} = a_{nn}^{1/2}
\end{aligned}
$$

The parallel version has been selected to introduce essentially all the new language extensions in HPJava.

In HPJava a base class `Group` describes a general group of processes. It has subclasses

`Procs1`, `Procs2`, ..., that represent one-dimensional process grids, two-dimensional process grids, and so on. In the example `p` is defined as a one-dimensional grid of extent `NP`. The *on* construct in the example acts like a conditional, excluding processors outside the group `p`. A *distributed range*, base class `Range`, defines a range of integer global subscripts, and specifies how they are mapped into a process grid dimension. In the example, the range `x` is initalized to a cyclically distributed range of extent `N`. `CyclicRange` is one of several subclasses of `Range` that define different *distribution formats*.

Now `a` and `b` are declared to be *distributed arrays*. In HPJava the type-signatures and constructors of distributed arrays use double brackets to distinguish them from ordinary Java arrays. If a particular dimension of an array has a distributed range, the corresponding slot in the type signature of the array should include a `#` symbol. Because `b` has no range distributed over the active process group (`p`) it is defined to be *replicated* across this group. The mapping of `a` and `b` is equivalent to the HPF declarations

```
!HPF$ PROCESSORS p(np)

!HPF$ TEMPLATE t(n)
!HPF$ DISTRIBUTE t(CYCLIC) ONTO p

      REAL a(n, n), b(n)
!HPF$ ALIGN a(i, *) WITH t(i)
!HPF$ ALIGN b(*)    WITH t(*)
```

with range `x` taking over the role of the one-dimensional template `t`.

Subscripting operations on distributed arrays are subject to a strict restriction. An access to an array element such as `a [s, k]` is legal, but *only* if the local process holds the element in question. The language provides syntax to alleviate the inconvenience of this restriction. The idea of a *location* is introduced. It can be viewed as an abstract element, or "slot", of a distributed range. Any location is mapped to a particular slice of a process grid. Locations are used to parametrize a new distributed control construct called the *at* construct. This works like *on*, except that its body is executed only on processes that hold the specified location. Locations can also be used directly as array subscripts, in place on integers (locations used as array subscripts must be elements of the corresponding ranges of the array). The array access above can be safely written in the context

```
Location l = x [k] ;
at(l)
  ... a [s, l] ...
```

(the first dimension of `a` is sequential, so we don't have to worry about the SPMD constraint for subscript `s`). In the main example, this syntax is used to ensure that the first block of code inside the loop only executes on the processor holding column `k`.

The example involves one communication operation. This is taken from the Adlib library: the function `remap` copies the elements of one distributed array or section to another of the same shape. The two arrays can have any, unrelated decompositions. Because `b` has replicated mapping, `remap` copies identical values to all processors—ie it implements a broadcast of the values in the array section `a [[k + 1 : , k]]`. The syntax for array sections in HPJava is almost identical to the syntax of sections in Fortran 90. Subscript triplets work in the same way as in Fortran 90.

```
Procs2 p = new Procs2(NP, NP) ;

on(p) {
  Range x = new BlockRange(N, p.dim(0), 1) ;  // ghost width 1
  Range y = new BlockRange(N, p.dim(1), 1) ;  // ghost width 1

  float [[#,#]] u = new float [[x, y]] ;

  int [] widths = {1, 1} ;        // Widths updated by 'writeHalo'

  // ... some code to initialise 'u'

  for(int iter = 0 ; iter < NITER ; iter++) {
    for(int parity = 0 ; parity < 2 ; parity++) {

      Adlib.writeHalo(u, widths) ;

      Index i, j ;
      over(i = x | 1 : N - 2)
        over(j = y | 1 + (x.idx(i) + parity) % 2 : N - 2 : 2)
          u [i, j] = 0.25 * (u [i - 1, j] + u [i + 1, j] +
                             u [i, j - 1] + u [i, j + 1]) ;
    }
  }
}
```

Figure C.3: Red-black iteration.

The last and most important distributed control construct in the language is called *over*. It is used to access all locally held locations in a particular range, and can therefore be used to access all locally held elements of arrays parametrized by that range. The *over* construct implements a distributed parallel loop. Its parameter is a member of the special class `Index` which is a subclass of `Location`. The `idx` member of `Range` can be used inside parallel loops to yield arithmetic expressions that depend on global index values. In the example the *over* construct is used to iterate over all columns of the matrix to the right of column `k`.

As promised, the Choleski example has introduced essentially all the important language ideas in HPJava. Further extensions are minor, or consist in adding new subclasses of `Range` or `Group`, rather than syntax extensions. Figure C.3 gives a parallel implementation of red-black relaxation in the same language. To support this important stencil-update paradigm, *ghost regions* are allowed on distributed arrays [26]. In our case the width of these regions is specified in a special form of the `BlockRange` constructor. The ghost regions are explicitly brought up to date using the library function `writeHalo`.

Note that the new range constructor and `writeHalo` function are *library* features (respectively from the base HPJava runtime and the Adlib communication library), *not* new language extensions. One new piece of syntax is involved: the addition and subtraction operators are overloaded so that integer offsets can be added or subtracted to locations, yielding new, shifted, locations. This kind of shifted access only works if the subscripted array has suitable ghost extensions.

