Community Grids Laboratory

December 2004 Summary

1. Overview

1.1 Mission

The Community Grids Laboratory (CGL) was established in July 2001 as one of Indiana University's Pervasive Technology Laboratories. It was funded by the Lilly Endowment and located in the Indiana University Research Park (Showers) in Bloomington. Its staff includes Director Geoffrey Fox, 5 Research Associates, 1 Professional Staff and 24 PhD candidates. There are also temporary employees and consultants.

The Laboratory is devoted to the combination of the best technology and its application to important problems. We believe that e-Business and e-Science will grow in importance and imply global virtual communities. Our technology focus, Grids, are the enabling infrastructure for virtual organizations. Communities are an important virtual organization and integrating peer-to-peer network ideas and people into the Grid is a key feature of our work. Much of our innovative research exploits the observation that computers and networks are now so fast that one can use new and more transparent architectures and protocols and move from inflexible hardware to modular flexible software solutions. We also see a blurring of computers and the Network as our systems get more and more distributed.

We need to understand mix of open source, "university proprietary" and commercial software in modern distributed systems. Currently we expect "core" technologies to be open source with packaged "solutions" and support coming from industry including proprietary software. The globalized workforce (outsourcing) will give us challenges and opportunities in both developing software and as target virtual organizations.

1.2 Publications and Presentations

In the two and half years since the founding of the laboratory, members have published over 100 papers and reports. Over 50% of these have been in the 2003 reporting period. A similar number of major presentations were also given both as invited talks and as conference contributions. This reflects our emphasis on outreach and on encouraging students to submit conference papers and attend if they are selected. We have a very high acceptance rate due to the interest in and quality of our research.
The highlight of our 2003 publications was the book "Grid Computing: Making the Global Infrastructure a Reality" co-edited with Fran Berman (head of the NSF Supercomputer center NPACI in San Diego) and Tony Hey (Director of the core UK e-Science program). This provides an in-depth coverage of the internationally most important Grid computing activities and we made major contributions to 6 of the 43 chapters. We followed this up by a survey with a colleague David Walker of Grids in action as a so-called "Gap Analysis" which required interviewing 100 academic and industry Grid leaders to delineate next steps. The interest in this book and the Gap Analysis report has prompted us to plan writing two further Grid books – one in the next three months as a shorter more popular survey of Grids and e-Science. The second longer effort would be a distributed computing textbook based on Grid and Web technologies. We note in November of 2002, we published a seminal book on parallel computing -- *The Sourcebook of Parallel Computing* – with as co-authors the recognized leaders in the field.

1.3 Activities in 2003

These are covered in separate document but cover the following major areas: portals, Grid technology for messaging, peer-to-peer networks and cellular Grids, fault tolerant data transport, metadata and Semantic Grid, collaboration and systematic use of web services, universal access including linkage of PDA's to the Grid, collaboration and a new approach to building applications in an explicit MVC – Model, View, Controller paradigm. We applied these technologies to earthquake science and the SERVOGrid (Solid Earth Research Virtual Observatory Grid). We added over the last few months work with the Undergroundfilm.org non profit organization to develop educational material in the multimedia area. Fox is co-chair of the Semantic Grid and Grid Computing Environments Research Groups of the Global Grid Forum. He worked closely with the UK e-Science program where he is a member of their TAG – Technical Advisory Group – and helped set up the OMII – Open Middleware Infrastructure Institute. Fox has

continued as lead editor of the well regarded journal Concurrency and Computation: Practice and Experience.

4 students obtained Ph.D.'s during 2003; three from Florida State and one from Syracuse.

We vigorously pursued technology transfer with NaradaBrokering and HPJava software packages being available open source. HPJava provides high level parallel programming to Java and is a long term NSF funded project completed in 2003. The Anabas Corporation is interested in licensing several CGL technologies and obtained venture funding during 2003 from the Lilly grant. A joint CGL-Anabas proposal to the 21st Century fund focused on Reverse Outsourcing and its application to education, entertainment, sports and apparel industries.

