Suggested Changes for Book (reducing size):

1) Remove sections 5.2,6,7.2 and possibly section 4

2) Remove current figures 2,3,5,8,9,10,11,16

3) Point out constraints on things like scheduling, QoS, security by elaborating requirements of integrating with commodity approaches

4) Add integration with technology chapters -- especially Kennedy, Gannon, Grimshaw, Dongarra, Johnson

5) Add integration with applications -- particularly clear with Moore (as stated in terms of requiring commodity services) and Messina as this is HLA application we discuss anyway

HPcc as High Performance Commodity Computing

Geoffrey C. Fox, Wojtek Furmanski

gcf@npac.syr.edu, furm@npac.syr.edu

World Wide Web: http://www.npac.syr.edu

http://www.npac.syr.edu/users/gcf/HPcc/HPcc.html

Abstract:

We review the growing power and capability of commodity computing and communication technologies largely driven by commercial distributed information systems. These systems are built from CORBA, Microsoft's COM, Javabeans, and less sophisticated web and networked approaches. One can abstract these to a three-tier model with largely independent clients connected to a distributed network of servers. The latter host various services including object and relational databases and of course parallel and sequential computing. High performance can be obtained by combining concurrency at the middle server tier with optimized parallel back end services. The resultant system combines the needed performance for large-scale HPCC applications with the rich functionality of commodity systems. Further the architecture with distinct interface, server and specialized service implementation layers, naturally allows advances in each area to be easily incorporated. We show that this approach can be applied to both metacomputing and to provide improved parallel programming environments. We describe exploitation issues within a CORBA context and illustrate how performance can be obtained within a commodity architecture. Examples are given from collaborative systems, support of multidisciplinary interactions, proposed visual HPCC ComponentWare, distributed simulation and the use of Java in high performance computing.

1: Introduction

We believe that industry and the loosely organized worldwide collection of (freeware) programmers is developing a remarkable new software environment of unprecedented quality and functionality. We call this DcciS - Distributed commodity computing and information System. We believe that this can benefit HPCC in several ways and allow the development of both more powerful parallel programming environments and new distributed metacomputing systems. In the second section, we define what we mean by commodity technologies and explain the different ways that they can be used in HPCC. In the third and critical section, we define an emerging architecture of DcciS in terms of a conventional 3 tier commercial computing model. The next section describes the expected steps in the CORBA model for establishing HPcc as a community framework and CORBA facility. In this and related papers (ref. [31]), we discuss several examples and the critical research issue: can high performance systems - called HPcc or High Performance Commodity Computing -be built on top of DcciS. Examples include integration of collaboration into HPcc; the natural synergy of distribution simulation and the HLA standard with our architecture (ref. [33]); and the step from object to visual component based programming in parallel and distributed computing. Finally we discuss the use of DcciS to build parallel programming environments and HPJava -- the use of Java on parallel or sequential machines for high performance computing (refs. [28,29]). There is a reading list and glossary of acronyms and buzzwords at the end of the article.

2: Commodity Technologies and their use in HPCC

The last three years have seen an unprecedented level of innovation and progress in commodity technologies driven largely by the new capabilities and business opportunities of the evolving worldwide network. The web is not just a document access system supported by the somewhat limited HTTP protocol. Rather it is the distributed object technology which can build general multi-tiered enterprise intranet and internet applications. CORBA is turning from a sleepy heavyweight standards initiative to a major competitive development activity that battles with COM and Javabeans to be the core distributed object technology.

There are many driving forces and many aspects to DcciS but we suggest that the three critical technology areas are the web, distributed objects and databases. These are being linked and we see them subsumed in the next generation of "object-web" technologies, which is illustrated by the recent Netscape and Microsoft version 4 browsers. Databases are older technologies but their linkage to the web and distributed objects, is transforming their use and making them more widely applicable.

In each commodity technology area, we have impressive and rapidly improving software artifacts. As examples, we have at the lower level the collection of standards and tools such as HTML, HTTP, MIME, IIOP, CGI, Java, JavaScript, Javabeans, CORBA, COM, ActiveX, VRML, new powerful object brokers (ORB's), dynamic Java servers and clients including applets and servlets. At a higher level collaboration, security, commerce, multimedia and other applications/services are rapidly developing using standard interfaces or frameworks and facilities. This emphasizes that equally and perhaps more importantly than raw technologies, we have a set of open interfaces enabling distributed modular software development. These interfaces are at both low and high levels and the latter generate a very powerful software environment in which large preexisting components can be quickly integrated into new applications. We believe that there are significant incentives to build HPCC environments in a way that naturally inherits all the commodity capabilities so that HPCC applications can also benefit from the impressive productivity of commodity systems. NPAC's HPcc activity is designed to demonstrate that this is possible and useful so that one can achieve simultaneously both high performance and the functionality of commodity systems.

Note that commodity technologies can be used in several ways. This article concentrates on exploiting the natural architecture of commodity systems but more simply, one could just use a few of them as "point solutions". This we can term a "tactical implication" of the set of the emerging commodity technologies and illustrate below with some examples:

- Perhaps VRML or Java3D are important for scientific visualization
- Web (including Java applets) front-ends provide convenient customizable interoperable user interfaces to HPCC facilities
- Perhaps the public key security and digital signature infrastructure being developed for electronic commerce, could enable more powerful approaches to secure HPCC systems
- Perhaps Java will become a common scientific programming language and so effort now devoted to Fortran and C++ tools needs to be extended or shifted to Java
- The universal adoption of JDBC (Java Database Connectivity) and the growing convenience of web-linked databases could imply a growing importance of systems that link large scale commercial databases with HPCC computing resources
- Javabeans, RMI, COM, CORBA, IIOP form the basis of the emerging "object web" which analogously to the previous bullet could encourage a growing use of modern object technology
- Emerging collaboration and other distributed information systems could allow new distributed work paradigms which could change the traditional teaming models in favor of those for instance implied by the new NSF Partnerships in Advanced Computation

However probably more important is the strategic implication of DcciS which implies certain critical characteristics of the overall architecture for a high performance parallel or distributed computing system. First we note that we have seen over the last 30 years many other major broad-based hardware and software developments -- such as IBM business systems, UNIX, Macintosh/PC desktops, video games -- but these have not had profound impact on HPCC software. However we suggest the DcciS is different for it gives us a world-wide/enterprise-wide distributing computing environment. Previous software revolutions could help individual components of a HPCC software system but DcciS can in principle be the backbone of a complete HPCC software system -- whether it be for some global distributed application, an enterprise cluster or a tightly coupled large scale parallel computer. In a nutshell, we suggest that "all we need to do" is to add "high performance" (as measured by bandwidth and latency) to the emerging commercial concurrent DcciS systems. This "all we need to do" may be very hard but by using DcciS as a basis we inherit a multi-billion dollar investment and what in many respects is the most powerful productive software environment ever built. Thus we should look carefully into the design of any HPCC system to see how it can leverage this commercial environment.