Adlib. The Adlib runtime library was initially designed to support HPF translation. Early development took place in the *shpf* [33] project at Southampton, UK. Subsequently the library was redesigned and reimplemented at Syracuse during in the PCRC project, and delivered as the NPAC PCRC runtime kernel [11]. It has been used as a foundation of two experimental HPF compilation systems [33, 42], (one in Europe and one at Syracuse), and is currently being used as a basis of the HPJava translator.

The Adlib kernel is C++ class library, built on MPI. Fortran, C++ and Java interfaces are available or under development. It provides communication primitives similar to Multiblock PARTI, plus the Fortran 90 transformational intrinsics for arithmetic on distributed arrays. It also provides some collective gather/scatter operations for irregular access. Benchmarks reported in [42] suggested Adlib provides superior performance to the then-current version of the commercial PGI HPF compiler.

The array descriptor of Adlib supports the full HPF 1.0 distributed array model—including all standard distribution formats, all alignment options including replicated alignment. The runtime array descriptor of the HPspmd languages will be an enhanced version of the Adlib descriptor. The Adlib collective communication library will provide initial library support for regular applications in HPspmd.

C.4.2 Related languages

F-- [35] is an extended Fortran dialect for SPMD programming. The approach is quite different to the one proposed here. In F--, array subscripting is local by default, or involves a combination of local subscripts and explicit process ids. There is no analogue of global subscripts, or HPF-like distribution formats. In F-- the logical model of communication is built into the language—remote memory access with intrinsics for synchronization—where we follow the philosophy of providing communication through separate libraries. While F-- and our approach share an underlying programming model, we believe that our framework offers greater opportunities for exploiting established library technologies.

Spar [40] is a Java-based language for array-parallel programming. Like our language it introduces multi-dimensional arrays, array sections, and a parallel loop. There are some similarities in syntax, but semantically Spar is very different to HPJava. Spar expresses parallelism but not explicit data placement or communication—in this sense it is a higher level language—closer to HPF.

ZPL [39] is a new programming language for scientific computations. Like Spar, it is an array language. It has an idea of performing computations over a *region*, or set of indices. Within a compound statement prefixed by a *region specifier*, aligned elements of arrays distributed over the same region can be accessed. This idea has certain similarities to our *overall* construct. Communication is more explicit than, say, Spar, but not as explicit as in the language discussed in this article.

Titanium [41] is another Java-based language for high-performance computing. It provides multi-dimensional arrays and a global address space, with an SPMD programming model. It does not provide any special support for distributed arrays, and the programming style is quite different to HPJava.

D. Bibliography

[1] A. Agrawal, A. Sussman, and J. Saltz. An integrated runtime and compile-time approach for parallelizing structured and block structured applications. *IEEE Transactions on Parallel and Distributed Systems*, 6, 1995.

[2] S. Balay, W. D. Gropp, L. C. McInnes, , and B. F. Smith. Efficient management of parallelism in object-oriented numerical software libraries. In E. Arge, A. M. Bruaset, and H. P. Langtangen, editors, *Modern Software Tools in Scientific Computing*, pages 163–202. Birkhauser Press, 1997.

[3] D. E. Bernholdt, E. Aprà, H. A. Früchtl, M. F. Guest, R. J. Harrison, R. A. Kendall, R. A. Kutteh, X. Long, J. B. Nicholas, J. A. Nichols, H. L. Taylor, A. T. Wong, G. I. Fann, R. J. Littlefield, and J. Nieplocha. Parallel computational chemistry made easier: The development of NWChem. *Int. J. Quantum Chemistry: Quantum Chem. Symposium*, 29:475–483, 1995.

[4] L. S. Blackford, J. Choi, A. Cleary, E. D'Azevedo, J. Demmel, I. Dhillon, J. Dongarra, S. Hammarling, G. Henry, A. Petitet, K. Stanley, D. Walker, and R. C. Whaley. *ScaLAPACK User's Guide*. SIAM, 1997.

[5] Bryan Carpenter. *Programming in ad++*, 1998. http://www.npac.syr.edu/projects/pcrc/doc.

[6] Bryan Carpenter, Yuh-Jye Chang, Geoffrey Fox, Donald Leskiw, and Xiaoming Li. Experiments with HPJava. *Concurrency: Practice and Experience*, 9(6):633, 1997.

[7] Bryan Carpenter, Geoffrey Fox, Donald Leskiw, Xinying Li, Yuhong Wen, and Guansong Zhang. Language bindings for a data-parallel runtime. In *Third International Workshop on High-Level Parallel Programming Models and Supportive Environments*, 1998. To appear. Also available at http://www.npac.syr.edu/projects/pcrc/doc.

[8] Bryan Carpenter, Geoffrey Fox, Xinying Li, and Guansong Zhang. A draft Java binding for MPI. http://www.npac.syr.edu/projects/pcrc/doc.

[9] Bryan Carpenter, Guansong Zhang, Geoffrey Fox, Xinying Li, and Yuhong Wen. Introduction to Java-Ad. http://www.npac.syr.edu/projects/pcrc/doc, November 1997.