1.4 Activities in 2004

We expect our work on portals to decrease as it is moving from a research to a production activity as indicated by CGL's leadership of the NSF funded National Middleware Initiative to harden portal technology for general deployment. We have several application areas where we hope to grow our collaborations and provide testbeds for the laboratory technology. These include particle physics with a new approved experiment looking for "Glueballs", crisis management (with the POLIS center at IUPUI and the openGIS consortium) and biocomplexity. CrisisGrid was the theme of a 21st Century grant proposal submitted to the State of Indiana in a Indianapolis, Bloomington and Purdue collaboration focused on flood management and emergency response. We expect our research in core Grid technology to grow with a focus on applications of our messaging environment NaradaBrokering to produce robust Grids. Most intriguing perhaps is new ideas on building client applications using Web services and explicit messaging. This could be important given the growing interest in alternatives to the traditional windows approach to clients.

We will continue our work with minority institutions – teaching a class spring 2004 over the Internet using our collaboration technologies to Jackson State University, an Historically Black University. We have designed an Education Grid with AIHEC (American Indian Higher Education Consortium) and Indiana's new School of Informatics. We intend to submit NSF proposals to support this new concept GESTALT – Grid for E-Science and Technology aimed at Learning and Teaching.

1.5 Contact Information

Geoffrey Fox: Phone 8122194643 email: gcf@indiana.edu
Community Grids Laboratory
501. N. Morton Street, Suite 224
Bloomington IN USA 47404-3730

http://www.infomall.org
http://grids.ucs.indiana.edu/ptliupages/

2. Grid Portals

Our major effort is the development of portals for users to use to access the capabilities of Grid and collaborative systems. We are collaborating with teams from Argonne National Laboratory, the University of Michigan, NCSA and the Texas Advanced Computing Center in developing core technology for this. The Grid supports a set of services providing a palette of capabilities for users. As shown on the left, we extended the Apache open source portlet model to provide a matching palette of user interface components that can be customized for each user and application domain. We have used our DoE, NASA, and NSF funding to develop specific application portals. Shown on left are our own portlet-based home page and the portals supporting nuclear fusion and earthquake simulation. The latter is an international collaboration with scientists from Australia, China and Japan sponsored by ACES (Asia-Pacific Cooperation for Earthquake Simulation). We are building a Grid iSERVO (International Solid Earth Research Virtual Observatory) linking their resources (portals, computers, sensors, databases). This work involves work in area of Grid approaches to GIS (Geographical Information Systems) where we have established local collaborations with the openGIS consortium. This area uses Semantic Grid technology and exploits our co-leadership of the Semantic Grid working group of the Global Grid Forum. This supports the addition of rich metadata to Grid services and we also are using these ideas to enrich descriptions of NASA satellite data to allow better use of it by researchers.