3: Three Tier High Performance Commodity Computing

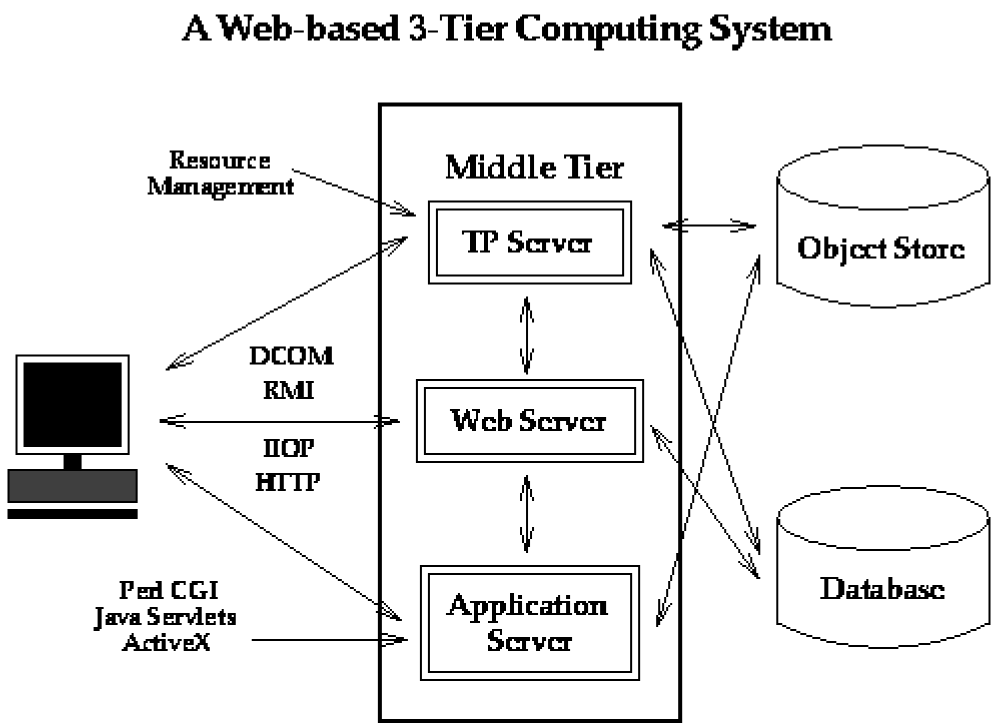

Fig. 1: Industry 3-tier view of enterprise Computing

We start with a common modern industry view of commodity computing with the three tiers shown in fig 1. Here we have customizable client and middle tier systems accessing "traditional" back end services such as relational and object databases. A set of standard interfaces allows a rich set of custom applications to be built with appropriate client and middleware software. As indicated on figure, both these two layers can use web technology such as Java and Javabeans, distributed objects with CORBA and standard interfaces such as JDBC (Java Database Connectivity). There are of course no rigid solutions and one can get "traditional" client server solutions by collapsing two of the layers together. For instance with database access, one gets a two tier solution by either incorporating custom code into the "thick" client or in analogy to Oracle's PL/SQL, compile the customized database access code for better performance and incorporate the compiled code with the back end server. The latter like the general 3-tier solution, supports "thin" clients such as the currently popular network computer.

The commercial architecture is evolving rapidly and is exploring several approaches which co-exist in today’s (and any realistic future) distributed information system. The most powerful solutions involve distributed objects. There are three important commercial object systems - CORBA, COM and Javabeans. These have similar approaches and it is not clear if the future holds a single such approach or a set of interoperable standards. CORBA is a distributed object standard managed by the OMG (Object Management Group) comprised of 700 companies. COM is Microsoft’s distributed object technology initially aimed at Window machines. Javabeans (augmented with RMI and other Java 1.1 features) is the "pure Java" solution - cross platform but unlike CORBA, not cross-language! Legion is an example of a major HPCC focused distributed object approach; currently it is not built on top of one of the three major commercial standards. The HLA/RTI standard for distributed simulations in the forces modeling community is another important domain specific distributed object system. It appears to be moving to integration with CORBA standards. Although a distributed object approach is attractive, most network services today are provided in a more ad-hoc fashion. In particular today’s web uses a "distributed service" architecture with HTTP middle tier servers invoking via the CGI mechanism, C and Perl programs linking to databases, simulations or other custom services. There is a trend toward the use of Java servers with the servlet mechanism for the services. This is certainly object based but does not necessarily implement the standards implied by CORBA, COM or Javabeans. However, this illustrates an important evolution as the web absorbs object technology with the evolution:

HTTP --> Java Sockets --> IIOP or RMI

(Low Level network standard) (High level network standard)

Perl CGI Script --> Java Program --> Javabean distributed object.

As an example consider the evolution of networked databases. Originally these were client-server with a proprietary network access protocol. Web linked databases shown in figs. 2 and 3, produced a three tier distributed service model with an HTTP server using a CGI program (running Perl for instance) to access the database at the backend. Today we can build databases as distributed objects with a middle tier Javabean using JDBC to access the backend database. Thus a conventional database is naturally evolving to the concept of managed persistent objects.

Multi Tier Client-Server-Service Examples I

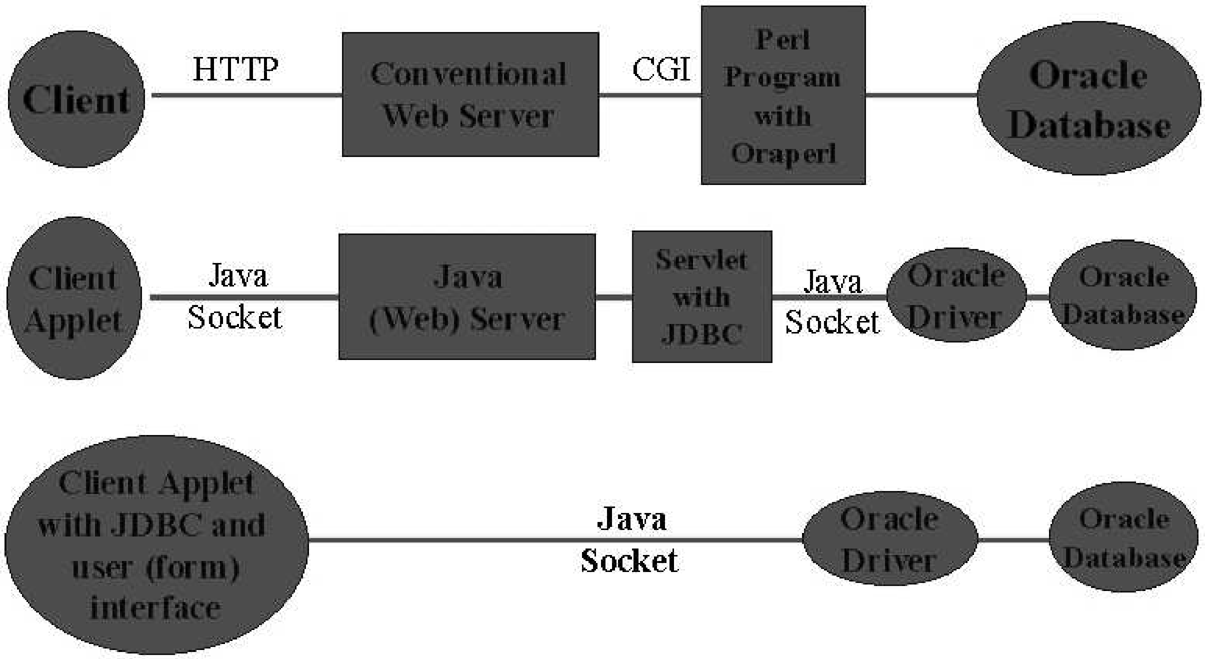

Fig. 2: Some example implementations of an Oracle Database in the 3-tier Approach

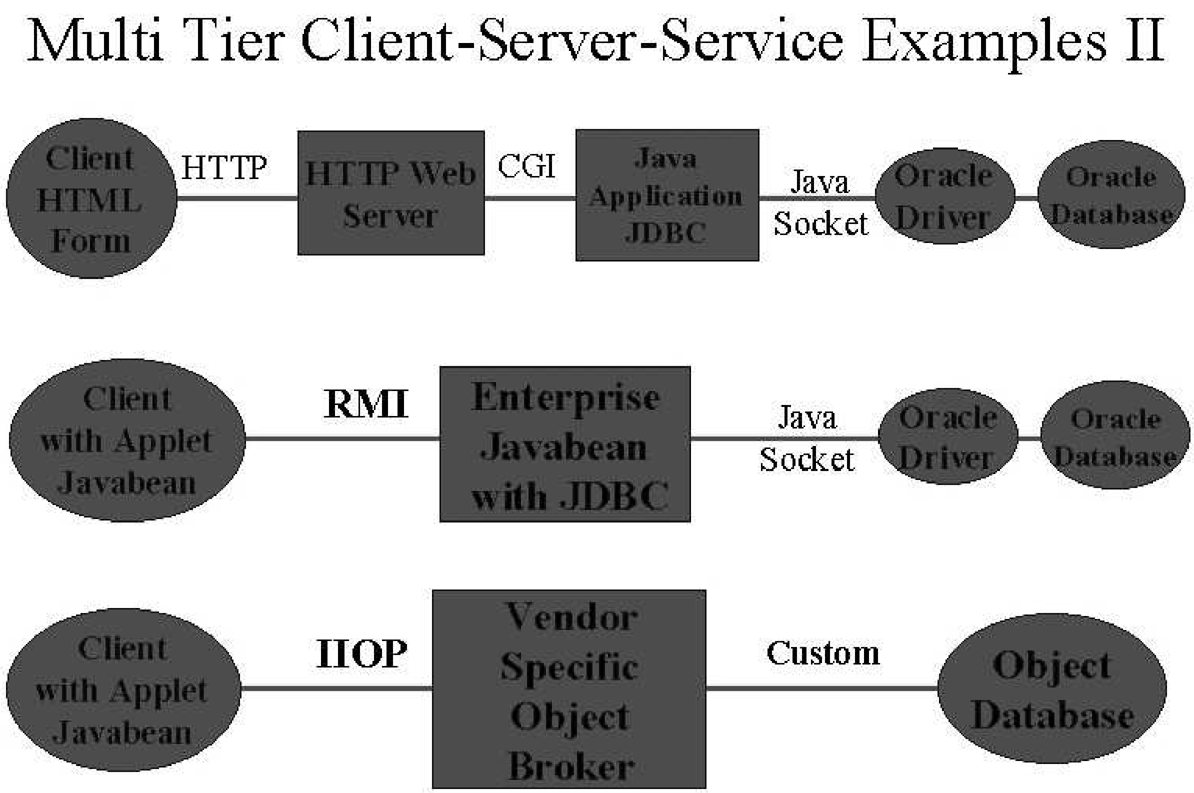

Fig. 3: More examples of the 3-tier Approach

Today as shown in fig. 4, we see a mixture of distributed service and distributed object architectures. CORBA, COM, Javabean, HTTP Server + CGI, Java Server and servlets, databases with specialized network accesses, and other services co-exist in the heterogeneous environment with common themes but disparate implementations. We believe that there will be significant convergence as a more uniform architecture is in everyone’s best interest.

