
Tinkertoy parallel programming: a case study with Zoltan

Karen Devine* and Bruce Hendrickson

Discrete Algorithms and Mathematics Department,
Sandia National Laboratories, Albuquerque, NM 87185-1111, USA
E-mail: kddevin@sandia.gov E-mail: bah@sandia.gov
*Corresponding author

Abstract: As the need for complex parallel simulation software grows, better strategies for efficient and effective software development become important. We advocate a toolkit- or ‘tinkertoy’-approach to parallel application development. By providing efficient implementations of basic services commonly needed by applications, toolkits allow application developers to benefit from others’ research, compare algorithms and save time for their own development. Unlike large frameworks, toolkits provide these services with light-weight interfaces and little or no restriction on application data structures, making them easy to use in both new and existing applications. In this paper, we describe features of effective toolkit design, using the Zoltan parallel, dynamic data management toolkit as an example.

Keywords: parallel computing; parallel programming; software engineering; software framework; load balancing.

Reference to this paper should be made as follows: Devine, K. and Hendrickson, B. (2005) ‘Tinkertoy parallel programming: a case study with Zoltan’, *Int. J. Computational Science and Engineering*, Vol. 1, Nos. 2/3/4, pp.64–72.

Biographical notes: Karen Devine earned her BS in Computer Science from Wilkes College, and her MS and PhD in Computer Science from Rensselaer Polytechnic Institute. She is a principal member of the technical staff in the Discrete Algorithms and Mathematics Department at Sandia National Laboratories in Albuquerque, New Mexico. She is the principal investigator for Zoltan, a toolkit of parallel data management and load-balancing algorithms. Her research interests include parallel algorithms, load balancing, adaptive finite element methods and software design.

Bruce Hendrickson received degrees in Math and Physics from Brown University, followed by a PhD in Computer Science from Cornell. He has been at Sandia National Laboratories in Albuquerque for the past 14 years where he holds the title of distinguished Member of technical staff and manages the Computational Math and Algorithms Department. He also has an appointment in the Computer Science Department at the University of New Mexico. He is an Editor of several leading journals in scientific and parallel computing and has helped to organise numerous international meetings. His research interests include combinatorial scientific computing, parallel algorithms, linear algebra, graph algorithms, scientific software and data mining.

1 Introduction

Developing software for parallel scientific simulations is always a challenge. Parallel simulations require a wide range of capabilities, from meshing tools and data managers to solvers and visualisation tools. Dynamic and/or adaptive simulations present an even greater challenge, as load redistribution, synchronisation and more complicated data structures must be managed. High parallel performance is always desired, requiring expertise by the software developer in efficient algorithm design and implementation. Development schedules are often tight. And parallel architectures change with each new generation of computers, requiring portability of codes and providing a ‘moving target’ for performance optimisation.

In such an environment, which is the best approach for the development of complex adaptive software?

Several approaches for parallel software development exist, each with its own advantages and disadvantages. Application developers could do all the software development themselves. This option is attractive because it gives developers total control of the software. This control, however, comes at a severe price. Do-it-yourself programming is time consuming, as much effort is spent writing code that is often available elsewhere – reinventing the wheel, so to speak. Moreover, developers must spend much of that time writing code in areas outside their areas of expertise and interest, resulting in nonexpert and, quite possibly, less efficient implementations of many parts of the software.

Software frameworks provide a second option for software development. A framework is an application or library providing a wide range of services and data structures for a specific class of applications; application developers use the framework's data structures and services in constructing their simulations. SIERRA (Edwards, 2002) and Overture (Brown et al., 1997) are examples of successful adaptive simulation frameworks. Such frameworks provide many capabilities in one package, which is a significant advantage to applications that need all the capabilities. However, frameworks typically are large and can have substantial overhead, which is a disadvantage to applications needing only a small subset of a framework's capabilities. Frameworks are also difficult to add to existing applications; instead, existing applications must be incorporated like new applications into the framework. To use a framework, application developers must learn both its interfaces and data structures, which is often a time-consuming task. In addition, framework use makes application developers highly dependent on the framework's developers, perhaps causing an undesirable loss of control in terms of enhancements and schedules for the application developers.

As an alternative, we advocate a toolkit- or 'tinkertoy'-approach to software development. The original tinkertoy, made for children by Hasbro, is a set of simple wooden pieces that can be interconnected in different ways to make surprisingly complex structures and machines. Similarly, in tinkertoy software development, applications are constructed of small, simple software parts with flexible, easy-to-use interfaces. These simple software parts are toolkits – libraries containing basic services commonly needed by applications. While frameworks provide all data structures and services for a specific type of application, toolkits provide only a small set of related services that can be used as parts of many different applications. Application developers can put together these services to create a larger application. For example, an application could be constructed from an adaptive meshing toolkit (e.g., Pyramid (Norton et al., 2001) and AOMD (Remacle et al., 2002, 2003)), a dynamic load-balancing toolkit (e.g., Zoltan (Devine et al., 1999a, 1999b, 2002), DRAMA (Maerten et al., 1999) and ParMETIS (Karypis et al., 2003)), a linear and non-linear solver library (e.g., Trilinos (Heroux et al., 2003; Heroux and Willenbring, 2003), Aztec (Hutchinson et al., 1995) and PETSc (Balay et al., 1997, 2002)) and some visualisation tools (e.g., VTK (Schroeder et al., 2003)). Application developers, then, can concentrate on the simulation details in which they are most interested (e.g., physics or engineering).