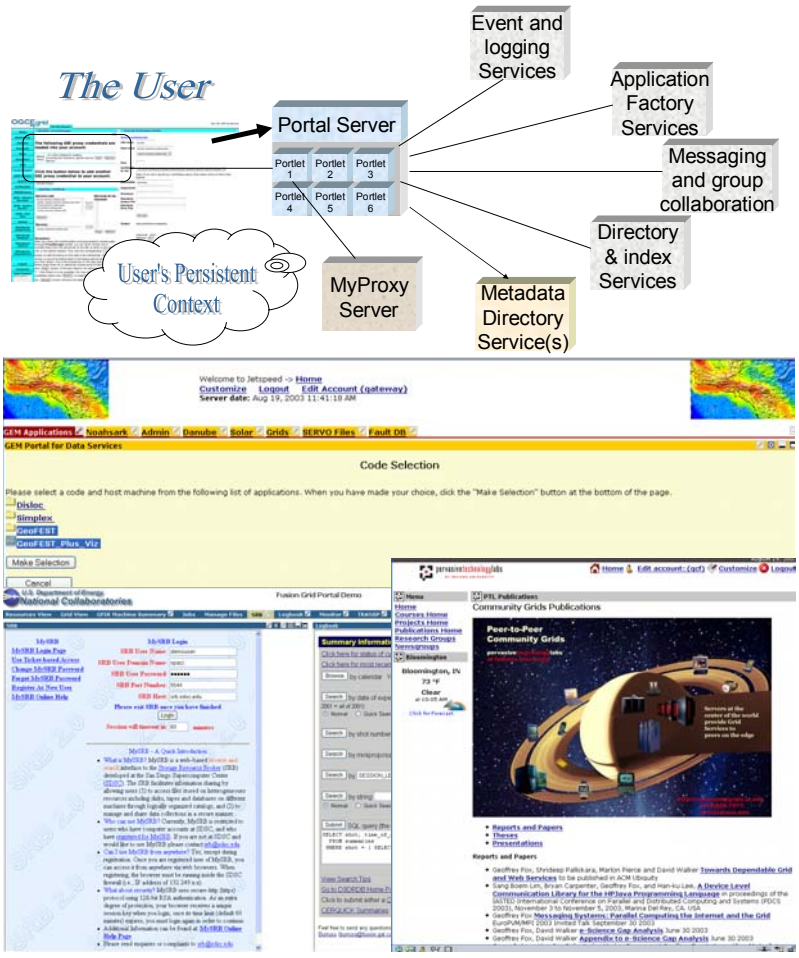

We have just started collaboration with a group – undergroundfilm.org – which is portal devoted to educational multimedia resources including 400 open source films. We are using this resource to test and refine our Grid technologies as its operation involves transferring several very large files. We are restructuring it using portlets and will provide collaboration support to allow discussion of the resources online. We expect multimedia applications to drive much of the commercial Grid deployment.

3. NaradaBrokering

NaradaBrokering is an open source technology supporting a suite of capabilities for reliable/robust flexible messaging; given the message based service architecture of Grids, this project is aimed at providing for the transport of messages between services and between services and clients. NaradaBrokering is designed around a scalable distributed network of cooperating message routers and processors.

Special features available in the current release include

- Publish-subscribe metaphor with general XML based topics
- Support for multiple protocols including TCP/IP (both blocking and non-blocking), UDP, Multicast, SSL, HTTP, RTP, HHMS (optimized for PDA and cell-phone access) and GridFTP with protocol chosen independently at each link
- The protocols can be chosen to tunnel through many firewalls and authenticating proxies.
- Interface to reliable storage and guaranteed order-preserving message delivery
- Support for message compression and decompression.
- Fragmentation and Coalescing of large files (>1 Gigabytes).
- JMS compliance and support for JXTA interactions.

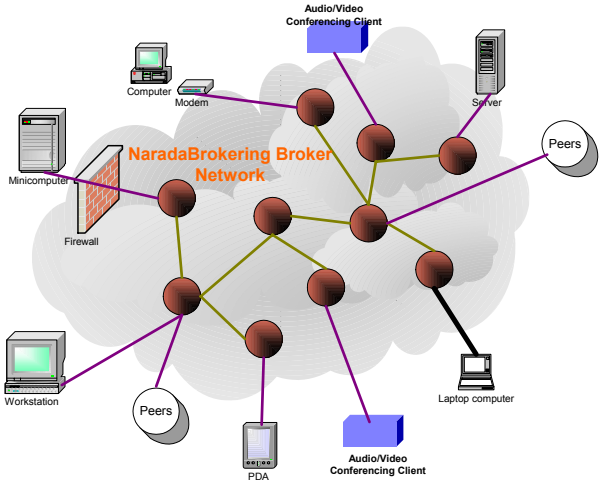

Functionality	WebSphereMQ	Pastry	NaradaBrokering
Max number of nodes	Medium	Very large	Very large
Robust Messaging	Yes	Yes	Yes
Robust delivery of >1 Gigabyte files	Yes	No	Yes
Support for routing P2P Interactions	No	Yes	Yes (JXTA)
A/V Conf. and raw RTP clients	NA	NA	Yes
Proxies/firewalls tunneling	Yes	No	Yes
Dynamic Topics	No	No	Yes
XPath queries/subscriptions	No	No	Yes
Performance Monitoring	No	No	Yes
End-to-End security	Yes	No	Yes. (Dec 2003)
Support for Workflow	Yes.	No.	No.
Support for distributed caching	No.	Yes (Squirrel)	No.
Maturity of Software	Extremely mature	Fair	Fair

The NaradaBrokering software has been used extensively in several projects described below.

4. GlobalMMCS

The GlobalMMCS technology is a new collaboration environment built around systematic use of Web (Grid) services and NaradaBrokering messaging. In particular it converts SIP, H323 and Access Grid protocols to a common Web service protocol XGSP. GlobalMMCS uses NaradaBrokering in all its communication whether it is in form of

