Trustless Grid Computing in ConCert

Bor-Yuh Evan Chang Karl Crary Margaret DeLap
Robert Harper Jason Liszka Tom Murphy VII
Frank Pfenning

June 2002
CMU-CS-02-152

School of Computer Science
Carnegie Mellon University
Pittsburgh, PA 15213

Also submitted to the Third International Workshop on Grid Computing.

Abstract

We believe that fundamental to the establishment of a grid computing framework where all (not just large organizations) are able to effectively tap into the resources available on the global network is the establishment of trust between grid application developers and resource donors. Resource donors must be able to trust that their security, safety, and privacy policies will be respected by programs that use their systems.

In this paper, we present a novel solution based on the notion of *certified code* that upholds safety, security, and privacy policies by examining *intrinsic* properties of code. Certified code complements authentication and provides a foundation for a safe, secure, and efficient framework that executes native code. We describe the implementation of such a framework known as the ConCert software. Finally, we discuss a sample application that runs on ConCert.

The ConCert Project is supported by the National Science Foundation under grant ITR/SY+SI 0121633: "Language Technology for Trustless Software Dissemination".

Keywords: grid computing, security, certification, certified code, grid architecture, ConCert

1 Introduction

In recent years, numerous organizations have been vying for donated resources for their grid applications. Potential resource donors are inundated with worthwhile grid projects such as discovering a cure for AIDS, finding large prime numbers, and searching for extraterrestrial intelligence. Part of the difficulty in obtaining resources is establishing trust between the grid application developer and the resource donors. Because resource donors often receive little or no direct reward for their contributions, they demand assurances of safety, security, and privacy to protect themselves from malicious as well as simply unreliable software. In an ideal grid framework where users are provided the abstraction of a single virtual machine that automatically distributes work and gathers results as proposed in Legion [GW97], this problem becomes even more salient as the use of donated resources is automatic.

Current grid frameworks, such as the Globus toolkit [FK99], focus on *authentication* to provide the basis for security on computational grids. Authentication provides a means for one entity to verify the identity of another. If this is combined with some form of *access control*, then resource donors are able to control *who* can access their resources. In essence, this seeks to address the question "how do I identify those I trust?"; however, there seems to be a more fundamental question: "whom can I trust?" Currently, resource donors are limited to relying on the reputation of the developer or possibly the reputation of some quality assurance team that endorses him to answer this question of trust. While this is a reasonable solution for large well-known projects, we would like a more automated means that would enable more people to utilize the grid.

To address this issue, the ConCert project [Con01] enforces safety, security, and privacy policies by verifying *intrinsic* properties of code. This is realized through the use of *certifying compilers* that allow software developers to produce efficient machine code along with checkable certificates that can be easily and automatically verified on behalf of the resource donor. Our vision is to empower *all* to utilize donated resources by establishing trust via rigorous, mathematical proof of the intrinsic safety, security, and privacy properties of software.

2 Certification and Security

In this section, we discuss current means of ensuring safety and security without authentication, provide an overview of the certified code technologies we employ, and discuss properties that can be certified.

2.1 Related Safety and Security Ideas

Mechanisms for safeguarding against faulty code produced by trusted users have been needed in systems long before the concern over malicious mobile code. For example, one method to safeguard memory between threads of computation is to simply place the processes in separate address spaces, as is done in most

modern-day operating systems. This and other operating system mechanisms provide a very coarse level of safety or fault isolation both for individual systems and any grid framework. The problem with basing safety solely on these coarse mechanisms is that there is very little control on what properties are enforced and there is often a large runtime overhead associated with these mechanisms (*e.g.* frequent context switching and copying between addresses spaces). Several technologies have been developed to lower the cost of fault isolation and provide finer control over which safety, security, and privacy properties are ensured. This, in essence, is also the goal of certified code. In this section, we describe a few of these related technologies, namely virtual machine techniques and methods that interpose between the process and the operating system, along with some shortcomings that we believe certified code addresses.

Virtual Machine Methods. The most prominent use of a virtual machine environment for providing a secure execution space is in the Java Development Kit (JDK). At the core of the JDK 1.2 security model [GMPS97] is the ability to type-check Java Virtual Machine (JVM) bytecode [LY99]. It is critical that untrusted code cannot bypass the runtime checks that enforce the high-level security policy. In addition, the JDK has authentication-based mechanisms for identifying the origin of code and restricting access based on this information.