This toolkit approach has a number of advantages. Because they typically provide a smaller number of basic services, toolkits are less cumbersome than frameworks; application developers can select and use only the functionality they need. Because toolkits generally use

library interfaces and private data structures, developers can incorporate them easily into both new and existing applications. Application developers need to learn only the toolkit interfaces, rather than all its internal data structures, so start-up time is shorter than with framework use. Well-designed toolkits can, like tinkertoys, be easily hooked together to build larger and more complex functionality, beyond the scope of any single library. And since most toolkits are developed by experts in the toolkits' capabilities, application developers benefit in terms of both development time and algorithmic efficiency by using toolkits.

Toolkits do, however, have some of the same trust and dependence issues as frameworks, although to a lesser degree. There are often competing toolkits providing similar functionality, so an application developer can switch if the need arises. Also, the simple interfaces associated with well-designed toolkits facilitate replacement if necessary. Toolkit users are dependent upon toolkit developers to provide correct algorithms and customer support. Open-source distribution used by many toolkits can increase reliability and trust by allowing users to inspect the implementations and by providing a broad testing community for the software. Toolkits also have some memory and performance overhead owing to separation of toolkit data structures from the applications, but with careful design, these costs can be kept acceptably low.

There are, of course, hybrids of all these strategies. The Common Component Architecture (CCA), for example, provides interfaces that allow plug-and-play interoperability of components, in line with the toolkit philosophy (Armstrong et al., 1999). The components, however, are launched within a framework (e.g., Ccaffeine (Allan et al., 2002)) that manages the components' operation. For this paper, however, we will focus on straightforward toolkit usage for parallel computing.

Is it really possible to build complex applications out of tinkertoys? It is widely accepted that linear solvers can be encapsulated as libraries, but what about the needs of complex, adaptive parallel applications like adaptive mesh calculations or particle simulations? It is often presumed that these kinds of dynamic applications require such intricate control over data structures that toolkits cannot easily be applied.

One goal of this paper is to give affirmative answers to these questions. We believe that even complicated, adaptive computations can be constructed efficiently and effectively from simple tools. As a second goal, the paper describes our attempt to instantiate this vision through Zoltan – a toolkit for adaptive parallel computation built with the tinkertoy philosophy. Finally, through Zoltan we have been exploring the possibilities and limits of tinkertoys. Specifically, what types of functionality can be delivered through application-independent toolkits, and what can be provided only by applications or frameworks? This paper reports on our current understanding of this important issue.

2 Zoltan overview

The Zoltan toolkit is a collection of data management services for parallel unstructured, adaptive and dynamic applications, available as open-source software from <http://www.cs.sandia.gov/Zoltan>. It is designed to simplify the load balancing, data movement, unstructured communication and memory usage difficulties that arise in dynamic applications such as adaptive finite element methods, particle methods and multiphysics simulations. Zoltan's data-structure neutral design allows it to be used by a wide range of applications without imposing restrictions on application data structures. Its object-based interface provides a simple and inexpensive way for application developers to use the library and researchers to make new capabilities available under a common interface.

As we detail in the subsequent sections of this paper, Zoltan provides tools that help application developers without imposing strict frameworks on them. For example, it includes parallel partitioning algorithms and data migration tools that help redistribute data to reflect, for example, changing processor workloads resulting from creation of elements in adaptive finite element methods. Zoltan also includes distributed data directories, dynamic memory-debugging tools and unstructured communication services that enable applications to perform complicated communication using only a few simple primitives. Zoltan is used in a variety of applications, including contact detection and crash simulations (Brown et al., 2003; Koteras and Gullerud, 2003), adaptive finite element methods (Edwards, 2002; Boucheron et al., 2002; Lawrence et al., 2003), parallel circuit simulations (Hutchinson et al., 2002), multiphysics simulations (Salinger et al., 1996) and linear solvers and preconditioners (Heroux et al., 2003; Heroux and Willenbring, 2003).

3 The promise and limitations of toolkits

The success of tinkertoy computing depends on a number of software design features:

- *functionality*: toolkits must solve problems that appear in multiple applications
- *portability*: toolkits must be portable across multiple parallel platforms
- *interfaces*: toolkits' software interfaces should be easy to use
- *added value*: toolkits should give application developers greater performance and flexibility
- *low overhead*: the overhead owing to toolkit use must be small, both in memory and in runtime
- *support*: toolkit developers should help application developers use their toolkits effectively.

Each of these issues provides challenges for advocates of tinkertoy parallel computing. We discuss below these challenges in the context of support for adaptive parallel

computations, and explain how we have tried to address them within Zoltan.