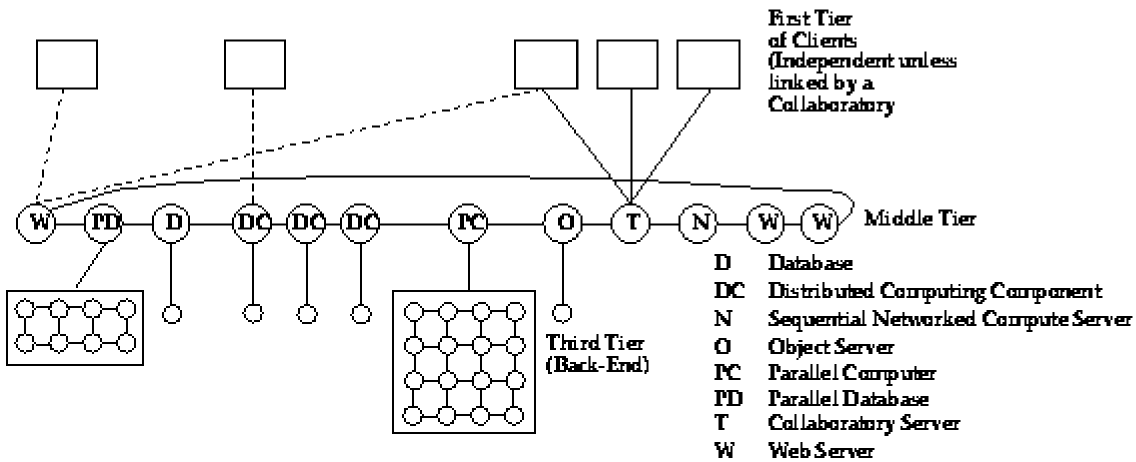

Fig. 4: Today's Heterogeneous Interoperating Hybrid Server Architecture. HPcc involves adding to this system, high performance in the third tier.

We also believe that the resultant architecture will be integrated with the web so that the latter will exhibit distributed object architecture shown in fig. 5.

More generally the emergence of IIOP (Internet Inter-ORB Protocol), CORBA2, and the realization that CORBA is naturally synergistic with Java is starting a new wave of "Object Web" developments that could have profound importance. Java is not only a good language to build brokers but also Java objects are the natural inhabitants of object databases. The resultant architecture in fig. 5 shows a small object broker (a so-called ORBlet) in each browser as in Netscape's current plans. Most of our remarks are valid for all these approaches to a distributed set of services. Our ideas are however easiest to understand if one assumes an underlying architecture which is a CORBA or Javabean distributed object model integrated with the web.

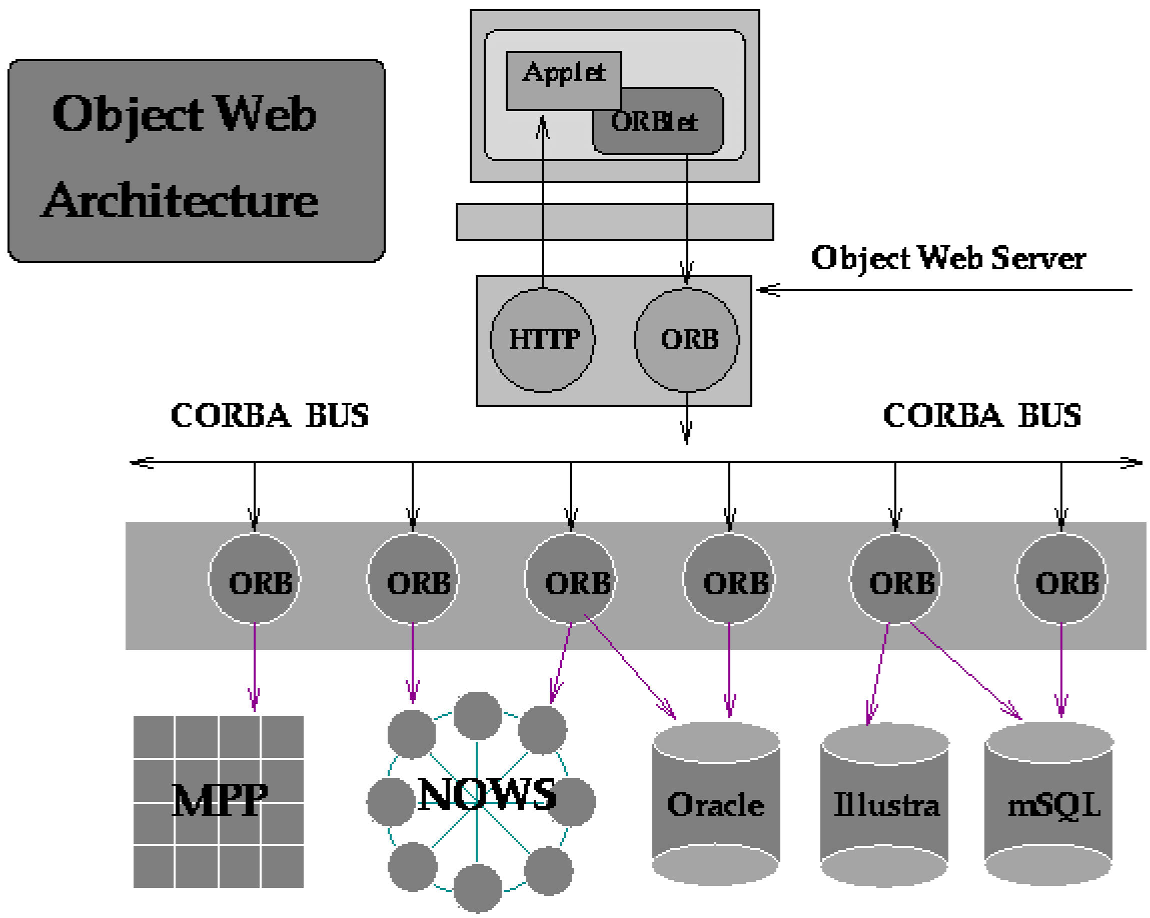

Fig. 5: Integration of Object Technologies (CORBA) and the Web

We wish to use this service/object evolving 3-tier commodity architecture as the basis of our HPcc environment. We need to naturally incorporate (essentially) all services of the commodity web and to use its protocols and standards wherever possible. We insist on adopting the architecture of commodity distribution systems as complex HPCC problems require the rich range of services offered by the broader community systems. Perhaps we could "port" commodity services to a custom HPCC system but this would require continued upkeep with each new upgrade of the commodity service. By adopting the architecture of the commodity systems, we make it easier to track their rapid evolution and expect it will give high functionality HPCC systems, which will naturally track the evolving Web/distributed object worlds. This requires us to enhance certain services to get higher performance and to incorporate new capabilities such as high-end visualization (e.g. CAVE's) or massively parallel systems where needed. This is the essential research challenge for HPcc for we must not only enhance performance where needed but do it in a way that is preserved as we evolve the basic commodity systems. We certainly have not demonstrated clearly that this is possible but we have a simple strategy that we will elaborate in ref. [31] and sec.5. Thus we exploit the three-tier structure and keep HPCC enhancements in the third tier, which is inevitability, the home of specialized services in the object-web architecture. This strategy isolates HPCC issues from the control or interface issues in the middle layer. If successful we will build an HPcc environment that offers the evolving functionality of commodity systems without significant re-engineering as advances in hardware and software lead to new and better commodity products.