Several grid frameworks have been developed using Java as a secure host for mobile code, such as IceT [GS99], Javelin [CCI+97], and Bayanihan [SH99]. The IceT project aims at allowing processes running in the IceT environment to freely flow amongst available computational resources. IceT code primarily consists of Java *application* bytecode that is dynamically loaded into a governing environment with some optional policy-based means to link with native code. Javelin focuses on creating a grid framework using Java-enabled web browsers to minimize the technical expertise needed to participate in grid computing. Code for Javelin must be in the form of Java *applets* rather than Java applications. Project Bayanihan also aims to utilize Java applets to make it easier to participate but also supports the loading of Java applications. Further, it investigates preventing *sabotage* by nodes who submit erroneous results.

At the same time, despite the additional safety and security guarantees afforded by a virtual machine environment for loading code, a number of grid frameworks load native code for performance reasons. Although concerns about the performance of interpreted Java bytecode are somewhat ameliorated by just-in-time (JIT) compilers, JIT compilation to native code occurs *after* bytecode verification, so errors in the JIT compiler may lead to security holes. In addition, since the JIT compiler is run at execution time, its compilation process must be fast, which limits the quality of code it is able to produce.

Interposition Methods. Even prior to grid computing, interposition mechanisms between an untrusted process and the operating system were used to provide finer control of safety properties. Interposition gives control of execution to a watchdog whenever the untrusted process performs a possibly unsafe

operation. One interposition method is known as Software-based Fault Isolation (SFI) [WLAG93], which is exhibited in the Omniware system for mobile code [LSW95]. One striking issue with interposition mechanisms is that they are limited in what properties they ensure. For example, SFI only ensures only memory safety. As with virtual machine methods, performance of interposition methods is often also a concern. In grid applications where we seek to make the best use of donated resources, avoiding the overhead of these methods is desirable.

Proof-Carrying Code. *Proof-carrying code* (PCC) is a form of certified code with arbitrary certification properties written as logic statements [Nec97]. In PCC, a *certifying compiler* generates a proof along with the object code, and this proof is verified by a simple proof-checker on the code recipient's computer. PCC has very little runtime overhead because many properties can be verified before the program is ever run. For instance, Necula and Lee show that PCC handily outperforms SFI on a network packet filter application [NL96].

Because the certified properties are, in principle, arbitrary, PCC is highly flexible. However, this flexibility also means that it can be difficult for programmers to know when the certifying compiler will be able to generate proofs. One response to this problem is to encode desired safety properties in a *type system* for the source programming language. In this sense, a type system is simply a syntactic realization of specific properties such that a well-formed program is guaranteed to have those properties. Verification then consists of type-checking rather than proof-checking. A well-designed type system makes it easy for the programmer to understand which properties are certifiable by a compiler.

Many properties can be encoded in a type system; however, all type systems rely on the basic notion of *type safety*. Type safety means that the program will not "go wrong" by violating abstraction boundaries, accessing memory outside its address space, or branching into unsafe code. Therefore, for our first version of a grid framework with certified code, we choose languages with just this basic property.

2.2 Enabling Technologies

In this section, we provide some background on the enabling technologies that allow us to develop a grid framework based on the notion of certified code. The implementation of our grid framework builds on the TALx86 realization [MCG+99] of the Typed Assembly Language (TAL) [MWCG99, MCGW02] developed by Morrisett *et al.* A certifying compiler for a type-safe C-like language called Popcorn that compiles to TALx86 has also been developed as part of the TAL project. These tools serve as a foundation for an implementation of the ConCert grid software and an example application that runs on it as described in later sections.

An Overview of Popcorn. The Popcorn programming language looks very much like C except without unsafe features like the unrestricted address-of op-

erator, pointer arithmetic, and pointer casts. At the same time, it has a number of advanced features akin to modern high-level programming languages like Java [GJSB00] and Standard ML (SML) [MTHM97], such as exceptions, garbage collection, tagged unions, and parametric polymorphism that mitigate the need for these unsafe features.

The primitive types supported by Popcorn include `bool`, `char`, `short`, `int`, `float`, `double`, `string`, and `unsigned` forms of the numeric types. Unlike C, arrays carry their size for bounds-checks, and strings are not null-terminated. Instead, strings are treated like arrays of `char`s, and a special `size` construct is used to extract the size of an array or string. Popcorn's basic control flow constructs (`if`, `for`, `while`, `do`, `break`, and `continue`) behave identically to their C counterparts except that test expressions must have type `bool`.