3.1 Functionality

Selection of services provided by a toolkit is a critical design step for toolkit designers. Toolkits should include services commonly needed by applications. Services should remain independent of each other as much as possible so that application developers can select and use only the tools that they want. In addition, services within a toolkit should be related to and complement each other. Toolkit developers should fight the urge to incorporate every possible service, so that toolkits do not become too large and difficult to use (and, indeed, start resembling frameworks).

As an example, efficient parallel implementation of adaptive applications requires dynamic load balancing to redistribute work to processors after adaptive refinement occurs. Dynamic load balancing involves both the computation of a new partitioning of data and workload, and movement of data to new processors. Moreover, dynamic data redistribution creates new needs for applications as they dynamically delete and insert data in their data structures, relocate needed off-processor data and build new communication patterns. The Zoltan toolkit includes functionality to address many of these related needs.

Zoltan's main utility is a suite of dynamic load-balancing algorithms that compute new distributions of data to processors. Since dynamic load balancers must run side-by-side with applications, Zoltan is implemented in parallel and is scalable in both execution and memory usage. For load balancing, it takes an existing distributed partition as input and computes a description of the new partition in terms of objects to be transferred between processors. Many of the partitioning algorithms are incremental, that is, small changes in processor workloads result in only small changes in the resulting decompositions. Zoltan's partitioning algorithms support non-uniform partition sizes and unequal numbers of partitions and processors. Additional utilities that compute which processors' partitions intersect a given point or region in space are provided for geometric partitioning methods; these utilities are key kernels of parallel contact detection simulations (Brown et al., 2003; Koteras and Gullerud, 2003).

After obtaining the map of a new decomposition, applications must move data from their old processors to their new processors. This data migration requires deletions and insertions from the application data structures, along with communication between the processors. A general-purpose toolkit like Zoltan can do little to help with the manipulation of application-specific data structures. However, because Zoltan has knowledge of both the old and new partitions, it can communicate object data among processors easily. In fact, by using user-supplied functions to pack and unpack data into communication buffers, Zoltan's data migration tools can perform all communication necessary to send data to their new location.

Zoltan's distributed data directories (based on the rendezvous algorithm of Pinar and Hendrickson (2001)) provide additional functionality related to dynamic data redistribution. After repartitioning, for example, a processor may need to rebuild ghost cells and lists of objects to be communicated; it may know which objects it needs, but may not know where they are located. By using Zoltan to locate this off-processor data, processors register data along with their processor numbers in a directory that is distributed evenly across processors in a predictable way (e.g., a linear decomposition of the data or a hashing of data to processors). Then, other processors obtain the processor number of a given object by sending a request for the information to the processor holding the directory entry. Thus, communication cost for look-ups is constant and total memory usage is linear in the amount of data (each registered object requires storage for approximately seven integers). Moreover, since the directory is distributed, no communication bottlenecks develop (as they would for a directory located completely on one processor).

The Zoltan toolkit provides further capability to dynamic applications with complicated and/or changing communication patterns. For example, multiphysics simulations and crash simulations may require complicated communication patterns to transfer data between decompositions for different simulation phases. To simplify this communication, Zoltan provides an unstructured communication package that generates a communication 'plan' with information about sends and receives for a given processor. The plan may be used and reused throughout the application, or it may be destroyed and rebuilt when communication patterns change. Simple communication primitives in the toolkit insulate users from details of sends and receives.

Similarly, memory usage in dynamic applications can change throughout the simulation. After repartitioning, for example, new memory is needed for imported data and exported data's memory is freed. Memory leaks are common in developing software. While there are many software development tools that enable users to track memory bugs (<http://www.rational.com>; Seward and Nethercote, 2003), these tools are often not available on state-of-the-art parallel computing platforms. Thus, Zoltan provides basic in-application memory-debugging tools that are simple wrappers around memory allocation routines. The wrappers record information (e.g., line number and file name) about memory operations, allowing developers to track memory leaks and print memory-usage statistics.

While these related tools operate well together, an important feature of Zoltan's toolkit design is separation between tools. Application developers can use only the tools they want; for example, they can use Zoltan to compute decompositions but perform all data migrations themselves. They can build Zoltan distributed data directories that are completely independent of load balancing. They can use Zoltan's unstructured communication tools within statically balanced applications or they can use Zoltan to perform

all the data management tasks associated with load balancing. In this way, Zoltan provides full service for dynamic partitioning, while allowing developers the flexibility to use Zoltan's tools in a variety of ways—even ways not originally envisioned by Zoltan's designers.

3.2 Portability

A toolkit is useful to a broad community only if it is portable across many platforms. In addition to allowing toolkit use on many current architectures, portability allows the toolkit to be used across generations of machines. Developers are more apt to be willing to use toolkits if they know that the software will continue to work as machines are upgraded or replaced.

To ensure portability, toolkits must rely on standards as much as possible. They should use only standard language features, to prevent compilation difficulties. Since many cutting-edge language features are not supported by older compilers, toolkits should include code that is as simple as functionally possible. Toolkit dependence on other libraries should be kept to a minimum; few things are more frustrating than trying to build a toolkit only to discover that many other libraries must be located, purchased, downloaded and/or installed first. Necessary dependencies should take advantage of 'standard' libraries (MPI, 1995; *MPI-2*, 1997; OpenGL (Segal and Akeley, 2003; BLAS, 2001), etc.) as much as possible.