[10] Bryan Carpenter, Guansong Zhang, Geoffrey Fox, Xinying Li, and Yuhong Wen. HPJava: Data parallel extensions to Java. In *ACM workshop on Java for High-performance Network Computing*, 1998. To appear in Concurrency: Practice and Experience. Extended version available at http://www.npac.syr.edu/projects/pcrc/doc.

[11] Bryan Carpenter, Guansong Zhang, and Yuhong Wen. NPAC PCRC runtime kernel definition. Technical Report CRPC-TR97726, Center for Research on Parallel Computation, 1997. Up-to-date version maintained at http://www.npac.syr.edu/projects/pcrc/doc.

[12] D. B. Carpenter. Adlib: A distributed array library to support HPF translation, 1995. 5th International Workshop on Compilers for Parallel Computers. http://www.npac.syr.edu/users/dbc/Adlib.

[13] B. Chapman, P. Mehrotra, and H. Zima. Programming in Vienna Fortran. *Scientific Programming*, 1(1):1–50, 1992.

[14] Parallel Compiler Runtime Consortium. Common runtime support for high-performance parallel languages. In *Supercomputing '93*. IEEE Computer Society Press, 1993.

[15] George Crawford III, Yoginder Dandass, and Anthony Skjellum. The jmpi commercial message passing environment and specification: Requirements, design, motivations, strategies, and target users, 1997. http://www.mpi-softtech.com/publications/.

[16] R. Das, M. Uysal, J.H. Salz, and Y.-S. Hwang. Communication optimizations for irregular scientific computations on distributed memory architectures. *Journal of Parallel and Distributed Computing*, 22(3):462–479, September 1994.

[17] G. Fox et al. Fortran D language specification. Technical Report CRPC-TR90079, Center for Research on Parallel Computation, Rice University, 1990.

[18] Stephen J. Fink and Scott B. Baden. Run-time data distribution for block-structured applications on distributed memory computers. In *Proceedings of the 7th SIAM Conference on Parallel Processing for Scientific Computing*, February 1995.

[19] High Performance Fortran Forum. High Performance Fortran language specification. *Scientific Programming*, special issue, 2, 1993.

[20] High Performance Fortran Forum. High Performance Fortran language specification, version 2.0, January 1997. http://www.crpc.rice.edu/HPFF/hpf2.

[21] Message Passing Interface Forum. *MPI: A Message-Passing Interface Standard.* University of Tenessee, Knoxville, TN, June 1995. http://www.mcs.anl.gov/mpi.

[22] Geoffrey Fox, Xiaoming Li, and Zheng Qiang. A prototype of Fortran-to-Java converter. *Concurrency: Practice and Experience*, 9(11):1047, 1997.

[23] Geoffrey C. Fox, editor. *ACM 1998 Workshop on Java for High-Performance Network Computing*, Concurrency: Practice and Experience (to appear). Palo Alto, California, February 28 and March 1, 1998. http://www.cs.ucsb.edu/conferences/java98.

[24] Geoffrey C. Fox, editor. *Java for Computational Science and Engineering—Simulation and Modelling*, volume 9(6) of *Concurrency: Practice and Experience*, June 1997.

[25] Geoffrey C. Fox, editor. *Java for Computational Science and Engineering—Simulation and Modelling II*, volume 9(11) of *Concurrency: Practice and Experience*, November 1997.

[26] Michael Gerndt. Updating distributed variables in local computations. *Concurrency: Practice and Experience*, 2(3):171–193, 1990.

[27] Vladimir Getov, Susan Flynn-Hummel, and Sava Mintchev. High-performance parallel programming in java: Exploiting native libraries. In *ACM workshop on Java for High-performance Network Computing*, 1998. To appear in Concurrency: Practice and Experience.

[28] Robert J. Harrison, Martyn F. Guest, Rick A. Kendall, David E. Bernholdt, Adrian T. Wong, Mark Stave, James Anchell, Anthony Hess, Rik Littlefield, George I. Fann, Jarek Nieplocha, Greg S. Thomas, David Elwood, Jeff Tilson, Ron L. Shepard, Albert F. Wagner, Ian T. Foster, Ewing Lusk, and Rick Stevens. High performance computational chemistry. II. A scalable SCF program. *J. Chem. Phys.*, 17:124, 1995.

[29] C. Koelbel and P. Mehrotra. Compiling global name-space parallel loops for distributed execution. *IEEE Transactions on Parallel and Distributed Systems*, 2(4):440–451, 1991.

[30] C.H. Koelbel, D.B. Loveman, R.S. Schreiber, G.L. Steel, Jr., and M.E. Zosel. *The High Performance Fortran Handbook.* MIT Press, 1994. ISBN: 0-262-61094-9.

[31] Scott R. Kohn and Scott B. Baden. A robust parallel programming model for dynamic non-uniform scientific computations. In *Proceedings of the 1994 Scalable High Performance Computing Conference*, March 1994.

[32] Xiaoming Li. HPFfe: a front-ed for HPF. Technical Report SCCS-771, Northeast Parallel Architectures Center, Syracuse University, October 1996. http://www.npac.syr.edu/projects/pcrc/doc.