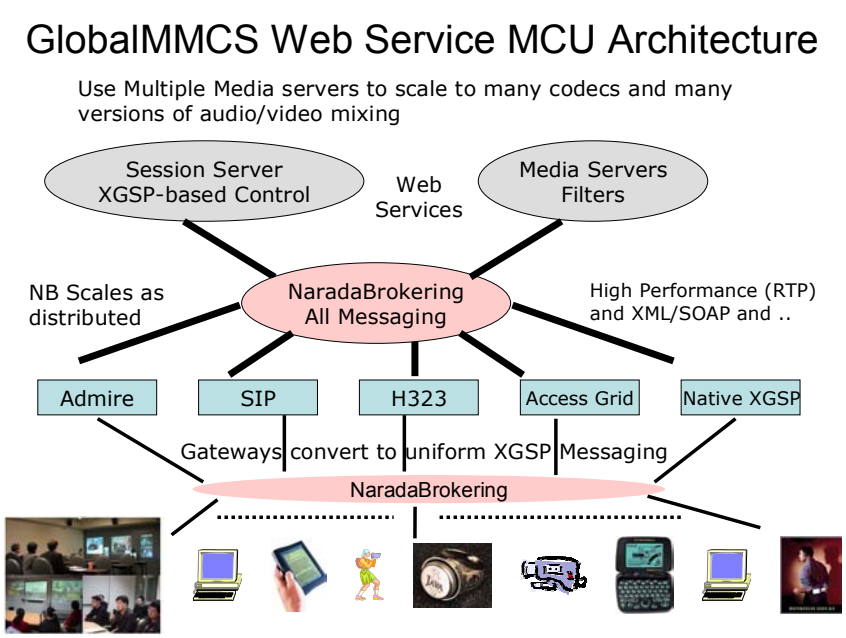

control packets or RTP based multimedia streams. It features Web/Grid services for session control, media type conversion, audio mixing and video mosaicing.

Functionality	Polycom (H.323)	AccessGrid	GlobalMMCS
Protocol Support	Only H.323 clients	Only MBONE clients	H.323, SIP, MBONE
Server Architecture	Expensive H.323 Hardware/Software MCU	No conf. servers, Needs multicast support in the network	Software Servers solution. Based on Web Services
Network and Communication Environment	Internet / ISDN Firewall transversal under the support of VPN	High performance with multicast support No firewall tunneling	Publish/Subscribe with Firewall transversal (VPN optional)
A/V Interaction Capability	Client shows a few video streams	A client shows up to 50 streams	A client shows up to 50 streams
Conferencing scalability	Usually less than 20~30 sites	Limited to the area of multicast networks	>1000 sites distributed over Internet
Integ. of streaming & conferencing	Only supports conferencing	Only supports conferencing	Any A/V source can be Streamed
Archive & Replay	No built in H.323 service	Voyager: allows replay of multiple streams	Multiple A/V streaming archive & replay
Other Collaboration Tools	Limited although T120 could support other apps	Limited to PowerPoint and chat	XGSP allows full integration of all tools

5. Universal Accessibility Service with NaradaBrokering

Universal access refers to the capability that all users are able to access information systems (grids) independent of their access device and their physical capabilities. The Carousel project addresses this for mobile devices such as smartphones and PDAs. We provide the collaborative service which is accessible from the heterogeneous types of mobile devices through NaradaBrokering.

5.1 Participating Collaborative Session with Mobile device

Carousel project allows users to participate a collaborative session with the mobile device wherever they are. The mobile users in the collaborative session share the application such as Microsoft PowerPoint or Web site browsed by Internet Explorer running on PCs. During the session, the mobile users also interact with other participants by exchanging text messages.

5.2 Support 3G Network Communication with NaradaBrokering
We are developing optimized protocol, HHMP (Hand-Held Message Protocol) for emerging 3G (Third-Generation) Network communication. This service provides seamless access from the mobile devices to NaradaBrokering.

5.3 MVC based dynamic sharing
Moreover, as an advanced sharing application technology, we support shared export which provides collaboration with more abstract data. Instead of sharing the framebuffer images as in shared display, the shared export exchanges abstract data such as vector graphics. We integrated W3C's 2D vector graphics format, SVG (Scalable Vector Graphics) and also provide the interface for mobile devices of collaborative SVG. Our collaborative SVG is designed based on the MVC (Model-View-Control) model therefore the presentation view of each user can be controlled individually. With the collaborative SVG service, mobile device also shares the display under its individual control in high-quality graphics.

6. NaradaBrokering-enhanced GridFTP

We have integrated GridFTP with the NaradaBrokering system. The advantage of this approach is that the source code of the GridFTP client need not be changed and existing client code can be used with NB system without any modification. Due to its integration with the NaradaBrokering system, files can be recovered after failures or even prolonged disconnects. Files can be partially transmitted and if there is a failure the transmission begins at a point closer to point where the failure occurred.