The aggregate data structures in Popcorn are similar to ones in high-level languages like Java and SML. First, Popcorn supports tuple types in addition to `struct`s and `union`s. Tuples (as well as `struct`s and `union`s) are created using the `new` construct and projected using `.1`, `.2`, ... as shown in figure 1A. There are two forms of structure definitions: `struct` and `?struct`. Values of types defined with `struct` cannot be `null` (a primitive in Popcorn) whereas values of types defined with `?struct` may. A value of a type defined with `?struct` is checked for `null` upon access to a field. If it is `null`, the program aborts immediately. Unions in Popcorn resemble SML datatypes more than C unions in that each variant has a tag and an associated type. In figure 1B, we declare a full binary tree with integer data at the leaves.

```
(A)     *(int,int) p   = new (0,1);
        int        sum = p.1 + p.2;
(B)     union tree { int Leaf; *(tree,tree) Node; }
(C)     int numleaves(tree t) {
          switch (t) {
            case Leaf(x):      return 1;
            case Node *(l,r):  return numleaves(l) + numleaves(r);
          }
        }
(D)     *(t2,t1) swap<t1,t2>(*(t1,t2) x) {
          return new (x.2,x.1);
        }
```

Fig. 1. Popcorn examples.

Notice that `union` types may be recursive. We often utilize values of a type defined with `union` by `switch`ing on them. For example, we can write a function that counts the number of leaves in a binary tree (figure 1C).

Parametric polymorphism provides a means to write data structures or algorithms where the use of values is independent of the types. The syntax of parametric polymorphism in Popcorn resembles that of templates in C++ or generic types in GJ [BOSW98]. However, the implementation is different from that of C++ templates. Popcorn is compiled using macro expansion so that any checking is done on all instances rather than on the template itself. In figure 1D,

File	*Produced By*	*Description*	*Shipped*
`main.pop`	developer	the Popcorn source file for `main`	
`main.tal`	`popcorn`	typed assembly language output	
`main_i.tali`	`talc`	`main`'s imports (any `extern` declarations)	√
`main_e.tali`	`talc`	`main`'s exports (non-static types and values)	√
`main.to`	`talc`	binary file with the typing annotations	√
`main.o`	`talc`	native object file in ELF	√

Table 1. TALx86/Popcorn Files for `main`.

we use parametric polymorphism to write a generic swap of the components of a pair. The symbols `t1` and `t2` are type variables that represent arbitrary types.

From this overview, we see the similarity of Popcorn to Java in that it has the look and feel of C but with strictures that enable the demonstration of safety properties. However, Popcorn differs from Java in that the output of the Popcorn compiler is machine code able to be run at full speed rather than bytecode that is interpreted or that requires just-in-time compilation upon execution.

An Overview of TALx86. TALx86 is a statically-typed assembly language for the Intel IA-32 architecture. Since it is not feasible to fully describe TALx86 here, we simply provide an overview of TALx86 in order to appreciate the underlying technology on which our grid framework builds. Further details about TALx86 can be found in Morrisett *et al.* [MCG+99] as well as about its theoretical basis in related works [MWCG99, MCGW02].

From the high-level source language, the Popcorn compiler (or potentially other certifying compilers) produce `.tal` files that contain TALx86 assembly with all the typing annotations for each source file along with a number of other files that are generated by the TALx86 assembler `talc`. Some of these files are shipped to the code recipient. These contain information such as the imports, exports, and typing annotations that are used for link-time verification. A `.tal` file is a realistic assembly language that is, in fact, compatible with the Microsoft Macro Assembler (MASM). Table 1 summarizes these files for an example program called `main` that is written in Popcorn.

The TALx86 assembler `talc` translates a `.tal` file into a native object file in either COFF (for Windows) or ELF (for Unix variants) and a `.to` file that contains typing annotations from the original `.tal` file in a binary format. The `talc` assembler is actually composed of a TALx86 type-checker, a link-verifier, a code assembler, and a code linker. Upon reading of `.tal` files, `talc` first type-checks each file individually. This type-check should never fail provided that the implementation of the certifying compiler is correct; however, by type-checking, we no longer need to assume the correctness of the certifying compiler. Contrast this with using the JVM with just-in-time compilation for ensuring type-safety where we have to assume that the JIT compiler is implemented correctly. Before creating the native object file, the link-verifier ensures that the multiple `.tal` files are safe after being linked together by verifying that they have the same assumptions about the values and types that they share. Technical details about link verification can be found in Glew and Morrisett [GM99].

The ConCert grid software simply ships the native object, `.to`, `_i.tali`, and `_e.tali` files. Before dynamically loading the code on the donor host, the grid software type-checks the code using a TAL verification library on the received files.