[33] John Merlin, Bryan Carpenter, and Tony Hey. shpf: a subset High Performance Fortran compilation system. *Fortran Journal*, pages 2–6, March 1996.

[34] J. Nieplocha, R.J. Harrison, and R.J. Littlefield. The Global Array: Non-uniform-memory-access programming model for high-performance computers. *The Journal of Supercomputing*, 10:197–220, 1996.

[35] R.W. Numrich and J.L. Steidel. F- -: A simple parallel extension to Fortran 90. *SIAM News*, page 30, 1997.

[36] Manish Parashar and J.C. Browne. Systems engineering for high performance computing software: The HDDA/DAGH infrastructure for implementation of parallel structured adaptive mesh. In *Structured Adaptive Mesh Refinement Grid Methods*, IMA Volumes in Mathematics and its Applications. Springer-Verlag.

[37] Ravi Ponnusamy, Yuan-Shin Hwang, Raja Das, Joel H. Saltz, Alok Choudhary, and Geoffrey Fox. Supporting irregular distributions using data-parallel languages. *IEEE Parallel and Distributed Technology*, Spring, 1995.

[38] Sanjay Ranka, Hon W Yau, Kenneth A Hawick, and Geoffrey C Fox. High-Performance Fortran for SPMD programming: An applications overview, 1996. URL: http://www.npac.syr.edu/hpfa/Papers/HPFforSPMD/.

[39] Lawrence Snyder. A ZPL programming guide. Technical report, University of Washington, May 1997. URL: http://www.cs.washington.edu/research/projects/zpl/.

[40] Kees van Reeuwijk, Arjan J. C. van Gemund, and Henk J. Sips. Spar: A programming language for semi-automatic compilation of parallel programs. *Concurrency: Practice and Experience*, 9(11):1193–1205, 1997.

[41] Kathy Yelick, Luigi Semenzato, Geoff Pike, Carleton Miyamoto, Ben Liblit, Arvind Krishnamurthy, Paul Hilfinger, Susan Graham, David Gay, Phil Colella, and Alex Aiken. Titanium: A high-performance java dialect. In *ACM workshop on Java for High-performance Network Computing*, 1998. To appear in Concurrency: Practice and Experience.

[42] Guansong Zhang, Bryan Carpenter, Geoffrey Fox, Xiaoming Li, Xinying Li, and Yuhong Wen. PCRC-based HPF compilation. In *10th International Workshop on Languages and Compilers for Parallel Computing*, 1997. To appear in Lecture Notes in Computer Science.

E. Biographical Sketches

Geoffrey Charles Fox

Address: 111 College Place, NPAC, SU, 13244 gcf@nova.npac.syr.edu , http://www.npac.syr.edu, Phone: (315) 443-2163, Fax: (315) 443-4741
Citizen Status: Permanent Resident Alien; Citizen of United Kingdom

Education: B.A. in Mathematics from Cambridge Univ., Cambridge, England (1961-1964) Ph.D. in Theoretical Physics from Cambridge University (1964-1967) M.A. from Cambridge University (1968)

Professional Experience:

1990- Professor of Computer Science, Syracuse University
1990- Professor of Physics, Syracuse University
1990- Director of Northeast Parallel Architectures Center
1979-1990 Professor of Physics, California Inst. of Tech.
1986-1988 Associate Provost for Computing, California Inst. of Tech.
1983-1985 Dean for Educational Computing, California Inst. of Tech.
1981-1983 Executive Officer of Physics, California Inst. of Tech.
1974-1979 Associate Professor of Physics, California Inst. of Tech.
1971-1974 Assistant Professor of Physics, California Inst. of Tech.
1970-1971 Millikan Research Fellow in Theoretical Physics, Caltech
1970 Visiting Scientist (April-May), Brookhaven National Laboratory
1969-1970 Research Fellow at Peterhouse College, Cavendish Lab.,Cambridge
1968-1969 Research Scientist, Lawrence Berkeley Lab., Berkeley, Calif.
1967-1968 Member of School of Natural Science, Inst. for Advanced Study, Princeton, New Jersey

Awards and Honors: Senior Wrangler, Part III Mathematics, Cambridge (1964) Alfred P. Sloan Foundation Fellowship (1973-75) Fellow of the American Physical Society (1990)

Journal Editorships:

Principal: Concurrency: Practice and Experience (John Wiley, Inc.) Physics and Computers (International Journal of Modern Physics C - World Scientific)

Associate: Journal of Supercomputing,

Selected List of Publications:

[1] Fox, G.C., Johnson, M.A., Lyzenga, G.A., Otto, S.W., Salmon, J.K., Walker, D.W., Solving Problems on Concurrent Processors, Vol. 1, Prentice-Hall, Inc. 1988; Vol. 2, 1990.

[2] Fox, G.C., Copty, N.,Ranka, S.,Shankar, R. "Solving the region growing problem on the Connection Machine," in Proceedings of the 22nd International Conference on Parallel Processing, volume 3, pages 102-105, 1993.