Returning to fig. 4, we see that it elaborates fig. 1 in two natural ways. Firstly the middle tier is promoted to a distributed network of servers; in the "purest" model these are CORBA/ COM/ Javabean object-web servers as in fig. 5, but obviously any protocol compatible server is possible. This middle tier layer includes not only networked servers with many different capabilities (increasing functionality) but also multiple servers to increase performance on an given service. The use of high functionality but modest performance communication protocols and interfaces at the middle tier limits the performance levels that can be reached in this fashion. However this first step gives a modest performance scaling, parallel (implemented if necessary, in terms of multiple servers) HPcc system which includes all commodity services such as databases, object services, transaction processing and collaboratories. The next step is only applied to those services with insufficient performance. Naively we "just" replace an existing back end (third tier) implementation of a commodity service by its natural HPCC high performance version. Sequential or socket based messaging distributed simulations are replaced by MPI (or equivalent) implementations, as in sec. 7, on low latency high bandwidth dedicated parallel machines. These could be specialized architectures or "just" clusters of workstations. Note that with the right high performance software and network connectivity, workstations can be used at tier three just as the popular "LAN" consolidation" use of parallel machines like the IBM SP-2, corresponds to using parallel computers in the middle tier. Further a "middle tier" compute or database server could of course deliver its services using the same or different machine from the server. These caveats illustrate that as with many concepts, there will be times when the relatively clean architecture of fig 4 will become confused. In particular the physical realization does not necessarily reflect the logical architecture shown in fig 4. For instance figs. 2 and 3 give examples of database and object services implemented in different ways in the commodity 3-tier approach.

4: HPcc as a CORBA Facility

CORBA is defined as a suite of software layers with the architecture illustrated in fig. 6.

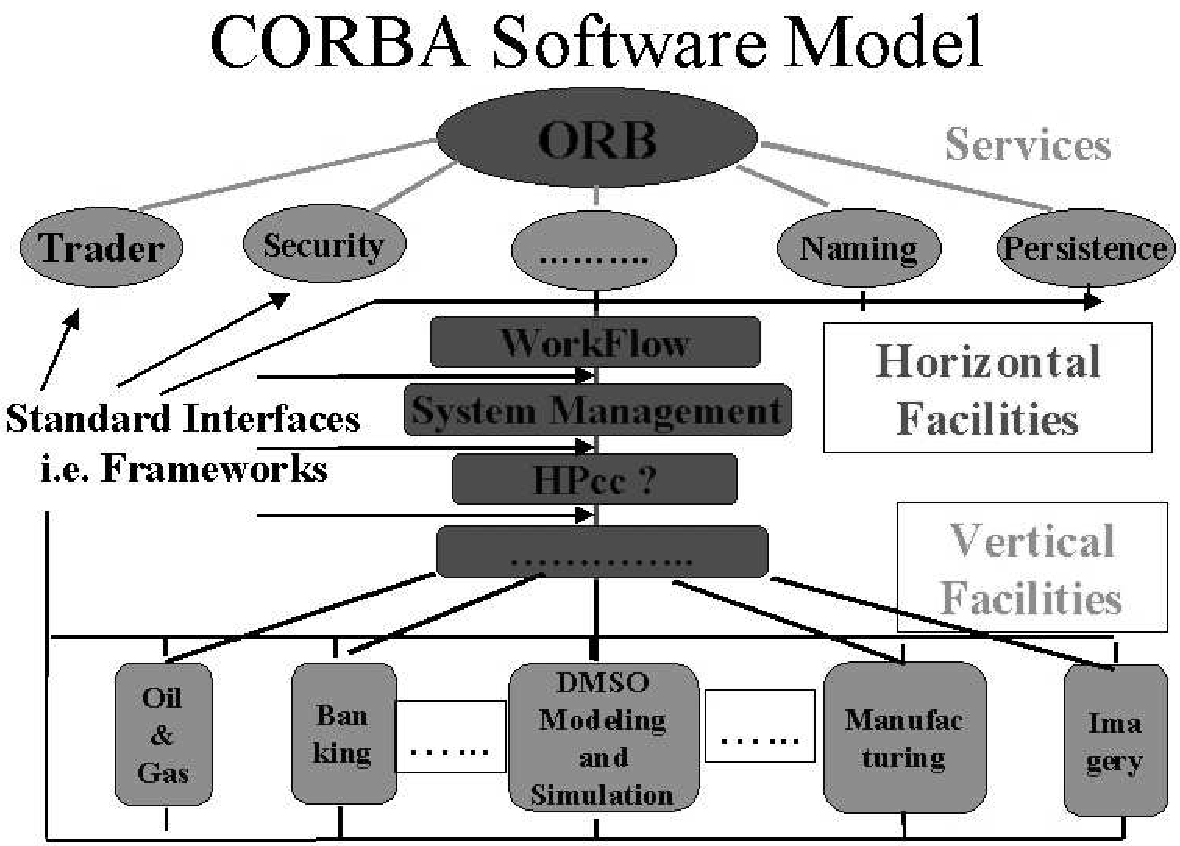

Fig. 6: Software Layers in CORBA

We see (currently 15) basic services such as naming and persistence layered below a set of general capabilities or horizontal facilities in the ComponentWare jargon. We suggest that HPcc is naturally placed here as it is designed to provide high performance to essentially any application area. The latter are seen as vertical or specialized facilities in fig. 6 that provide domain-specific distributed object support. Here we see mainstream commercial applications such as manufacturing, banking and mapping. The vertical and horizontal facilities are associated with frameworks that are the universal interfaces with which they are linked together and to user programs (or objects).

Note that CORBA currently supports only relatively simple computing models including "embarrassingly parallel" as in transaction processing or dataflow as in the CORBA workflow facility. The modeling and simulation community appears likely to evolve their HLA standard to a new vertical CORBA facility. HPcc therefore fills a gap and is defined as the HPCC (here we are using capital C's) CORBA horizontal facility. In the following paragraph, we point out that this allows us to define a commercialization strategy for high performance computing technologies.

Academia and Industry should now experiment with the design and implementations of HPcc as a general framework for providing high performance CORBA services. Then one or more Industry led groups proposes HPcc specifications as a new horizontal facility. A process similar to the MPI or HPF forum activities leads to consensus and the definition of an HPcc facility standard. This then allows industry (and academia) to compete in the implementation of HPcc within an agreed interoperable standard. HPcc can either include or be accompanied by another CORBA facility -- namely that involved in user interfaces to (scientific) computers. This would include interfaces necessary for performance tools and resource managers as well as file systems, compilation, debugging and visualization. Such a seamless interface is the subject of a Birds of the Feather Session at SC97 and a recent workshop in England (Ref. [32]).

The remarks we make above for CORBA have analogies in the Java and COM object models. In particular CORBA facilities are logically equivalent to the different frameworks being developed for Java. HPcc would become a Java framework for high performance computing and used in 100% Java solutions -- something that is quite realistic as we show in section 8.2 that Java is a potentially excellent language for Scientific computing. Of course not all code will be written in Java and it is unlikely to be wise to convert existing Fortran and C++ to Java. However putting Java (or more generally CORBA) wrappers around existing code seems to us a good way of preserving old codes. This can both document their capability (through the CORBA trader and Javabean Information services) and allow definition of methods that allow such codes to be naturally incorporated into larger systems. In this way a Java framework for HPcc can be used in general computing solutions. In particular the CORBA seamless interfaces discussed above are likely to be built in a 100% Java environment anyway as the natural implementation is a client side (see ref. [36]) Java Applet linking to a Java Web server acting as the proxy for the high performance computer. Finally we note that the discussion in section 8.2 of HPJava -- the collection of libraries, compilers, coding practices and environments to support high performance scientific computing naturally form a Java framework for Scientific Computing. Alternatively high performance scientific computing is another CORBA facility that is more specialized than HPcc and hence becomes perhaps a vertical facility.

5: A Possible Approach for High Performance Commodity Systems

5.1: Multidisciplinary Applications

We can illustrate the commodity technology strategy with a simple multidisciplinary application involving the linkage of two modules A and B -- say CFD and structures applications respectively. Let us assume both are individually parallel but we need to link them. One could view the linkage sequentially as in fig. 7(a), but often one needs higher performance and one would "escape" totally into a layer which linked decomposed components of A and B with high performance MPI (or PVMPI). Here we view MPI as the "machine language" of the higher-level commodity communication model given by approaches such as WebFlow from NPAC.