Certifiable Properties. Type safety is the key property that allows us to control how a program behaves. As a baseline, a type-safe program must be memory safe (no illegal reads or writes) and control-flow safe (no jumps to illegal addresses). In addition, it cannot violate abstraction boundaries that could lead to such errors. We are further able to provide the distributed code with an arbitrary "safe" subset of the system library. It is simple, for instance, to allow the code to modify or delete only files that it created (without run-time checks), or to simply not allow access to any files at all. The grid volunteer can be given a choice of several policies on these sorts of issues, each of which is enforced by verifying that the candidate code type-checks with the supplied view of the system library.

There are, however, other properties that we may wish to certify that we are investigating in ongoing research and that have not yet been implemented. To express these, we enhance our type systems so that type safety implies adherence to the property in question. For instance, one primary concern for grid computing is the resource usage of the distributed code. Much work is in progress regarding *resource bound certification* [NL97, CW00], which allows the type system to include bounds on how much CPU, memory, and disk resources a piece of code will use.

Finally, users may wish to give code access to private information (such as the computer's configuration), so long as this information does not make its way back onto the network. Such requirements are known as *information-flow* properties, and work is being carried out on type systems that certify them [HR98, ZM00, Mye99].

3 ConCert Architecture

The main feature of our architecture is decentralization. A decentralized grid architecture allows us to use idle resources while avoiding bottlenecks that could overload hosts or networks and degrade performance—possibly even for users who are not involved in grid computation. The need for failure tolerance also has major effects on our programming and scheduling models. At present, our architecture is completely peer-to-peer. Although this may make our network topology as a whole more fault-tolerant, whether it is the best model overall requires further investigation. In the future, we expect to consider whether specialized nodes for peer discovery or for code and result caching would be beneficial.

Every participating node in our network has three components: a *locator*, a *worker*, and a *conductor*. The *locator* finds a host's peers at runtime as is necessitated by the decentralized architecture. Every participating node in our

network can both serve work and run it. The *conductor* component keeps track of what work is available to be done and bundles files as necessary to provide work to peers. The *worker* acquires code to run, **verifies** it using the TAL verifier, and runs it. We now provide an overview of each of these components. Further details can be found in DeLap's honors thesis [DeL02].

Node Discovery. Because of decentralization participants should find their peers dynamically. In our current implementation of the *locator* component, we use a protocol similar to Gnutella's [Cli00] to discover nodes. Using this protocol, hosts build tables of peers as they run. As in Gnutella, a host that wishes to participate joins the network by sending an initial ping message to some number of predetermined peers. If any such peers are participating, they will forward that message to the peers of which they are aware, and so forth within a limited number of hops. Anyone receiving a ping message notifies the originating host that it is alive by sending a response.

We do not share all of Gnutella's privacy goals, so we are able to reduce network traffic for certain exchanges. Nonetheless, since it is known that Gnutella may have scalability issues [Rit01], we may wish to change our method of node discovery in the future. Components other than the locator remain unaware of how the list of contacts is found. Therefore, it would be feasible for us to change our entire network topology without affecting the serving and running of work itself.

Parallelism Model. Since the grid is likely to contain slow, unreliable network links and hosts, we do not support the use of either fine-grained, high-communication threads or shared memory, both for performance reasons and for failure tolerance. At the same time, our model of parallelism should allow for inexpensive scheduling. We therefore base our model of parallelism on that of Cilk-NOW [BL97] and, more generally, dataflow computer architectures [CSG99]. In our model, programs are split into segments called *cords*, on which we impose the three invariants. Although we do not at present enforce these invariants via certification, they are necessary for clients to write correct and reliable code for our framework. In future work we plan to consider if they can be verified, not as a protection of the resource donors, but as an aid to the application developers.

First, once a cord is ready to run, it is able to execute continuously to completion without waiting for data from other cords or otherwise blocking. In other words, a cord does not communicate with other cords while executing. This makes it feasible, where failure due to the grid's unreliability makes rescheduling necessary, to reschedule cords independently from one another.

Second, any possible run of a cord is, as far as the developer is concerned, "as good as" any other. This is evident for deterministic cords. In the case of nondeterministic or randomized algorithms, the programmer must carefully consider how the program will map to cords (if indeed it will map to cords well), since subsequent runs of a nondeterministic cord may produce different results. Again, multiple runs of the same cord may be necessary in the case of failure.