[3] Fox, G. C., Messina, P., Williams, R., Parallel Computing Works!, Morgan Kaufmann, San Mateo Ca, 1994.
[4] Fox G.C., Mills K., "InfoMall: an Innovative strategy for high-performance computing and communications application development", Internet Research, 4:31- 45, 1994.
[5] Fox, G.C., Hiranadani, S., Kennedy, K., Koelbel, C., Kremer, U., Tseng, C.W., Wu, M.Y., "FortranD Language Specifications", Rice COMP TR90079, December 1990, Revised, April 1991.
[6] Fox, G. C. "Approaches to Physical Optimization," in Proceedings of 5th SIAM Conference on Parallel Processes for Scientific Computation, pp 153-162, March 25-27, 1991, Houston, TX, J. Dongarra, K. Kennedy, P. Messina, D. Sorensen, R. Voigt, editors, SIAM, 1992. C3P-959, CRPC-TR91124
[7] Fox G.C., Mansour N., "Parallel Physical Optimization Algorithms for allocating data to multicomputer nodes", Journal of Supercomputing, 8:53-80,1994.
[8] Fox, G. C. "Parallel Computing and Education," Daedalus, Journal of the American Academy of Arts and Sciences, Vol. 121, No. 1, pps 111-118, Winter 1992. C3P-958, CRPC-TR91123.
[9] Fox, G, Bozkus, Z., Choudhary, A., Haupt, T., and Ranka, S. ""A compilation approach for Fortran 90D/HPF compilers on distributed memory MIMD computers," in Proceedings of the Sixth Annual Workshop on Languages and Compilers for Parallel Computing. Lecture Notes in Computer Science, Springer-Verlag, pp. 200–215. U. Banerjee, D. Gelernter, A. Nicolau, and D. Padua (editors).
[10] Fox, G and Furmanski, W. "Computing on the Web – New Approaches to Parallel Processing – Petaop and Exaop Performance in the Year 2007", Submitted to IEEE Internet Computing, http://www.npac.syr.edu/users/gcf/petastuff/petaweb/

Summary of Interests: Fox is an internationally recognized expert in the use of parallel architectures and the development of concurrent software and algorithms. His activities include high performance Java and Fortran compilers and their runtime support. Fox has established a community activity to investigate value of Java in large scale networked computing. He is also a leading proponent for the development of computational science as an academic discipline and a scientific method. He has established at Syracuse University both graduate and undergraduate programs which cover both simulation and information technologies. All courses have been made available on the Web and his research includes HPCC technology to support education at both K-12 and University level. His research on parallel computing has focused on development and use of this technology to solve large scale computational problems with recent application foci including numerical relativity, earthquake prediction and financial modeling. Fox directs InfoMall, which is focused on accelerating the introduction of high speed communications and parallel computing into New York State industry and developing the corresponding software and systems industry. Much of this activity is in educational area where Fox is leading developments of new K-12 curricula material built using VRML, Java and other new technology. A recent set of activities center on Web collaboration technology and its application to synchronous distance education

Ph.D Advisor: Dr. Richard Eden Cambridge University, England

David Bryan Carpenter

Northeast Parallel Architectures Center	Phone: 315 443 5068
Syracuse University	Fax: 315 443 1973
111 College Place	E-mail: dbc@npac.syr.edu
Syracuse, NY 13244-4100	URL: http://www.npac.syr.edu/users/dbc

Education

Ph D. in Physics, University of London (1979-1983).
B.S. in Physics, University of London (1976-1979).

Professional Experience

1996– Research Scientist, NPAC, Syracuse University **1994–1995** Programmer, High Performance Computing Centre, Southampton, UK. **1989–93** Research Fellow in Department of Electronics and Computer Science and Department of Physics, Southampton, UK. **1989** Employed at "Transputer Technology Solutions", Southampton, UK. **1985–1988** Research Fellow in Department of Physics, Southampton, UK. **1985** Royal Society Overseas Fellowship at DESY, Hamburg. **1983–1984** Research Fellow in Theoretical Physics Department, Edinburgh University.

Selected List of Publications:

[1] Bryan Carpenter, Geoffrey Fox, Donald Leskiw, Xinying Li, Yuhong Wen and Guansong Zhang "Language Bindings for a Data-parallel Runtime", To appear in proceedings Third International Workshop on High-Level Parallel Programming Models and Supportive Environments, 1998.

[2] Bryan Carpenter, Guansong Zhang, Geoffrey Fox, Xinying Li and Yuhong Wen "HP-Java: Data Parallel Extensions to Java", 1998 ACM workshop on Java for High-performance Network Computing, To appear in Concurrency: Practice and Experience.

[3] Bryan Carpenter, Guansong Zhang and Yuhong Wen "NPAC PCRC Runtime Kernel Definition", Center for Research on Parallel Computation, CRPC-TR97726, 1997.

[4] Bryan Carpenter, Yuh-Jye Chang, Geoffrey Fox and Xiaoming Li, "Java as a Language for Scientific Parallel Programming" 10th Int'l Workshop on Languages and Compilers for Parallel Computing (Aug, 1997). To appear in Lecture Notes in Computer Science.