The Zoltan toolkit is implemented in ANSI C, with an optional Fortran90 interface available. It uses MPI for all communication; it can use any version of MPI and has been tested using MPICH, LAM and system-specific MPI libraries for IBM, DEC and Intel architectures. To use Zoltan's graph partitioners, applications must link with Zoltan and either ParMETIS (Karypis et al., 2003) or Jostle (Walshaw et al., 1997; Walshaw, 2002); a compatible version of ParMETIS is distributed with Zoltan to simplify building and configuring both libraries. However, if graph partitioning is not needed, dependence on ParMETIS and/or Jostle can be excluded from Zoltan.

3.3 Easy-to-use interfaces

Toolkits' capabilities should be easily accessible by many different applications. To accomplish this goal, several features are needed. There should be separation between the application and toolkit data structures so that toolkit use is not restricted to a particular application. Toolkits should have simple interfaces that do not require extensive programming by the application developer. And toolkits should fit easily into both existing and new applications, allowing application developers to retrofit and update their existing codes.

Separation between application and toolkit data structures is achieved through data-structure neutral toolkit design. In a data-structure neutral design, details of the toolkit data structures are hidden from the application and vice versa. Thus, the toolkit does not impose data structures upon an application as frameworks do. This separation can

be achieved in several ways. Some toolkits (e.g., ParMETIS (Karypis et al., 2003) and Jostle (Walshaw, 2002)) require the application to build specific data structures (e.g., graphs) for them to use. While this requirement is acceptable, it has the drawback that data-structure changes in the toolkit require changes to both the toolkit interface and application. It also burdens the application programmer with the task of creating a complex data structure, and it may incur a significant memory overhead if the library creates yet another copy of the data structure. Other toolkits (e.g., Trilinos (Heroux et al., 2003; Heroux and Willenbring, 2003)) provide an object interface (e.g., a matrix) and methods for performing operations on the objects (e.g., transposition). This interface allows greater code hiding than the previous approach.

The Zoltan toolkit uses a callback function interface in which Zoltan calls user-supplied functions to obtain needed application data. The functions answer questions like, ‘How many data items are owned by this processor?’ and ‘What are the geometric coordinates of the data items on this processor?’ The application developer must provide simple functions that answer these queries. Then, Zoltan calls the functions to build appropriate data structures for the particular tool requested. This approach has several advantages in terms of both ease of use and ease of maintenance. First, once application developers implement the callback functions, they can access all technology within Zoltan without additional construction of data structures; as new capabilities are added to the toolkit, users can access them with little effort. Second, by not requiring users to build data structures for them, Zoltan developers can use the most efficient data structures for their algorithms and can improve them without impacting the applications.

Third, the user interface remains unchanged regardless of any internal changes in Zoltan, allowing users to upgrade versions of Zoltan with no change to their applications. Finally, at no time does the application developer have to build (or debug!) complicated data structures for use within Zoltan.

Toolkit interfaces should be simple to understand and utilise for users with various levels of interest and expertise. Only a small set of functions should be needed to invoke the toolkit’s basic capabilities; additional functions can be provided to support more advanced features. Parameters can be used to control toolkit functionality, but reasonable default values should be provided. With this layered design, amateur users can benefit from the basic toolkit functionality, while more advanced and interested users can experiment with a broader range of options. Zoltan uses only a small set of callback functions and makes them easy to write by requesting only information that is, in general, easily accessible to applications. For the most basic partitioning algorithms, Zoltan requires only four callback functions; these functions return the number of objects owned by a processor, a list of weights and names for owned data, the dimensionality of the problem and the coordinates of a given owned object. More sophisticated graph-based partitioning and matrix-ordering algorithms

require only two additional callback functions, returning the number of edges per data object and edge lists for data objects. All algorithms have parameters that can alter the algorithms’ performance and results; default values are set to reflect the most common scenarios for algorithm use.

Toolkits should be easy to use in both new applications and existing ones. When toolkits allow individual tools to be used independently, application developers can incorporate the toolkits incrementally into their applications. For example, an application developer may replace a load-balancing scheme in an existing dynamic application with a partitioning algorithm from Zoltan, but continue using the data migration code previously written in the application.

3.4 *Added value*

Since toolkits are most often implemented by researchers in the areas addressed by the toolkit, they can provide high performance implementations of state-of-the-art algorithms. Thus, by using toolkits, application developers can focus on their particular areas of interest, rather than concern themselves with every detail of the parallel simulation. Instead of trying to understand the state of the art in every field, they can concentrate on research in their own field. Likewise, they can provide valuable user feedback to toolkit developers, creating a synergy that benefits both application developers and toolkit researchers.