Finally, cords do not produce outside effects on which other cords rely. For example, it would be unsound for one cord to write to a file on a participating machine and another cord to depend upon the contents of that file. Such effects are forms of communication with other cords. As such, they would be hidden from the scheduler; if any cords were later rescheduled, they would almost certainly not behave correctly.

How, then, does information travel between the parts of a program? Each cord's I/O consists of the *arguments* to the segment of the entire program it contains and the *result* it produces. We may therefore represent a program as a graph in which nodes denote cords and edges indicate the data flowing into them as arguments. A cord is not ready to run until its inward edges have the necessary data available. Figure 2(A) gives a simple example of this model.

To make our model more powerful, we collect sets of dependencies into groups of *and*- and *or*-dependencies. In an *and*-dependency set, all of the dependencies are required—they must all contain data before the set of dependencies is considered complete. Alternatively, in an *or*-dependency set, only one dependency result is necessary. As soon as any one of the dependencies is filled, that set is considered completed. Such sets can be chained into trees of dependencies that are to be simplified and collapsed as dependency data (results) are filled in. Figure 2(B) provides a simple example of a dependency tree.

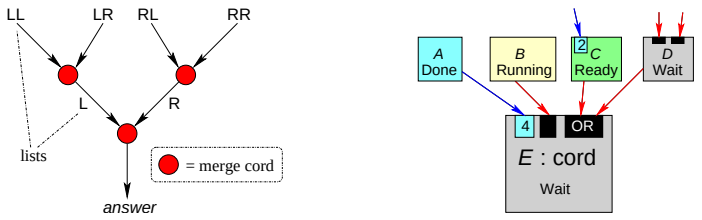

(A) (B)

Fig. 2. (A) A simple cord graph example. This graph represents an execution of a program that mergesorts four sorted list segments into one sorted list. The list segments, some of which are themselves results of cords, are arguments to the cords. Merged list segments travel downward in the graph until the final result (the entire sorted list) is obtained. (B) Dependency resolution. Here, cord A has recently terminated; its result, 4, becomes one of the arguments to E. Cord E now awaits results from B and either C or D. Meanwhile, cord C is ready to run, having received the result 2 from the cord on which *it* depends. To summarize, the set $C + D$ is a set of *or*-dependencies for E; $AB(C + D)$ is its overall set of *and*-dependencies.

Work Stealing. Whenever a participating workstation becomes idle, it may pull work from its peers. Hosts never try to foist work on others. We have three important reasons for this.

First, we assume that the worker itself knows best when the host it is running on is idle. If peers try to send out work without knowing which machines are

idle, they will contact hosts uselessly and may irritate users of non-idle machines by loading their networks. Alternatively, trying to keep track of which peers are idle would require more protocol messages, and again, more network traffic. For all that, it would still likely be less foolproof than letting the host itself decide when to run code.

Second, this enables us to customize the amount of work sent to worker hosts. Since worker hosts are doing the most computationally intensive work, we want to control how much work they do at any given time. For example, a machine with one processor will probably want to request only one cord at a time, whereas a multiprocessor machine may request more. Conversely, hosts serving out work will generally be happy to have connections to many peers willing to do work for them. Work-stealing allows worker machines to actively request the exact amount of work they want at any given time.

Third, having workers make the requests allows conductors to match security policies in advance. In our work-stealing model, a host can send a representation of its security policy along with its requests for work. Hosts serving work can check this policy against available work *before* sending it out over the network so as to avoid serving code the worker will not run.

Work stealing has major implications for scheduling. In our case, it means the conductor's scheduler can run *on demand only*, *i.e.*, whenever it receives a work request or result from a worker. Since users may wish to serve cords even when their machines are not idle, reducing the scheduler's use of resources is especially advantageous.

To steal work from a conductor, a worker initiates a sequence of messages to acquire the appropriate cord and its arguments. It then verifies and runs the code, sending the result back to the conductor. If, on the other hand, the code fails to verify, the worker does not run it.

Failure Tolerance and Recovery. Of course, computation will not always proceed smoothly on a failure-prone network such as the grid. Given that grid computing attempts to solve large problems on possibly unreliable networks, tolerance of failure is important. In particular, we need to be able to checkpoint programs at some reasonable granularity and to restart them with partial results, so that earlier computation is not lost if some part of the program fails (*e.g.* due to a downed node). Since cords do not communicate among themselves as they run, restarting them is not especially difficult provided the code for the cord in question is kept. With respect to checkpointing, results are cached on conductors so that if a cord *is* rescheduled, preceding cords on which it depends do not have to be re-run.