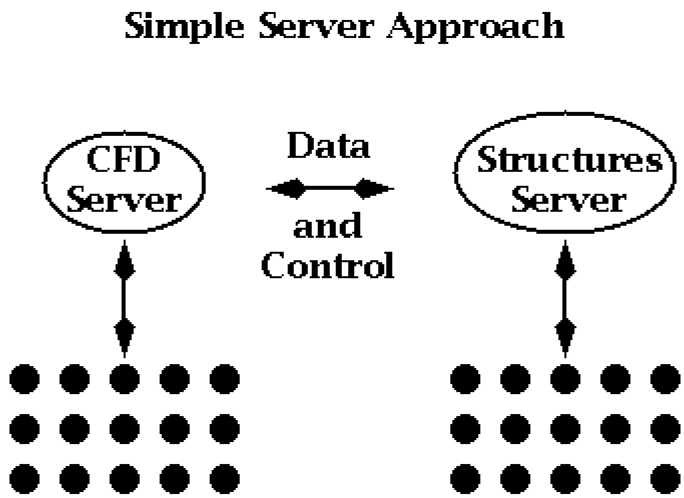

Fig. 7(a): Simple sequential server approach to Linking Two Modules

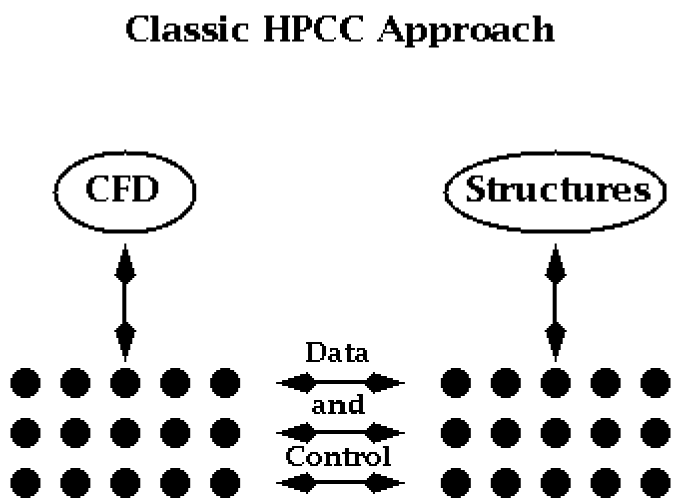

Fig. 7(b): Full HPCC approach to Linking Two Modules

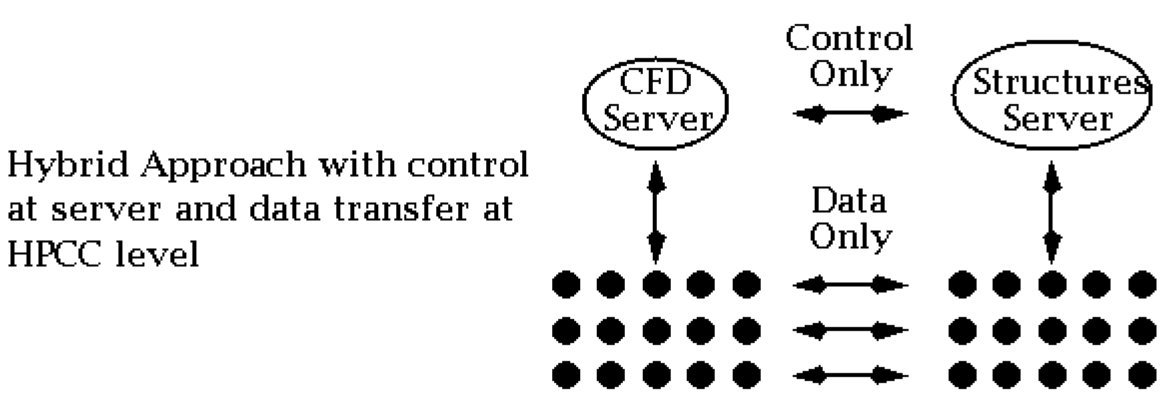

Fig. 7(c): Hybrid approach to Linking Two Modules

There is the "pure" HPCC approach of fig. 7(b), which replaces all commodity web communication with HPCC technology. However there is a middle ground between the implementations of figs. 7(a) and 7(b) where one keeps control (initialization etc.) at the server level and "only" invokes the high performance back end for the actual data transmission. This is shown in fig. 7(c) and appears to obtain the advantages of both commodity and HPCC approaches for we have the functionality of the Web and where necessary the performance of HPCC software. As we wish to preserve the commodity architecture as the baseline, this strategy implies that one can confine HPCC software development to providing high performance data transmission with all of the complex control and service provision capability inherited naturally from the Web.

5.2: Javabean Communication Model

We note that Javabeans (which are one natural basis of implementing program modules in the HPcc approach) provide a rich communication mechanism, which supports the separation of control (handshake) and implementation. As shown below in fig. 8, Javabeans use the JDK 1.1 AWT event model with listener objects and a registration/call-back mechanism.

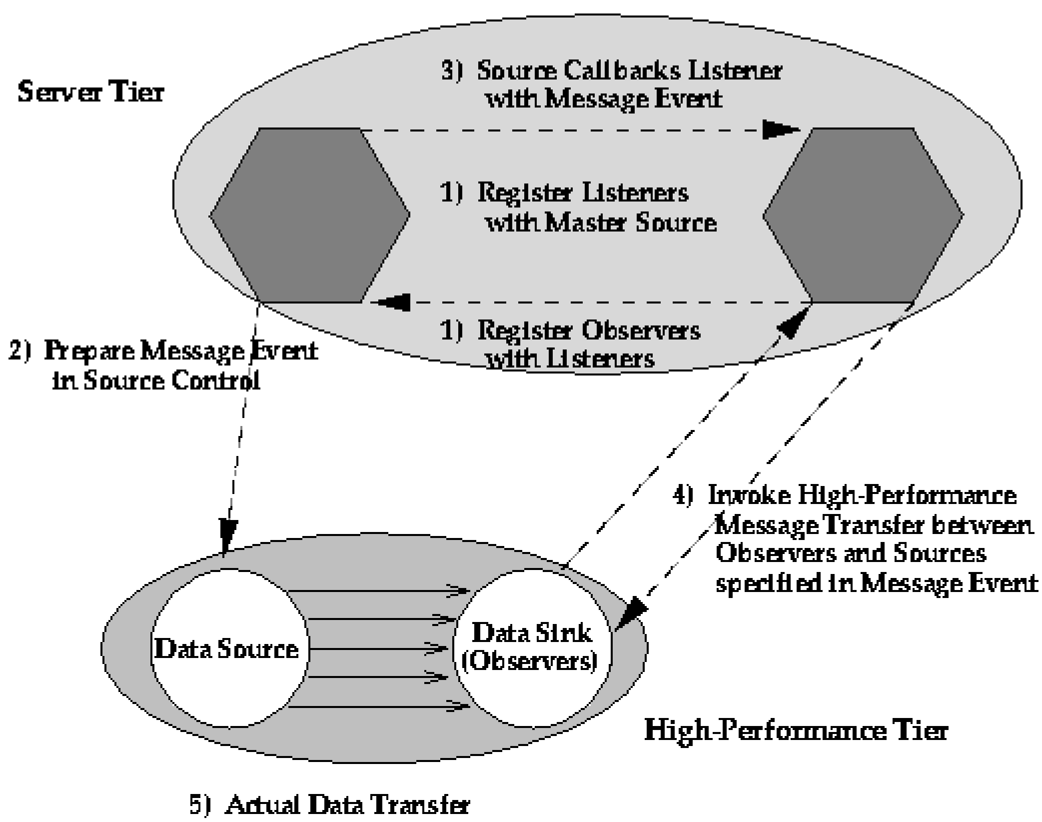

Fig. 8: JDK 1.1 Event Model used by (inter alia) Javabeans

Javabeans communicate indirectly with one or more "listener objects" acting as a bridge between the source and sink of data. In the model described above, this allows a neat implementation of separated control and explicit communication with listeners (a.k.a. sink control) and source control objects residing in middle tier. These control objects decide if high performance is necessary or possible and invoke the specialized HPCC layer. This approach can be used to advantage in "run-time compilation" and resource management with execution schedules and control logic in the middle tier and libraries such as MPI, PCRC and CHAOS implementing the determined data movement in the high performance (third) tier. Parallel I/O and "high-performance" CORBA can also use this architecture. In general, this listener model of communication provides a virtualization of communication that allows a separation of control and data transfer that is largely hidden from the user and the rest of the system. Note that current Internet security systems (such as SSL and SET) use high functionality public keys in the control level but the higher performance secret key cryptography in bulk data transfer. This is another illustration of the proposed hybrid multi-tier communication mechanism.

6: Commodity Services in HPcc

We have already stressed that a key feature of HPcc is its support of the natural inclusion into the environment of commodity services such as databases, web servers and object brokers. Here we give some further examples of commodity services that illustrate the power of the HPcc approach.