Toolkits can also add value to applications by providing a number of different algorithms whose effectiveness can be compared within an application. For example, there is no single partitioning strategy that is effective for all parallel computations. Some applications require partitions based on only the workloads and geometry of the problem; others benefit from explicit consideration of dependencies between objects. Some applications require the highest quality partitions possible, regardless of the cost to generate them; others can sacrifice some quality as long as new partitions can be generated quickly. Most importantly, an application developer may not know in advance which strategy works best in his application. By providing a collection of algorithms and a convenient way to compare them, toolkits can significantly improve application performance with little additional effort required by application developers. By facilitating easy algorithmic comparisons, toolkits also help advance algorithmic research.

In the Zoltan library, we have included a suite of parallel partitioning algorithms. Three classes of algorithms are provided: geometric bisection, space-filling curves and graph partitioning. Within each class, several different algorithms are implemented. Geometric algorithms include Recursive Coordinate Bisection (Berger and Bokhari, 1987) and Recursive Inertial Bisection (Taylor and Nour-Omid, 1994). Space-filling curve partitions are generated via a binned Hilbert Space-Filling Curve algorithm (Devine et al., 2005; Bauer, 2002), Octree partitioning (Loy, 1998; Gervasio, 1998; Campbell et al., 2003) or a Refinement Tree Partitioning algorithm designed especially

for adaptive mesh refinement applications (Mitchell, 1995, 1998). Graph partitioning is provided through easy-to-use interfaces to ParMETIS (Karypis et al., 2003) and Jostle (Walshaw, 2002). Once users write the callback functions for each class, switching between classes and methods requires only a single-parameter change with the new algorithm name. In this way, developers can compare algorithms easily within their applications to find the strategy that works best for them.

3.5 Low overhead

The performance obtained using a toolkit will almost never be as high as the performance of an equivalent algorithm embedded directly within an application. Data separation and general application interfaces require additional memory use and computation time for creating toolkit data structures. However, with careful design, this overhead can be kept acceptably low; the additional cost can be tolerated when the application benefits from toolkit functionality.

The amount of overhead that can be tolerated depends, of course, on the application's use of the toolkit. Greater overhead can be tolerated if the tools are invoked infrequently. However, since we expect dynamic load balancing and its related functionality to be executed frequently during a simulation, low overhead is important in Zoltan. Multiple versions of many callback functions (e.g., list-based functions that return arrays of data vs. iterator functions that return one data item at a time) are provided to applications, allowing application developers to pick the interface most suitable for their data structures. Callbacks allow Zoltan's algorithms to directly build the data structures they need; no intermediate data structure is used. The callbacks also allow the algorithms to obtain only the data they need; for example, geometric algorithms do not require graph information and, thus, do not obtain that information from the application.

Table 1 provides evidence that Zoltan's callback function interface adds only a small overhead to a simulation. Experiments using several Zoltan partitioning algorithms were run using the mesh-based driver application distributed with Zoltan. Two types of overhead are measured: the overhead associated with calling the user-specified callback functions (e.g., to obtain coordinate or graph information) and the overhead associated with building the data structures needed for load balancing. While the callback function overhead would not exist if the load-balancing algorithms were embedded directly in the applications, most of the overhead needed to build data structures for load balancing would still be incurred. Cases where the application could use exactly the same data structures for computation and partitioning are rare; for example, few applications' computations use as their native data structure the compressed, distributed graph structure commonly used in graph partitioning. Thus, the overhead of building the data structures cannot be completely disregarded for embedded partitioning algorithms.

Table 1 Overheads incurred when using Zoltan in a 3D mesh-based application. Overhead is reported both in seconds and as a percentage of total partitioning time. Callback overhead is the time spent in callback functions; this time is unique to Zoltan. Build overhead includes time spent constructing load-balancing data structures; this time would most likely be incurred by applications using embedded load balancing

	<i>RCB</i>	<i>HSFC</i>	<i>ParMETIS PartKWay</i>
Callback time	0.038 (4.4%)	0.037 (8.7%)	0.096 (3.6%)
Data-structure build time	0.066 (7.7%)	0.059 (14.0%)	0.288 (10.7%)
Total partitioning time	0.865	0.423	2.674

In the experiments, a three-dimensional, unstructured finite element mesh with about one million elements was randomly distributed on 16 processors. New decompositions were computed using three of Zoltan's partitioning algorithms: Recursive Coordinate Bisection (RCB), Hilbert Space-Filling Curves (HSFC) and a graph-based method (ParMETIS PartKWay). The time reported in the table is just that associated with the partitioning itself. Once the partitioning has been computed, additional time will be required to migrate data and to update the data structures on each processor. In our experience, these latter operations are several times as costly as the partitioning itself, yet they are independent of the use of a toolkit. This further reduces the significance of the overhead times in the table.

The time spent in callback functions, time spent building data structures (including the callback-function time) and total partitioning time (including both callbacks and data-structure construction) were measured. The geometric algorithms, RCB and HSFC have similar amounts of overhead, since they both use only geometric information about data to be partitioned. Because the graph-based partitioner requires more application data (e.g., the edge lists for each graph vertex) and more complicated data structures (i.e., a distributed graph), its overheads are somewhat higher. Still, for all algorithms, the callback overhead is less than 9% of the total partitioning time. Similarly, the time required to build data structures is less than 15% of the partitioning time for all three algorithms. And as discussed above, this time is likely to be required even by embedded partitioning implementations.