[5] G. Zhang, B. Carpenter, G. Fox, X. Li, X. Li and Y. Wen "PCRC-based HPF Compilation" 10th Int'l Workshop on Languages and Compilers for Parallel Computing (Aug, 1997). To appear in Lecture Notes in Computer Science.

[6] B. Carpenter, Y.-J. Chang, G. Fox, D. Leskiw and X. Li, "Experiments with 'HP Java' ", Concurrency: Practice and Experience, Vol 9, num 9 (1997), p633.

[7] J. Merlin, B. Carpenter and Tony Hey, "shpf: a Subset High Performance Fortran compilation system," Fortran Journal, (1996), pp 2-6.

[8] D.B. Carpenter, "Adlib: A Distributed Array Library to Support HPF Translation" 5th Workshop on Compilers on Parallel Computer, Malaga (1995).

[9] D.B. Carpenter and H. Glaser, "Some Lattice-Based Scientific Problems, expressed in Haskell" University of Southampton, Department of Electronics and Computer Science, preprint CSTR 93-06 (1993). To appear, Journal of Functional Programming (1996).

[10] As. Abada, C.R. Allton, Ph. Boucard, D.B. Carpenter, M. Crisafulli, S. Güsken, P. Hernadez, V. Lubicz, G. Martinelli, O. Pène, C.T. Sachrajda, K. Schilling, G. Siegert and R. Sommer, "Semi-leptonic Decays of Heavy Flavours on a Fine-grained Lattice" Nucl.Phys. B416 (1994) p675.

Summary of Interests:

Carpenter has worked in parallel computing since 1985, when he was one of the first scientists to exploit the Inmos transputer for simulations of physical systems. Working at the University of Southampton, UK he became involved with design of libraries to support parallel computing. In a fruitful collaboration with John Merlin, then also at Southampton, he was codeveloper of the *shpf* system, an early implementation of subset HPF. This work produced the first implementation of the Adlib runtime library. After moving to Syracuse in 1996, Carpenter worked in the PCRC project. As part of this work the Adlib library was substantially redesigned to meet the requirements of a new HPF compiler, and reimplemented. Eventually it was delivered as the NPAC PCRC runtime kernel. Work on C++ interfaces to Adlib was a formative influence in the current work on HPJava and HPspmd.

Ph.D. Advisor

Elliot Leader

Guansong Zhang

Address:

NPAC,Syracuse University, 111 College Place, Syracuse, NY, 13244, U.S.A.

Education:

B.E. in Computer Science and Engineering from Harbin Institute of Technology, Harbin, China, (1986-1990).

Ph. D. in Computer Science from Harbin Institute of Technology, Harbin, China, (1986-1990).

Professional Experience:

1995-: Research scientist, NPAC, Syracuse University

1992-1995: Teaching Assistant, CS Dept., Harbin Institute of Technology

Awards and Honors:

Anchongen scholarship for outstanding postgraduate student (HIT, 1995)

Earlier Research Projects:

Data Parallel Processing Based on Networked Workstations (1994-1995, Grant from the Commission of National Defence Industry)

Multiprocessor Performance Evaluation Techniques (1991-1993, Grant from the Commission of National Defence Industry, China)

Distributed Computing Architecture (1989-1992, Grant from the Ministry of Aerospace, China)

Research Projects in progress:

Frontend system implementation and integration.

HPF compiler and PCRC run time library.

Development of the HPJava translator.

Selected List of Publications:

[1] Xiaoming Li, Guansong Zhang, et al, "HPFfe: a Front-end for HPF", Technical Report, SCCS-771, NPAC at Syracuse University, 1996.4.

[2] Guansong Zhang, "Partitioning Data Parallel Program for Workstation Cluster", Ph.D thesis, Harbin Institute of Technology., 1995

[3] Guansong Zhang, Xiaoming Li, "Barrier Synchronization and Pipeline Synchronization in Distributed Memory Machines", Science Bulletin, 1996

[4] Guansong Zhang, Xiaoming Li, "Predicting Execution Time of Parallel Program Based on Simulation Software", Chinese Journal on Computers., 1995

[5] Guansong Zhang, Xiaoming Li, "Analysis of Interconnection Functions Needed for a New Nonlinear Skewing Scheme", Chinese Journal on Computers., 1994

[6] Guansong Zhang, et al, "Software System for ABC-90jr., an Array Based Computer", IEEE TENCON'93, 1993

Ph.D Advisor:

Dr. Xiaoming Li, Computer Science Department, Peking University, Beijing 100871, People's Republic of China.

David E. Bernholdt

Northeast Parallel Architectures Center
Syracuse University
111 College Place
Syracuse, NY 13244-4100

Phone: 315 443 3857
Fax: 315 443 1973
E-mail: bernhold@npac.syr.edu

Education

Ph.D. in Chemistry, minors in Physics and Mathematics, and Certification in Chemical Physics, Univ. of Florida, Gainesville, FL, 1993.