6.1:Distributed Collaboration Mechanisms

The current Java Server model for the middle tier naturally allows one to integrate collaboration into the computing model and our approach allow one to "re-use" collaboration systems built for the general Web market. Thus one can without any special HPCC development, address areas such as computational steering and collaborative design, which require people to be integrated with the computational infrastructure. In fig. 9, we define collaborative systems as integrating client side capabilities together. In steering, these are people with analysis and visualization software. In engineering design, one would also link design (such as CATIA or AutoCAD) and planning tools. In both cases, one would need the base collaboration tools such as white-boards, chat rooms and audio-video conferencing.

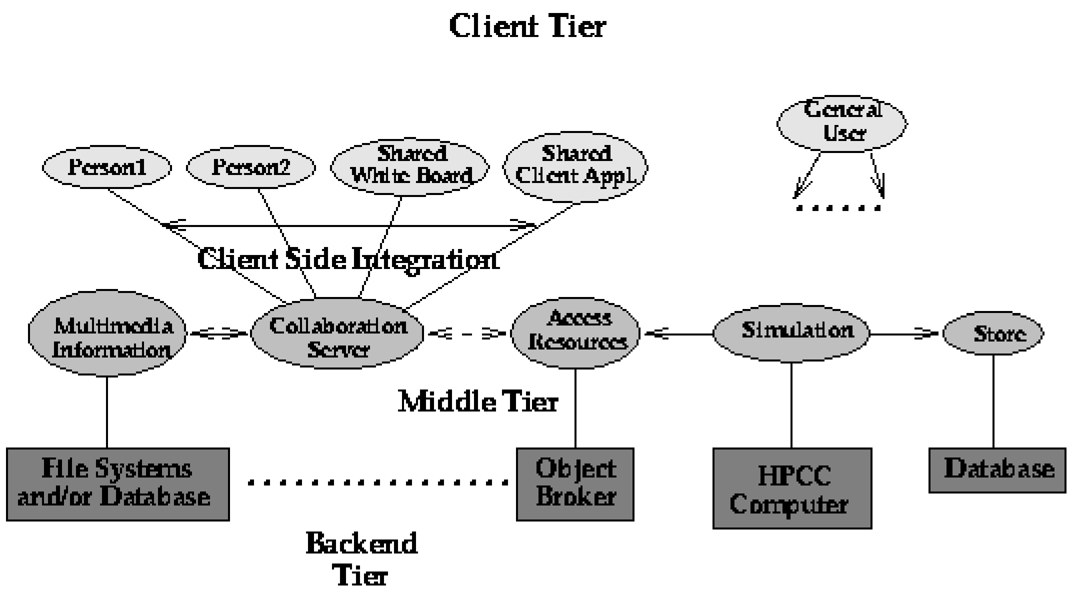

Fig. 9: Collaboration in today's Java Web Server implementation of the 3 tier computing model. Typical clients (on top right) are independent but Java collaboration systems link multiple clients through object (service) sharing.

If we are correct in viewing collaboration as sharing of services between clients, the 3 tier model naturally separates HPCC and collaboration and allows us to integrate into the HPCC environment, the very best commodity technology which is likely to come from larger fields such as business or (distance) education. Currently commodity collaboration systems are built on top of the Web and although emerging facilities such as Work Flow imply approaches to collaboration, are not yet defined from a general CORBA point of view. We assume that collaboration is sufficiently important that it will emerge as a CORBA capability to manage the sharing and replication of objects. Note CORBA is a server-server model and "clients" are viewed as servers (i.e. run Orb's) by outside systems. This makes the object-sharing view of collaboration natural whether application runs on "client" (e.g. shared Microsoft Word document) or on back-end tier as in case of a shared parallel computer simulation.

NPAC has developed two systems TANGO (ref. [35]) and WebFlow (ref. [34]), which illustrate the differences between collaborative and computational sharing. Both NPAC systems use Java servers for their middle tier level. TANGO, which can be considered here as equivalent to NCSA's Habanero, provides collaboration services as in fig. 9. Client side applications are replicated using an event distribution model and to put a new application into TANGO one must be able to define both its absolute state and changes therein. Using Java object serialization or similar mechanisms, this state is maintained identically in the linked applications. On the other hand, WebFlow can be thought of a Web version of AVS or Khoros and as in fig.7(a) integrates program modules together using a dataflow paradigm. Now the module developer defines data input and output interfaces and builds methods to handle data I/O. However there would typically be no need to replicate the state of module in a WebFlow application.

6.2: Object Web and Distributed Simulation

The integration of HPCC with distributed objects provides an opportunity to link the classic HPCC ideas with those of DoD's distributed simulation DIS or Forces Modeling FMS community. The latter do not make extensive use of the Web these days but they have a commitment to CORBA with their HLA (High Level Architecture) and RTI (Runtime Infrastructure) initiatives. The Naval postgraduate school at Monterey has built some interesting web-linked DIS prototypes. However distributed simulation is traditionally built with distributed event driven simulators managing C++ or equivalent objects. We suggest that the object web (and parallel and distributed ComponentWare described in sec. 6.3) is a natural convergence point for HPCC and DIS/FMS. This would provide a common framework for time stepped, real time and event driven simulations. Further it will allow one to more easily build systems that integrate these concepts as is needed in many major DoD projects -- as exemplified by the FMS and IMT DoD computational activities which are part of the HPCC Modernization program.

6.3: HPCC ComponentWare

HPCC does not have a good reputation for the quality and productivity of its programming environments. Indeed one of the difficulties with adoption of parallel systems, is the rapid improvement in performance of workstations and recently PC's with much better development environments. Parallel machines do have a clear performance advantage but this for many users, this is more than counterbalanced by the greater programming difficulties. We can give two reasons for the lower quality of HPCC software. Firstly parallelism is intrinsically hard to find and express. Secondly the PC and workstation markets are substantially larger than HPCC and so can support a greater investment in attractive software tools such as the well-known PC visual programming environments. The DcciS revolution offers an opportunity for HPCC to produce programming environments that are both more attractive than current systems and further could be much more competitive than previous HPCC programming environments with those being developed by the PC and workstation world. Here we can also give two reasons. Firstly the commodity community must face some difficult issues as they move to a distributed environment, which has challenges where in some cases the HPCC community has substantial expertise. Secondly as already described, we claim that HPCC can leverage the huge software investment of these larger markets.

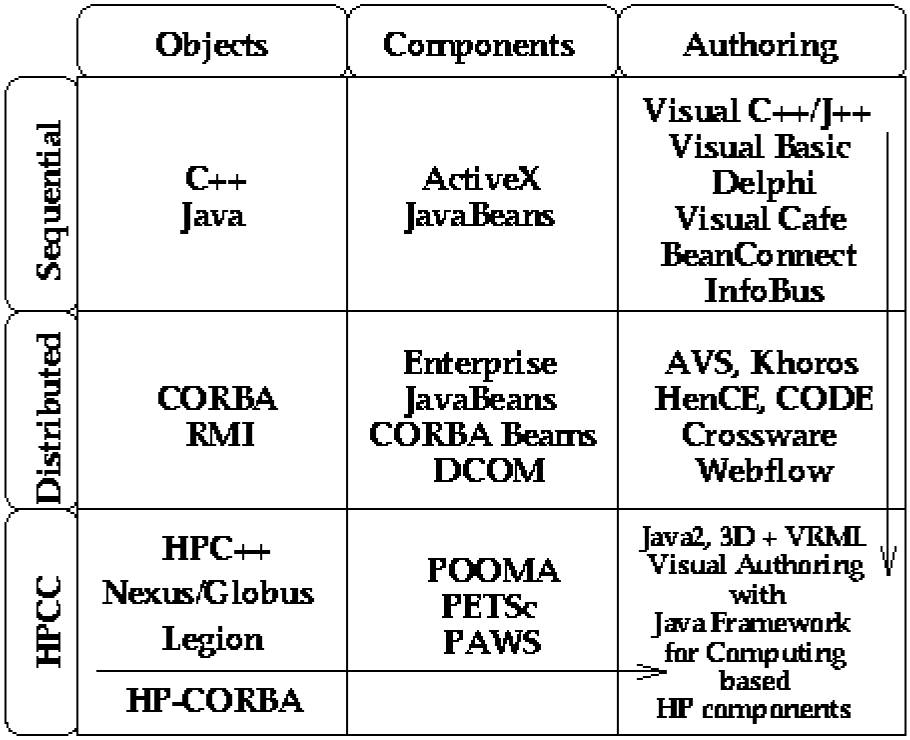

Fig. 10: System Complexity (vertical axis) versus User Interface (horizontal axis) tracking of some technologies

In fig. 10, we sketch the state of object technologies for three levels of system complexity -- sequential, distributed and parallel and three levels of user (programming) interface -- language, components and visual. Industry starts at the top left and moves down and across the first two rows. Much of the current commercial activity is in visual programming for sequential machines (top right box) and distributed components (middle box). Crossware (from Netscape) represents an initial talking point for distributed visual programming. Note that HPCC already has experience in parallel and distributed visual interfaces (CODE and HenCE as well as AVS and Khoros). We suggest that one can merge this experience with Industry's Object Web deployment and develop attractive visual HPCC programming environments. Currently NPAC's WebFlow system uses a Java graph editor to compose systems built out of modules. This could become a prototype HPCC ComponentWare system if it is extended with the modules becoming Javabeans and the integration with CORBA. Note the linkage of modules would incorporate the generalized communication model of fig. 7. Returning to fig. 1, we note that as industry moves to distributed systems, they are implicitly taking the sequential client-side PC environments and using them in the much richer server (middle-tier) environment which traditionally had more closed proprietary systems.