As the table indicates, geometric partitioners are often faster than graph-based partitioners. However, there may be compensating differences in partition quality. The goal of this table is to merely quantify the overhead associated with the use of Zoltan and not to compare the merits of different partitioning strategies.

We must emphasise that the overhead incurred in using Zoltan can vary depending upon the application and its implementation of the callback functions. If the application gives Zoltan expensive implementations of callback functions, the callback-function overhead will, of course, increase.

3.6 Support

Convincing application developers to use someone else's code is often a difficult proposition. Their reluctance is understandable, as they cannot always ascertain the quality of outside code, and they have justifiable worries about long-term maintenance and development. In this regard, toolkits are more attractive than frameworks, because the commitment of an application developer is less with the former than the latter. But, although reduced in significance, support issues are still important. Toolkit designers can help ease these concerns in several ways. First, open-source distribution of toolkit software allows users to study and experiment with the software. A number of open-source licenses are available with varying levels of protection for toolkit developers (<http://www.opensource.org/licenses/index.php>). Zoltan is distributed using the GNU Lesser General Public License (<http://www.gnu.org/copyleft/lesser.html>), which allows free use of Zoltan software and guarantees that redistributions of Zoltan are also free. Second, toolkit developers must provide documentation for their software, with user's guides describing functionality and options. User's guides should provide detailed descriptions of capabilities, options, interfaces and configurations, and should include usage examples whenever appropriate. Zoltan's hyperlinked, web-based User's Guide (Devine et al., 1999a) has proved to be useful both to users and Zoltan developers. Third, simple codes using the toolkit can be distributed with the toolkit. While allowing users to verify their toolkit installation, these codes serve as examples that application developers can study in learning to use the toolkit. Most toolkit developers have such programs available for their own testing; including the examples with the toolkit distribution takes little additional effort. For example, code, instructions and sample inputs and outputs for the Zoltan regression test programs are included in the Zoltan distribution. And finally, the promise of customer support from toolkit developers encourages application developers to give toolkits a try.

4 Future work

The difficulty of software development continues to be a principal impediment in the adoption of high-performance computing. It is our belief that well-executed toolkits will play a growing role in addressing this problem. Although they provide less support than full-fledged frameworks, toolkits have the advantage of greater flexibility and incremental adoption. But toolkits never can address the detailed manipulation of application-specific data structures. Thus, we anticipate a continuing need for frameworks and hope for a decline in the heroic but inefficient practice of all-in-one application development. We look forward to a day when the community has built a diverse set of toolkits with clean interfaces, and application developers can mix and match them like tinkertoys to build complex yet high-performing applications quickly.

We feel that Zoltan offers an attractive model for toolkit development. It provides a diverse set of related services via a simple interface and low overhead in both memory and runtime. We continue to add functionality to Zoltan including support for heterogeneous parallel architectures and alternative models of load balancing. However, in doing so, we are constantly facing choices about the appropriate breadth of functionality of the toolkit and about the complexity of the interface. In our experience, good toolkit development requires both discipline and humility. We want to avoid adding functionality capriciously, and to focus on the areas in which we are best able to provide novel support to application developers.

We hope that the proven success of Zoltan will attract attention to the promise of tinkertoy parallel programming. More specifically, we hope to persuade others to build additional toolkits with complementary functionality, and to convince application developers of the promise this approach has in the development of more, and more complex, simulation codes.

Acknowledgements

The ideas and opinions in this paper have grown out of our many years of experience developing supporting infrastructure for parallel computations. Along the way, we have learned from many colleagues including Steve Plimpton, Mike Heroux, Erik Boman, John Shadid, Bob Heaphy, Courtenay Vaughan and Bill Mitchell.

Sandia is a multiprogram laboratory operated by Sandia Corporation, a Lockheed Martin Company, for the US Department of Energy under Contract DE-AC04-94AL85000.

References

- Allan, B.A., Armstrong, R.C., Wolfe, A.P., Ray, J., Bernholdt, D.E. and Kohl, J.A. (2002) 'The CCA core specification in a distributed memory spmd framework', *Concurrency Computat*, Vol. 14, pp.1–23.
- Armstrong, R., Gannon, D., Geist, A., Keahey, K., Kohn, S., McInnes, L., Parker, S. and Smolinski, B. (1999) 'Toward a common component architecture for high-performance scientific computing', *Proceedings of the 1999 Conference on High Performance Distributed Computing*, Redondo Beach, CA, August.
- Balay, S., Buschelman, K., Gropp, W.D., Kaushik, D., Knepley, M., McInnes, L.C., Smith, B.F. and Zhang, H. (2002) *PETSc Users Manual*, Argonne National Laboratory, Tech. Rep. ANL-95/11 – Revision 2.1.5.
- Balay, S., Gropp, W.D., McInnes, L.C. and Smith, B.F. (1997) 'Efficient management of parallelism in object oriented numerical software libraries', in Arge, E., Bruaset, A.M. and Langtangen, H.P. (Eds.): *Modern Software Tools in Scientific Computing*, Birkhauser Press, pp.163–202.