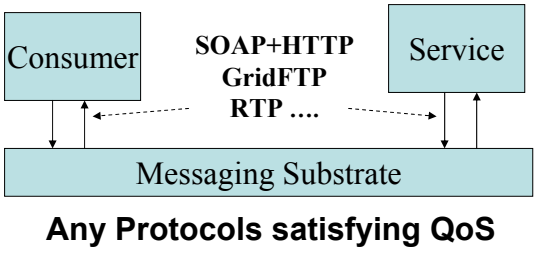

7. Anabas Web Conferencing

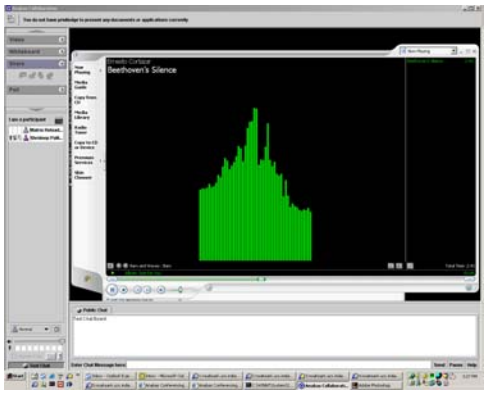

NaradaBrokering has also been deployed in real-time settings by providing back end support for the Anabas conferencing software running in JMS emulation mode. Several online seminars have been conducted using this Anabas-NaradaBrokering combination and the number of users collaborating concurrently has been in the excess of 30 users several times. Anabas uses NaradaBrokering to support shared display, text chats and all session control functions. Note shared display involves sharing large – sometime multi-megabit files specifying complete frame-buffer. The messaging system must handle this, the very dynamic video frames and the short text buffers of control and text chat.

8. Applications as Web Services

One can use the Web service model to build "ordinary" applications in the Model-View-Control paradigm with NaradaBrokering playing the role of the "control". Conventionally a user-interaction generates an interrupt to which one posts listeners. This is replaced by a "view" module sending messages to a "model"

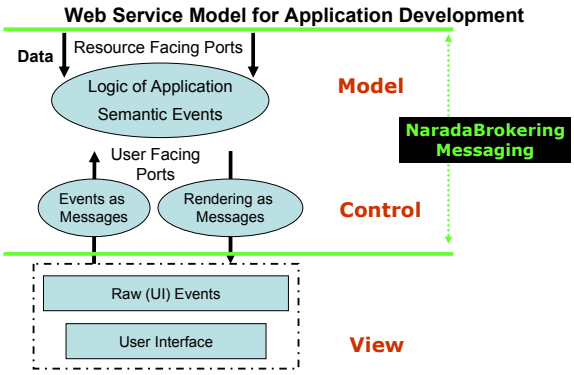

Interrupts in traditional monolithic applications become "real messages" not directly method calls
Natural for collaboration and universal access

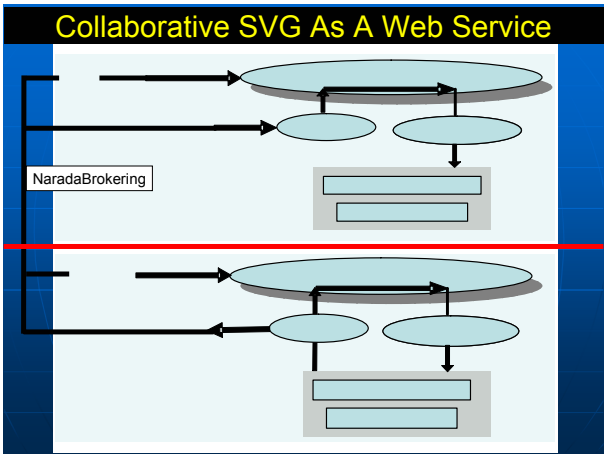

Web service where a subscriber plays role of listener. This way of building application has significant advantages over traditional bundled applications. We can retarget applications to multiple clients as with the PDA and universal access described in sec. 5. We can support multiple O/S more easily as only the view need be ported. Further it allows us to "automatically" build collaborative applications by sharing

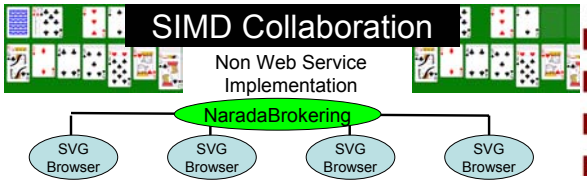

Identical Programs receiving identical events.Token determines if browser is moving, waiting for opponent or an observer