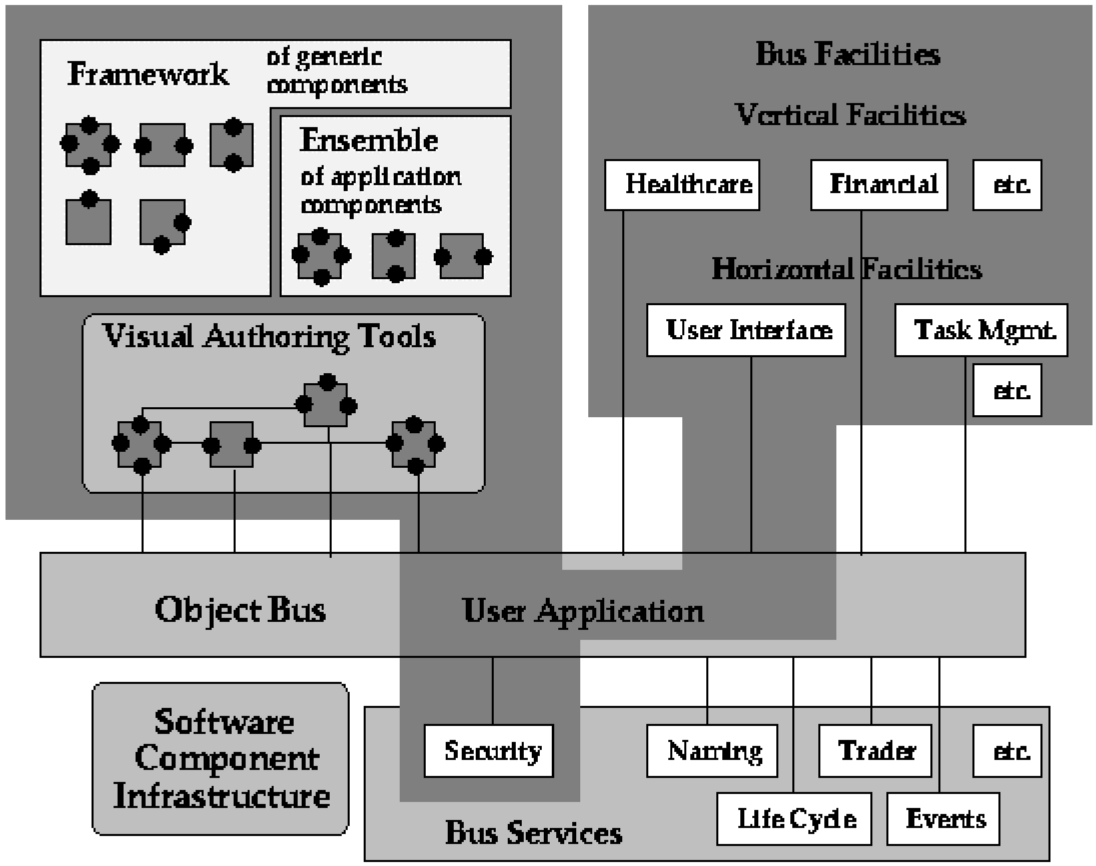

Fig. 11: Visual Authoring with Software Bus Components

We will then generate an environment such as fig. 11 including object broker services, and a set of horizontal (generic) and vertical (specialized application) frameworks. These represent a combination of library and interfaces and are illustrated by the discussion in section 4 of the CORBA facilities and Java frameworks. We do not have much experience with an environment such as fig. 11, but suggest that HPCC could benefit from its early deployment without the usual multi-year lag behind the larger industry efforts for PC's. Further the diagram implies a set of standardization activities (establish frameworks) and new models for services and libraries that could be explored in prototype activities.

7: Commodity Technologies for Parallel Computing Environments

7.1 High Performance Commodity Communication

Most of the discussion in this paper has been devoted to the use of commodity technologies for computational grids or the field that is sometimes termed HPDC (High Performance Distributed

Computing). However as we started to explore in the last subsection, we believe that one can also use commodity technologies to build parallel computing environments which combine both high functionality and high performance. As usual the functionality comes from inheriting commodity services and the challenge is to combine these with high performance.

First compare the two views of a parallel computer in figs. 12 and 13.

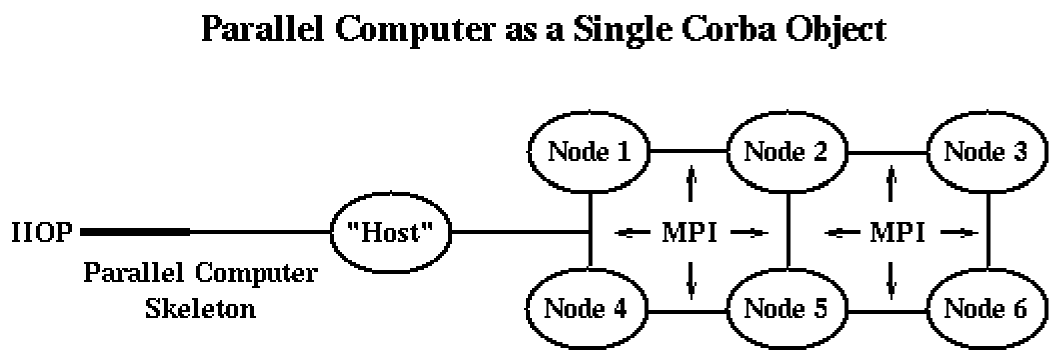

Fig. 12: A Parallel Computer viewed as a single CORBA object in a classic "Host-node computing model". Logically the "Host" is at tier-2 and the nodes are at tier 3. The Physical architecture could differ from the logical architecture.

In the above figure, we see a simple multi-tier view with commodity protocols (HTTP, RMI, COM or the IIOP pictured) used to access the parallel computer as a single entity. This entity (object) delivers high performance in an obvious way by running classic HPCC technologies (such as HPF, PVM or the pictured MPI) in the third tier. This has been successfully implemented by many groups (ref. [32]) to provide web interfaces to parallel computing systems. Although this provides important commodity services to the user based on Java and JavaScript client interfaces, the approach only addresses the parallel computer as a single object and essentially implies the "host-node" model of parallel programming. Note that in figs. 12 and 13, we draw various nodes and the host as separate entities. These represent logically distinct functions but the physical implementation need not reflect the distinct services. In particular, two or more capabilities can be implemented on the same sequential or SMP system.

In fig. 12, we are not exploiting the distributed computing support of commodity technologies for parallel programming. However in fig. 13, we view the parallel computer as a distributed computer with a particularly fast network and integrated architecture.