- Basic Linear Algebra Subprograms Technical (BLAST) Forum Standard, (2001) Basic Linear Algebra Subprograms Technical (BLAST) Forum, University of Tennessee, Knoxville, Tennessee, <http://www.netlib.org/blas/blast-forum/>.
- Bauer, A.C. (2002) *Efficient Solution Procedures for Adaptive Finite Element Methods – Applications to Elliptic Problems*, PhD Dissertation, State University of New York at Buffalo.
- Berger, M.J. and Bokhari, S.H. (1987) 'A partitioning strategy for nonuniform problems on multiprocessors', *IEEE Trans. Computers*, Vol. 36, pp.570–580.
- Boucheron, E.A., Brown, K.H., Budge, K.G., Burns, S.P., Carroll, D.E., Carroll, S.K., Christon, M.A., Drake, R.R., Garasi, C.G., Haill, T.A., Peery, J.S., Petney, S.V., Robbins, J., Robinson, A.C., Summers, R., Voth, T.E. and Wong, M.K. (2002) *ALEGRA: User Input and Physics Descriptions Version 4.2*, Sandia National Laboratories, Albuquerque, NM, Tech. Report SAND2002-2775.
- Brown, D.L., Chesshire, G.S., Henshaw, W.D. and Quinlan, D.J. (1997) 'OVERTURE: An object-oriented software system for solving partial differential equations in serial and parallel environments', *Eighth SIAM Conf. on Parallel Processing for Scientific Computing*, SIAM, March, Minneapolis, MN.
- Brown, K.H., Glass, M.W., Gullerud, A.S., Heinstejn, M.W., Jones, R.E. and Voth, T.E. (2003) *ACME Algorithms for Contact in a Multiphysics Environment, API Version 1.3*, Sandia National Laboratories, Albuquerque, NM, Tech. Report SAND2003-1470.
- Campbell, P.M., Devine, K.D., Flaherty, J.E., Gervasio, L.G. and Teresco, J.D. (2003) *Dynamic Octree Load Balancing Using Space-Filling Curves*, Williams College Department of Computer Science, Tech. Rep. CS-03-01.
- Devine, K., Boman, E., Heaphy, R., Hendrickson, B. and Vaughan, C. (2002) 'Zoltan data management services for parallel dynamic applications', *Computing in Science and Engineering*, Vol. 4, No. 2, pp.90–97.
- Devine, K., Hendrickson, B., Boman, E., St. John, M. and Vaughan, C. (1999a) *Zoltan: A Dynamic Load Balancing Library for Parallel Applications; User's Guide*, Sandia National Laboratories, Albuquerque, NM, Tech. Report SAND99-1377, <http://www.cs.sandia.gov/Zoltan/ug/html/ug.html>.
- Devine, K., Hendrickson, B., Boman, E., St. John, M. and Vaughan, C. (1999b) *Zoltan: A Dynamic Load Balancing Library for Parallel Applications; Developer's Guide*, Sandia National Laboratories, Albuquerque, NM, tech. Report SAND99-1376 <http://www.cs.sandia.gov/Zoltan/dev/html/dev.html>.
- Devine, K.D., Boman, E.G., Heaphy, R.T., Hendrickson, B., Teresco, J.D., Faik, J., Flaherty, J.E. and Gervasio, L.G. (2005) 'New challenges in dynamic load balancing', *Appl. Numer. Math.*, Vol. 52, Nos. 2–3, pp.133–152.
- Edwards, H.C. (2002) *SIERRA Framework Version 3: Core Services Theory and Design*, Sandia National Laboratories, Tech. Rep. SAND2002-3616, Albuquerque, NM.
- Gervasio, L.G. (1998) *Octree Load Balancing Techniques for the Dynamic Load Balancing Library*, Master's Thesis, Computer Science Dept., Rensselaer Polytechnic Institute, Troy.
- GNU Lesser General Public License', Free Software Foundation, <http://www.gnu.org/copyleft/lesser.html>.
- Heroux, M., Bartlett, R., Howle, V., Hoekstra, R., Hu, J., Kolda, T., Lehoucq, R., Long, K., Pawlowski, R., Phipps, E., Salinger, A., Thornquist, H., Tuminaro, R., Willenbring, J. and Williams, A. (2003) *An Overview of Trilinos*, Sandia National Laboratories, Albuquerque, NM, Tech. Rep. SAND2003-2927.
- Heroux, M.A. and Willenbring, J.M. (2003) *Trilinos Users Guide*, Sandia National Laboratories, Albuquerque, NM, Tech. Report SAND2003-2952.
- Hutchinson, S.A., Keiter, E.R., Hoekstra, R.J., Waters, L.J., Russo, T.V., Rankin, E.L. and Wix, S.D. (2002) *Xyce Parallel Electronic Simulator User's Guide, Version 1.0*, Sandia National Laboratories, Albuquerque, NM, Tech. Report SAND2002-3790.
- Hutchinson, S.A., Shadid, J.N. and Tuminaro, R.S. (1995) *Aztec User's Guide: Version 1.0*, Sandia National Laboratories, Albuquerque, NM, Tech. Rep. SAND95-1559.
- IBM Rational PurifyPlus, IBM, <http://www.rational.com>.
- Karypis, G., Schloegel, K. and Kumar, V. (2003) *ParMETIS: Parallel Graph Partitioning and Sparse Matrix Ordering Library Version 3.1*, University of Minnesota, Minneapolis.
- Koteras, J.R. and Gullerud, A.S. (2003) *Presto User's Guide: Version 1.05*, Sandia National Laboratories, Albuquerque, NM, Tech. Report SAND2003-1089.
- Lawrence, S.J.P., Musson, C. and Schmidt, R.C. (2003) 'MEMS fabrication modeling with ChISELS: A massively parallel 3D level-set based feature scaled modeler', *Proc. 2003 Nanotechnology Conference and Trade Show*, Vol. 3, San Francisco, CA, February, also Sandia National Laboratories, Tech. Rep. SAND2002-3994C.
- Loy, R.M. (1998) *Adaptive Local Refinement with Octree Load-balancing for the Parallel Solution of Three-dimensional Conservation Laws*, PhD Dissertation, Computer Science Dept., Rensselaer Polytechnic Institute, Troy.
- Maerten, B., Roose, D., Basermann, A., Fingberg, J. and Lonsdale, G. (1999) 'DRAMA: a library for parallel dynamic load balancing of finite element applications', *Proc. Ninth SIAM Conference on Parallel Processing for Scientific Computing*, San Antonio, TX.
- Mitchell, W.F. (1995) 'Refinement tree based partitioning for adaptive grids', *Proc. Seventh SIAM Conf on Parallel Processing for Scientific Computing*, SIAM, pp.587–592.
- Mitchell, W.F. (1998) 'The refinement-tree partition for parallel solution of partial differential equations', *NIST Journal of Research*, Vol. 103, No. 4, pp.405–414.
- MPI (1995) *A Message-Passing Interface Standard*, Message Passing Interface Forum, University of Tennessee, Knoxville, Tennessee, <http://www.mpi-forum.org/docs/mpi-11-html/mpi-report.html>.
- MPI-2 (1997) *Extensions to the Message-Passing Interface*, 1st ed., Message Passing Interface Forum, University of Tennessee, Knoxville, Tennessee, <http://www.mpi-forum.org/docs/mpi-20-html/mpi2-report.html>.
- Norton, C.D., Lou, J.Z. and Cwik, T. (2001) 'Status and directions for the PYRAMID parallel unstructured AMR library', *Eighth MI Workshop on Solving Irregularly Structured Problems in Parallel (15th IPDPS)*, IPDPS, San Francisco, CA.