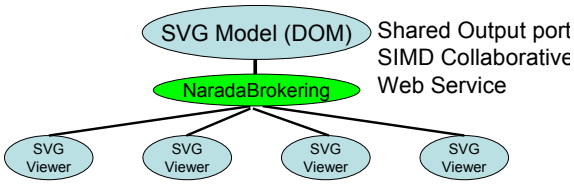

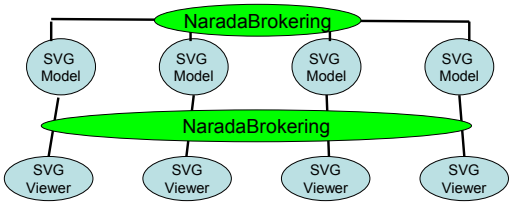

either the input or output ports of the web service. We have applied this to the Java SVG (Scalable Vector Graphics) browser and to PowerPoint; the former is a complete Web service implementation using user-interaction events; the latter uses higher-level semantic events like "slide change". We believe a complete Web Service architecture desktop can be built this way and that it will revolutionize client computing.

9. Federation in Peer-to-Peer Grids

NaradaBrokering can be used for federation between different distributed systems by acting as a dsitributed gateway intercepting messages between different syles of Grids and peer-to-peer networks. We have demonstrated this with the JXTA P2P network where we have modified their rendezvous peers to act as such gateways. This allows NaradaBrokering to federate many JXTA peer groups together.

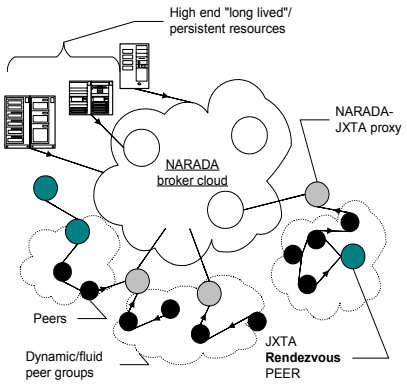

10. Java for High Performance Computing

The *HPJava* project was started several years ago by researchers now working at CGL. It was conceived as part of the community-wide *Java Grande* process. Both these activities have been promoting large scale computing—notably parallel and scientific computing—using software-engineering approaches based on the Java platform. An early product of the HPJava project was the widely-used *mpiJava* MPI binding for Java. This package continues to be downloaded and supported, and is being used in classroom settings.

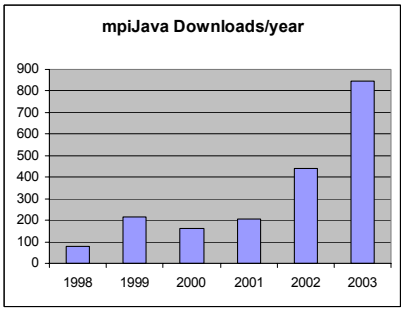

Another major milestone was achieved earlier this year when we made the first release of the *HPJava Development Kit*. This includes a translator (i.e. compiler) for a syntactic extension of Java with support for Fortran-like multidimensional arrays and HPF-like distributed arrays. The lack of these has been an obstacle to uptake of Java for scientific computing. The release package also includes standard libraries for operating on the new data structures.

10.1 mpiJava

The mpiJava package provides a bridge between Java and various platform-specific implementations of the *Message Passing Interface* for high-performance computing. Supported native platforms include MPICH, LAM, SunHPC and the IBM SP implementation of MPI. So mpiJava allows development of portable Java programs that can be executed efficiently on important distributed-memory parallel computers. This software continues to be popular (see chart—about 2000 downloads so far) and we actively maintain and update it.

10.2 HPJava Development Kit

In April this year we released the first version of the HPJava Development Kit, as open source software. It supports development of scientific software in an extended version of the Java language, called HPJava. HPJava builds on mpiJava and provides a much higher level collective communication library called Adlib. HPJava programs are compiled to standard Java class files that can execute on a distributed collection of standard Java Virtual Machines, interacting through a native MPI or TCP/IP sockets. It can also run in parallel on SMPs using Java thread parallelism. Compilers and high-level libaries are pure Java. A new portability layer called *mpjdev* encapsulates any unavoidable platform dependence.