B.S. in Chemistry *Cum Laude* with Highest Distinction, Univ. of Illinois, Urbana, IL, 1986

Professional Experience

1996– Research Assistant Professor, Department of Chemistry, Syracuse University, Syracuse, NY. **1995–** Alex G. Nason Fellow and Research Scientist, Northeast Parallel Architectures Center, Syracuse University, Syracuse, NY. **1995–** Affiliate Staff Scientist, Pacific Northwest National Laboratory, Richland, WA. **1993–1995** AWU Postdoctoral Fellow, High Performance Computational Chemistry Group, Environmental Molecular Sciences Laboratory, Pacific Northwest National Laboratory, Richland, WA. **1986–1993** Teaching or Research Assistant, University of Florida, Gainesville, FL. **1985–1986** Teaching Assistant, University of Illinois, Urbana, IL.

Related and Significant Publications

[1] J. Anchell, E. Apra, D. Bernholdt, P. Borowski, T. Clark, D. Clerc, H. Dachsel, M. Deegan, M. Dupuis, K. Dyall, G. Fann, H. Früchtl, M. Gutowski, R. Harrison, A. Hess, J. Jaffe, R. Kendall, R. Kobayashi, R. Kutteh, Z. Lin, R. Littlefield, X. Long, B. Meng, J. Nichols, J. Nieplocha, A. Rendell, M. Stave, T. Straatsma, H. Taylor, G. Thomas, K. Wolinski, and A. Wong. *NWChem, A Computational Chemistry Package for Parallel Computers, Version 3.1.* Pacific Northwest National Laboratory, Richland, Washington 99325-0999 USA, 1997.

[2] D. E. Bernholdt, E. Aprà, H. A. Früchtl, M. F. Guest, R. J. Harrison, R. A. Kendall, R. A. Kutteh, X. Long, J. B. Nicholas, J. A. Nichols, H. L. Taylor, A. T. Wong, G. I. Fann, R. J. Littlefield, and J. Nieplocha. Parallel computational chemistry made easier: The development of NWChem. *Int. J. Quantum Chemistry: Quantum Chem. Symposium*, 29:475–483, 1995.

[3] David E. Bernholdt and Robert J. Harrison. Orbital invariant second-order many-body perturbation theory on parallel computers: An approach for large molecules. *J. Chem. Phys.*, 102(24):9582–9589, 22 June 1995.

[4] David E. Bernholdt and Robert J. Harrison. Large-scale correlated electronic structure calculations: The RI-MP2 method on parallel computers. *Chem. Phys. Lett.*, 250:477–484, 8 March 1996.

[5] David E. Bernholdt and Robert J. Harrison. Fitting basis sets for the RI-MP2 approximate second-order many-body perturbation theory method. *J. Chem. Phys.*, submitted.

[6] David Feller, Edoardo Aprà, Jeff A. Nichols, and David E. Bernholdt. The structure and binding energy of K^+-ether complexes: A comparison of MP2, RI-MP2 and density functional methods. *J. Chem. Phys.*, 105(5):1940–1950, 1 August 1996.

[7] M. F. Guest, E. Aprà, D. E. Bernholdt, H. A. Früchtl, R. J. Harrison, R. A. Kendall, R. A. Kutteh, X. Long, J. B. Nicholas, J. A. Nichols, H. L. Taylor, A. T. Wong, G. I. Fann, R. J. Littlefield, and J. Nieplocha. Advances in parallel distributed data software; computational chemistry and NWChem. In *Applied Parallel Computing. Computations in Physics, Chemistry and Engineering Science*, volume 1041 of *Lecture Notes in Computer Science*. Springer, Heidelberg, 1996.

[8] Robert J. Harrison, Martyn F. Guest, Rick A. Kendall, David E. Bernholdt, Adrian T. Wong, Mark Stave, James Anchell, Anthony Hess, Rik Littlefield, George I. Fann, Jarek Nieplocha, Greg S. Thomas, David Elwood, Jeff Tilson, Ron L. Shepard, Albert F. Wagner, Ian T. Foster, Ewing Lusk, and Rick Stevens. High performance computational chemistry. II. A scalable SCF program. *J. Computat. Chem.*, 17:124, 1995.

[9] Meenakshi A. Kandaswamy, Mahmut T. Kandemir, Alok N. Choudhary, and David E. Bernholdt. An experimental study to analyze and optimize Hartree-Fock application's I/O with PASSION. *Int. J. Supercomputer Appl.*, in press.

[10] M.F.Guest, E.Apra, D.E.Bernholdt, H.A.Fruchtl, R.J.Harrison, R.A.Kendall, R.A.Kutteh, X.Long, J.B.Nicholas, J.A.Nichols, H.L.Taylor, A.T.Wong, G.I.Fann, R.J.Littlefield, and J.Nieplocha. High-performance computing in chemistry; nwchem. *Future Generation Computer Systems*, 12(4):273–289, December 1996.