Failure detection is also an issue. If a worker notices and notifies the corresponding conductor that its cord has died or failed to match its host's security policy, the conductor can simply reschedule it and hand it off to any subsequent requester. Otherwise, we need to decide if and when to reschedule apparently failed cords. In some cases, grid programmers may wish to approximate the result of a failed cord rather than reschedule the cord itself. We expect that

failure detection will involve some form of heartbeat between cords and originating conductors, but further work on implementation is required. Related work has already been done on failure detection for distributed computing, including [CT96].

Our framework also entails a new mode of failure, where a binary does not pass the type-checking phase once transmitted to the worker's hosts. There are several reasons such failure may occur. For example, it may fail because the certificate does not match the code, due to corruption of the files, a bad copy of the certifying compiler, malice, or similar causes. Second, it may fail because, while correct, it does not match the host's security policy. Presumably this would occur only if the conductor serving the code did not check the worker's policy against it in advance, or if the worker host's policy changed in the midst of the negotiation. In those cases, however, explicit failure notices can be generated to reduce the impact on the overall grid performance.

Grid Clients. External clients can connect to conductors on local machines to submit work and view or interpret results. To do this, they use a socket interface to send messages to the conductor specifying the cord and the arguments with which to run it. If all goes well, the conductor will return a handle to that cord, which may be used to query its status and retrieve its result through the same socket interface. On multiuser systems, this socket should be protected with appropriate file permissions. Using a socket allows for easy interaction with clients written in various languages—we have written demonstration clients in C and Standard ML. The ConCert *protocol* software is also language-independent. Note, however, that the cords themselves must be available in the TAL-produced format of an object file and typing annotations, and the verifier clearly must check this format.

4 Sample Application: Ray Tracer

Because of the strong interplay between certification technology and the source programming language, we are interested in how applications are developed for the grid and especially in what programming issues arise. Therefore we have developed a sample application for our grid framework: a parallel ray tracer called Lightharp. The ray tracer back-end, which does the actual tracing work (and is distributed across the grid), is written in the Popcorn language described in section 2.2. The front-end, which divides the scene into small chunks, submits the tracing jobs, and displays the results, is a separate program written in Standard ML.

Lightharp implements the specification for the ICFP 2000 Programming Contest [ICF00]. Scenes to be rendered are described by a simple stack-based language called GML. GML supports basic diffuse and specular lighting, constructive solid geometry, and several advanced features such as procedural textures. The GML representation of a scene is typically quite compact because it can create named objects and duplicate them to populate the scene. Therefore,

rather than create an intermediate representation of the scene we communicate a (slightly) modified version of the original GML code on the grid.

The back-end is a function of type `string -> string`, which takes a GML program as input, and returns the rendered scene as a sequence of RGB colors. The color of a particular pixel is computed by shooting a "ray" from the point of view through the image plane and recursively tracing its path through the scene using standard ray-tracing algorithms. Note that the color of each pixel or block of pixels can be calculated independently, which is the only source of parallelism in our implementation.

The front-end works by parsing the GML scene and modifying it to instruct the back-end to render only a small part of the image, that is, only trace rays through a certain range of pixels. It then submits cords into the grid to cover the entire image, and simply waits for the results to arrive (currently by polling the ConCert software).

For an application with such a simple parallelism model as a ray tracer, we found that this implementation strategy was adequate. Popcorn is a powerful enough language to implement the back-end tracing functions without much pain. It is also relatively easy to manage the cords from the front-end and only mildly tedious to manually marshal between strings and RGB data.

Our ray tracer application's parallel behavior is extremely simple; for instance, the back-end cords have no need to spawn other cords, nor do they have any dependencies. Soon, we intend to tackle some more difficult parallel computations, such as parallel game tree search in chess. We hope to push the limits of what is possible with cords, and discover interesting research problems in programming techniques for the grid.

5 Conclusion

We have presented a framework and technologies for grid computing in a trustless setting based on the idea of certified code. Certification compares favorably to similar technology; it verifies rich properties of native code with low run-time overhead. We have also described a peer-to-peer network for fetching and running work, and presented a sample ray-tracing application that runs on the grid.

Future Work. At present, we certify only type safety (which entails memory safety and control flow safety). In section 2.2, we discussed other properties that might be certified. When more certification options are available, we will need to encode policies in such a way that code can be tested against them easily and efficiently. We also plan to provide an accessible interface for users to specify their security policies in a transparent manner.