- Pinar, and Hendrickson, B. (2001) 'Communication support for adaptive computation', *Proc. 10th SIAM Conf. Parallel Processing for Scientific Computing*, Portsmouth, VA, March.
- Remacle, J-F. and Shephard, M.S. (2003) 'An algorithm oriented mesh database', *Int. J. Numer. Meth. Engng.*, Vol. 58, pp.349–374.
- Remacle, J-F., Klaas, O., Flahery, J.E. and Shephard, M.S. (2002) 'Parallel algorithm oriented mesh database', *Eng. Comput.*, Vol. 18, No. 3, pp.274–284.
- Salinger, A., Devine, K., Hennigan, G., Moffat, H., Hutchinson, S. and Shadid, J. (1996) *MPSalsa A Finite Element Computer Program for Reacting Flow Problems, Part 2 – User's Guide*, Sandia National Laboratories, Albuquerque, NM, Tech. Report SAND96-2331.
- Schroeder, W., Martin, K. and Lorensen, B. (2003) *The Visualization Toolkit: An Object-Oriented Approach to 3D Graphics*, 3rd ed., Kitware, Inc., <http://public.kitware.com/VTK>.
- Segal, M. and Akeley, K. (2003) *The OpenGL Graphics System: A Specification (Version 1.5)*, Silicon Graphics, Inc., <http://www.opengl.org/developers/documentation/specs.html>.
- Seward, J. and Nethercote, N. (2003) *Valgrind, stable release 20031012*, <http://devel-home.kde.org/~sewardj/>.
- Taylor, V.E. and Nour-Omid, B. (1994) 'A study of the factorization fill-in for a parallel implementation of the finite element method', *Int. J. Numer. Meth. Engng.*, Vol. 37, pp.3809–3823.
- The Approved Licenses*, Open Source, links to numerous open source licenses at <http://www.opensource.org/licenses/index.php>.
- Walshaw, C. (2002) *The ParallelJOSTLE Library User's Guide, Version 3.0*, University of Greenwich, London, UK.
- Walshaw, C., Cross, M. and Everett, M.G. (1997) 'Parallel dynamic graph partitioning for adaptive unstructured meshes', *J. Parallel Distrib. Comput.*, Vol. 47, No. 2, pp.102–108.