Recent Collaborators

JL Anchell,[1] E Apra,[2] RJ Bartlett,[3] RR Birge,[4] AN Choudhary,[5] PD Ellis,[6] D Elwood,[6] GI Fann,[6] DF Feller,[6] IT Foster,[7] GC Fox,[4] HA Früchtl,[8] MF Guest,[9] RJ Harrison,[6] AC Hess,[6] MA Kandaswamy,[4] MT Kandemir,[4] RA Kendall,[6] RA Kutteh,[6] RJ Littlefield,[6] X Long,[6] E Lusk,[7] JB Nicholas,[6] JA Nichols,[6] J Nieplocha,[6] SA Perera,[3] RL Shepard,[7] MS Stave,[10] R Stevens,[7] HL Taylor,[6] GS Thomas,[6] JL Tilson,[7] AF Wagner,[7] AT Wong.[11]

Institutional Affiliations: [1]Schrödinger, Inc., [2]Center for Study of the Relations between Structure and Chemical Reactivity, Nat'l Research Council of Italy, [4]Syracuse Univ., [5]Northwestern Univ., [6]Pacific Northwest Nat'l Lab., [7]Argonne Nat'l Lab., [8]Fujitsu European Centre for Information Technology, [9]Daresbury Lab., UK, [3]Univ. of Florida, [10]Oxford Molecular, Inc., [11]Nat'l Energy Research Supercomputer Center.

Students and Postdoctoral Scholars

None

Advisors

Graduate: Rodney J. Bartlett, University of Florida
Postgraduate: Robert J. Harrison, Pacific Northwest National Laboratory

Scott Alan Klasky

Northeast Parallel Architectures Center	Phone: 315 443 1690
Syracuse University	Fax: 315 443 1973
111 College Place	E-mail: scott@npac.syr.edu
Syracuse, NY 13244-4100	URL: http://www.npac.syr.edu/users/scott

Education

Ph.D. in Physics from University of Texas at Austin (1989-1994).
B.S. in Physics from Drexel University, Philadelphia PA (1984-1989).

Professional Experience

1996– Senior Research Scientist, NPAC, Syracuse University **1995–1996** Post. Doc. Fellow in Relativity, The University of Texas at Austin

Awards and Honors

Phi Beta Kappa Honors Society 1991

Selected List of Publications:

[1] "Java based Collaborative Scientific Visualization" (w/ B. Ki), Concurrency: Practice and Experience, accepted 1998.

[2] "Schwarzschild-Perturbative gravitational wave extraction and outer boundary conditions" (w./ Abrahams et. al.), submitted to Phys. Rev. Letters 1997.

[3] "Boosted three-dimensional black-hole evolutions with singularity excision" (w/ Cook et. al.), accepted: Phys. Rev. Letters 1997.

[4] "The Binary Black Hole Grand Challenge ADM code", (w/ Cook et. al.), to be submitted to Phys. Rev. D, 1998.

[5] "Collaborative Scientific Visualization" (w/ B. Ki), Concurrency: Practice and Experience, November 1997.

[6] "Multigrid- An Approach in HPF" (w/ U. Dittmer), NPAC technical report (SCCS-772), 1996.

[7] "Multigrid support with the DAGH package: Specifications and Applications" (w/ M. Choptuik et al.), Site report, 1995.

[8] "Visualizing Complex Patterns in the Spread of Head and Neck Cancers," (w/ L. Gray et al.), The International Journal of Supercomputer Applications 7, 167 (1993).

[9] "Three-dimensional initial data for the collision of two black holes," (w/ G. Cook et al.), Physical Review D47, 1471 (1993).

[10] "Properties of gravitational "solitons" ", (w/ J. Centrella et al.), Physical Review D43, 379 (1991).

Summary of Interests:

For the last nine years Klasky has designed several major computer codes in the areas of physics, computer science, and numerical analysis. His main area of expertise is in designing large scale codes in the area of computational science. Klasky has lead teams of researchers to develop state-of-the-art computer codes in the area of high performance scientific computing/physics. He has expertise in solving large scale Partial Differential Equations (PDE's), particularly for numerical relativity, using state-of- the-art techniques (Adaptive Mesh Refinement) as well as in designing collaborative visualization tools, to be used over the Internet. He has also designed Monte Carlo codes to price derivatives using a Path Integral Monte Carlo Approach.

Ph.D. Advisor

Richard Matzner

F. Summary Proposal Budget

Replace with NSF Form 1030s.

G. Current and Pending Support

Replace with NSF Form 1239.

H. Facilities, Equipment, and Other Resources

H.1 Local Computer Facilities

NPAC has a substantial computer infrastructure, including networked workstations, Mac and PC equipment, and various parallel computers. Parallel and distributed platforms include an 8-CPU SGI Power Challenge, a cluster of about 14 Sun UltraSPARCs interconnected with an ATM network (used for a mix of database, educational and computational resources), and a cluster of 10 PC NT systems. These clusters are shared by all NPAC users, but available in dedicated parallel mode on a reservation basis.

NPAC has a professional system support group which handles day to day operation and maintenance of the computer facilities.

H.2 Office

NPAC provides adequate office facilities for all participants.

H.3 Major Equipment

The computer systems described above constitute the major equipment for this effort.

H.4 Other Resources

NPAC provides basic secretarial services, duplicating facilities, and the like.

I. Letters of Cooperation and Support

1. Professor Xiaoming Li, Dept of Computer Science and Technology, Peking University, Beijing, 100871 China (1 page).

2. Dr Emil J. Sarpa, Manager, External Research, Sun Microsystems, Inc (1 page).