Our framework protects cycle donors from broken or malicious code, but it does not protect developers from false answers. Malicious workers might, instead of actually running the code presented, return a fabricated result. We plan to investigate how techniques proposed elsewhere [Sar98, SH99] can be adapted to the ConCert architecture.

So far, we have concentrated on certification of safety properties. In practice, this should be combined with methods for peer-to-peer authentication. Here, again, the ideas behind proof-carrying code provide novel solutions [AF99]. We plan to investigate how they may be applied in the ConCert framework.

Although we have implemented Lightharp (a ray tracer application) using the ConCert grid software, we noted in section 4 that the parallel structure of Lightharp is exceedingly simple. To support more sophisticated uses of parallelism, we will need a high-level programming model for supporting some notion of starting remote computation and gathering of results *within* a grid application. Consequently, we will need the means to map high-level programming languages to the simplistic low-level interface provided by ConCert. Although it has yet to be demonstrated, we believe we can leverage several programming language techniques to achieve this goal. Preliminary work in this direction is discussed in Chang's honors thesis [Cha02] and a companion paper on work in progress [CDL+02].

Acknowledgments. Our implementation and protocols, especially the work-stealing protocol and code to support TAL verification, are based largely on Joshua Dunfield's initial implementation of the ConCert framework. We would also like to acknowledge the other members of the ConCert project and Guy Blelloch for their helpful comments.

References

[AF99] Andrew W. Appel and Edward W. Felten. Proof-carrying authentication. In G. Tsudik, editor, *Proceedings of the 6th Conference on Computer and Communications Security*, pages 52–62, Singapore, November 1999. ACM Press.

[BL97] Robert D. Blumofe and Philip A. Lisiecki. Adaptive and reliable parallel computing on networks of workstations. In *USENIX 1997 Annual Technical Conference on UNIX and Advanced Computing Systems*, pages 133–147, Anaheim, California, 1997.

[BOSW98] Gilad Bracha, Martin Odersky, David Stoutamire, and Philip Wadler. Making the future safe for the past: Adding genericity to the Java™ programming language. In *Object Oriented Programming: Systems, Languages, and Applications (OOPSLA)*, pages 183–200, Vancouver, British Columbia, October 1998.

[CCI+97] Peter Cappello, Bernd Christiansen, Mihai F. Ionescu, Michael O. Neary, Klaus E. Schauser, and Daniel Wu. Javelin: Internet-based parallel computing using Java. In *ACM Workshop on Java for Science and Engineering Computation*, Las Vegas, Nevada, June 1997.

[CDL+02] Bor-Yuh Evan Chang, Margaret DeLap, Jason Liszka, Tom Murphy VII, Karl Crary, Robert Harper, and Frank Pfenning. Towards a functional library for fault-tolerant grid computing. Submitted to GRID 2002, June 2002. To appear on the ConCert project webpage, URL: http://www.cs.cmu.edu/~concert.

[Cha02] Bor-Yuh Evan Chang. Iktara in ConCert: Realizing a certified grid computing framework from a programmer's perspective. Technical Report CMU-CS-02-150, Carnegie Mellon University, June 2002. Undergraduate honors thesis.

[Cli00] Clip2 Distributed Search Services. The Gnutella protocol specification v0.4, September 2000. URL: http://www.gnutella.co.uk/library/pdf/gnutella_protocol_0.4.pdf.

[Con01] ConCert. Certified code for grid computing, project webpage, 2001. URL: http://www.cs.cmu.edu/~concert.

[CSG99] David E. Culler, Jaswinder Pal Singh, and Anoop Gupta. *Parallel Computer Architecture: A Hardware/Software Approach*. Morgan Kaufmann, San Francisco, California, 1999.

[CT96] Tushar Deepak Chandra and Sam Toueg. Unreliable failure detectors for reliable distributed systems. *Journal of the ACM*, 43(2):225–267, March 1996.

[CW00] Karl Crary and Stephanie Weirich. Resource bound certification. In *Twenty-Seventh ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, pages 184–198, Boston, Massachusetts, January 2000.

[DeL02] Margaret DeLap. Implementing a framework for certified grid computing. Technical Report CMU-CS-02-151, Carnegie Mellon University, June 2002. Undergraduate honors thesis.

[FK99] Ian Foster and Carl Kesselman. The Globus toolkit. In Ian Foster and Carl Kesselman, editors, *The Grid: Blueprint for a New Computing Infrastructure*, chapter 11, pages 259–278. Morgan Kaufmann, San Francisco, California, 1999.

[GJSB00] James Gosling, Bill Joy, Guy Steele, and Gilad Bracha. *The Java™ Language Specification*. Addison-Wesley, second edition, 2000.

[GM99] Neal Glew and Greg Morrisett. Type-safe linking and modular assembly language. In *Twenty-Sixth ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, pages 250–261, San Antonio, Texas, January 1999.

[GMPS97] Li Gong, Marianne Mueller, Hemma Prafullchandra, and Roland Schemers. Going beyond the sandbox: An overview of the new security architecture in the Java Development Kit 1.2. In *USENIX Symposium on Internet Technologies and Systems*, Monterey, California, December 1997.

[GS99] Paul A. Gray and Vaidy S. Sunderam. Metacomputing with the IceT system. *International Journal of High Performance Computing Applications*, 13(3):241–252, 1999.

[GW97] Andrew S. Grimshaw and William A. Wulf. Legion: The next logical step toward the world-wide virtual computer. *Communications of the ACM*, 40(1):39–45, January 1997.

[HR98] Nevin Heintze and Jon G. Riecke. The SLam calculus: Programming with secrecy and integrity. In *Twenty-Fifth ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, pages 365–377, San Diego, California, January 1998.

[ICF00] ICFP. The third annual ICFP programming contest, 2000. URL: http://www.cs.cornell.edu/icfp/.

[LSW95] Steven Lucco, Oliver Sharp, and Robert Wahbe. Omniware: A universal substrate for web programming. In *Fourth International World Wide Web Conference*, pages 359–368, Boston, Massachusetts, December 1995.

[LY99] Tim Lindholm and Frank Yellin. *The Java™ Virtual Machine Specification*. Addison-Wesley, second edition, 1999.

[MCG+99] Greg Morrisett, Karl Crary, Neal Glew, Dan Grossman, Richard Samuels, Frederick Smith, David Walker, Stephanie Weirich, and Steve Zdancewic. TALx86: A realistic typed assembly language. In *1999 ACM SIGPLAN Workshop on Compiler Support for System Software*, pages 25–35, Atlanta, Georgia, May 1999.

[MCGW02] Greg Morrisett, Karl Crary, Neal Glew, and David Walker. Stack-based typed assembly language. *Journal of Functional Programming*, 12(1):43–88, January 2002.

[MTHM97] Robin Milner, Mads Tofte, Robert Harper, and David MacQueen. *The Definition of Standard ML (Revised)*. MIT Press, Cambridge, Massachusetts, 1997.

[MWCG99] Greg Morrisett, David Walker, Karl Crary, and Neal Glew. From System F to typed assembly language. *ACM Transactions on Programming Languages and Systems*, 21(3):527–568, May 1999.

[Mye99] Andrew C. Myers. Jflow: Practical mostly-static information flow control. In *Twenty-Sixth ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, pages 228–241, San Antonio, Texas, January 1999.

[Nec97] George C. Necula. Proof-carrying code. In *Twenty-Fourth ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, pages 106–119, Paris, France, January 1997.

[NL96] George C. Necula and Peter Lee. Safe kernel extensions without run-time checking. In *Second Symposium on Operating Systems Design and Implementation*, pages 229–243, Seattle, Washington, October 1996.

[NL97] George C. Necula and Peter Lee. Safe, untrusted agents using proof-carrying code. In Giovanni Vigna, editor, *Special Issue on Mobile Agent Security*, volume 1419 of *Lecture Notes in Computer Science*, pages 61–91. Springer-Verlag, October 1997.

[Rit01] Jordan Ritter. Why Gnutella can't scale. No, really., February 2001. URL: http://www.darkridge.com/~jpr5/doc/gnutella.html.

[Sar98] Luis F. G. Sarmenta. Bayanihan: Web-based volunteer computing using Java. In *Second International Conference on World-Wide Computing and its Applications*, pages 444–461, March 1998.

[SH99] Luis F. G. Sarmenta and Satoshi Hirano. Bayanihan: Building and studying web-based volunteer computing systems using Java. *Future Generation Computer Systems*, 15(5-6):675–686, 1999. Special Issue on Metacomputing.

[WLAG93] Robert Wahbe, Steven Lucco, Thomas E. Anderson, and Susan L. Graham. Efficient software-based fault isolation. In *Fourteenth ACM Symposium on Operating Systems Principles*, pages 203–216, December 1993.

[ZM00] Steve Zdancewic and Andrew C. Myers. Confidentiality and integrity with untrusted hosts. Technical Report 2000-1810, Cornell University, 2000.