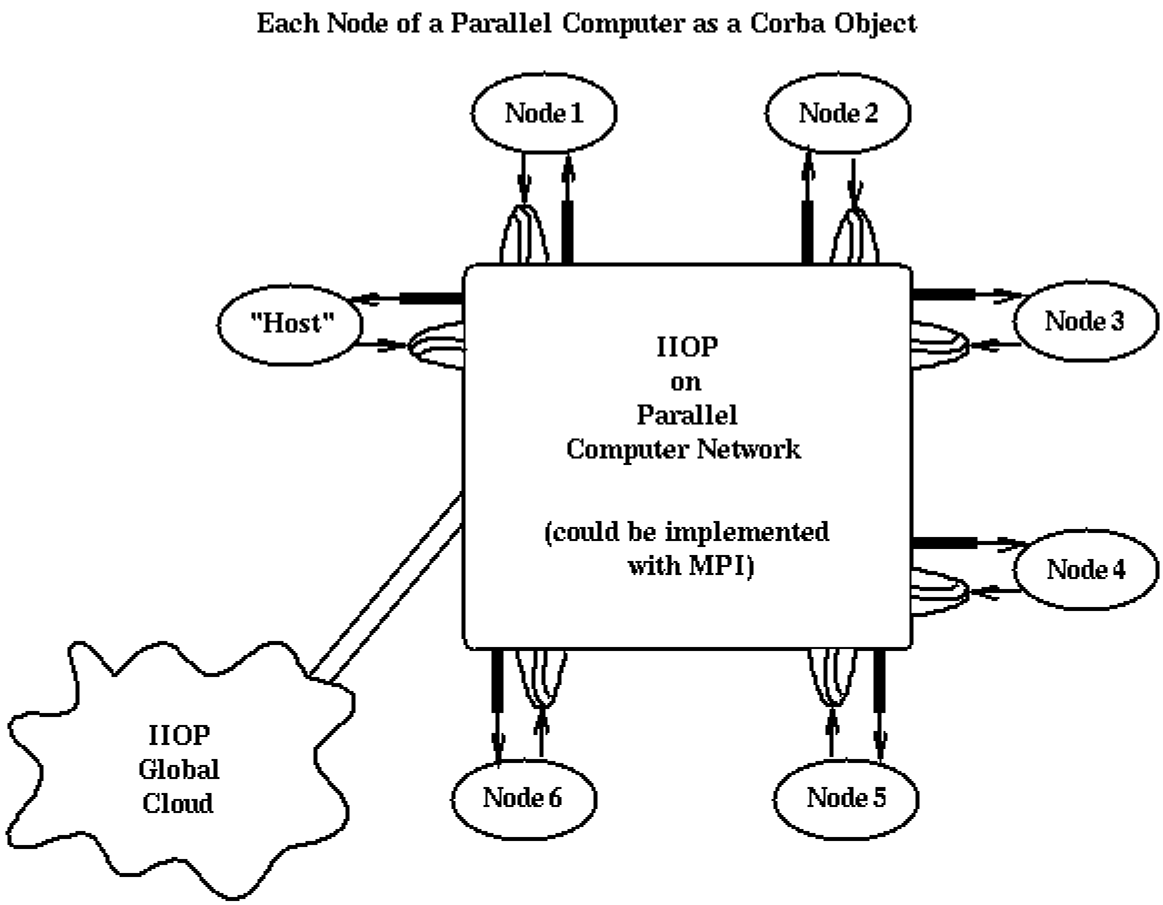

Fig. 13: Each node of a parallel computer instantiated as a Corba object. The "Host" is logically a separate Corba object but could of course be instantiated on the same computer as one or more of the nodes. Using the protocol bridge of fig. 15, one could address objects using Corba with local parallel computing nodes invoking MPI and remote accesses using Corba where its functionality (access to very many services) is valuable.

In fig. 13, each node of the parallel computer runs a Corba ORB (or perhaps more precisely a stripped down ORBlet), Web Server or equivalent commodity server. Now commodity protocols can operate both internally and externally to the parallel machine. This allows a particularly powerful environment where one can uniformly address the full range of commodity and high performance services. Further tools such the visual environment os sec. 7.3 can now be applied to parallel as well as distributed computing. Obviously one should be concerned that this flexibility has been accompanied by a reduction in communication performance from that of fig. 12. Indeed most messaging protocols such as RMI, IIOP and HTTP have unacceptable performance for most parallel computing applications. However we can use the ideas of sec. 5 to obtain good performance with a suitable binding of MPI and PVM to the commodity protocols. In fig. 14, we redraw figs 4 and 5 in a fashion to demonstrate the analogy to fig 13.

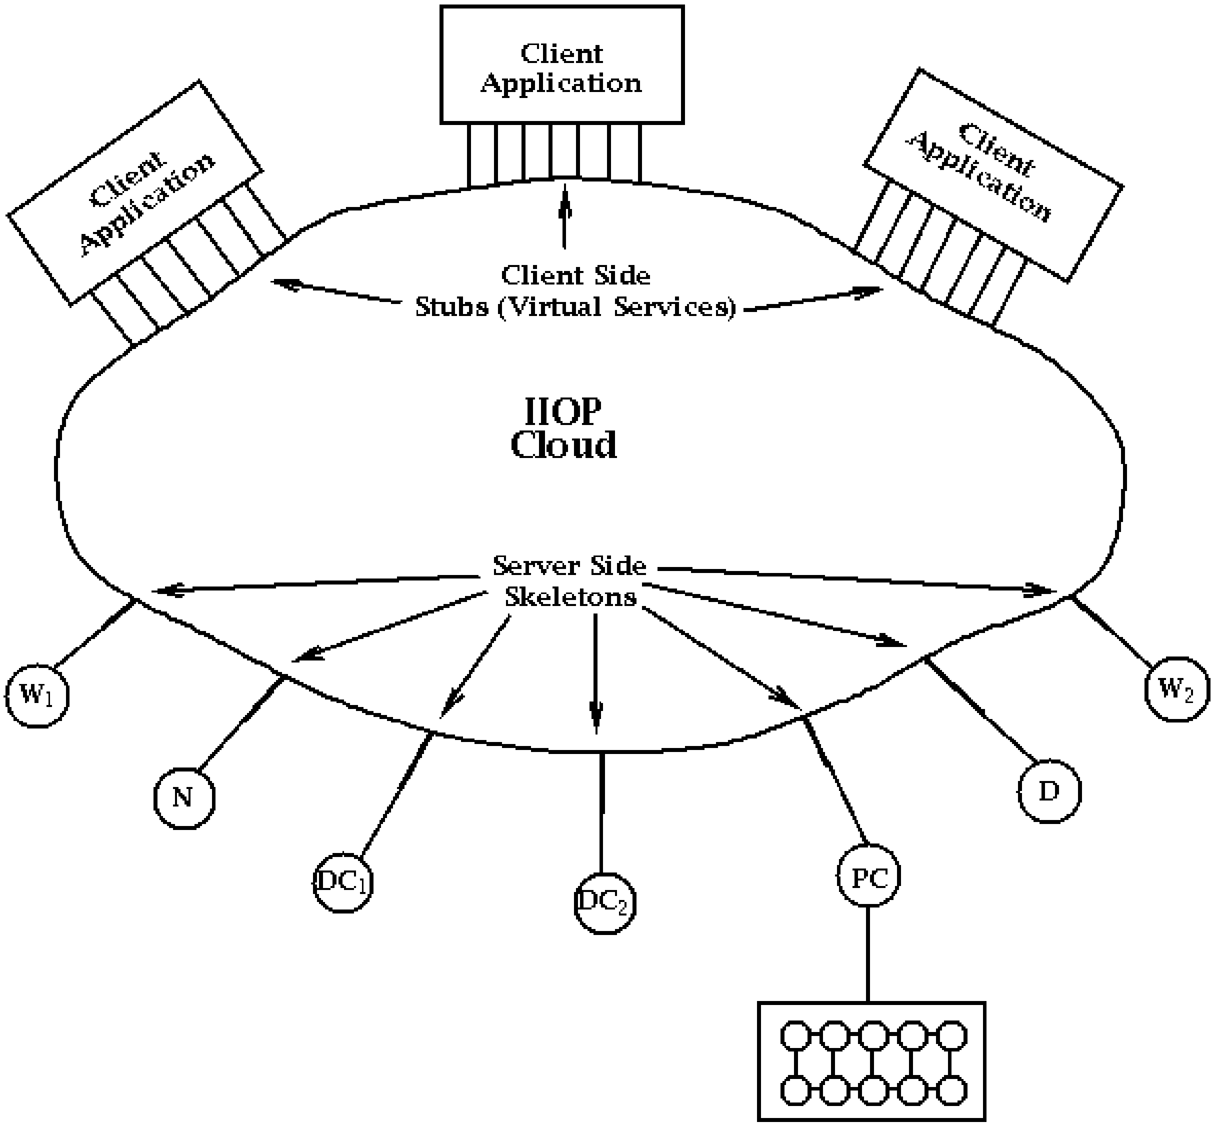

Fig. 14: Pure Corba architecture for the heterogeneous DcciS services of fig. 2. There is a similar Java version of this using RMI and JDBC with of course the linked application being restricted to Java code. Corba and the analogous COM solution are cross-language solutions.

In fig. 15, we extend the previous figure to show an approach to high performance which uses a separation between messaging interface and implementation. The bridge shown in this figure, allows a given invocation syntax to support several messaging services with different performance-functionality tradeoffs.

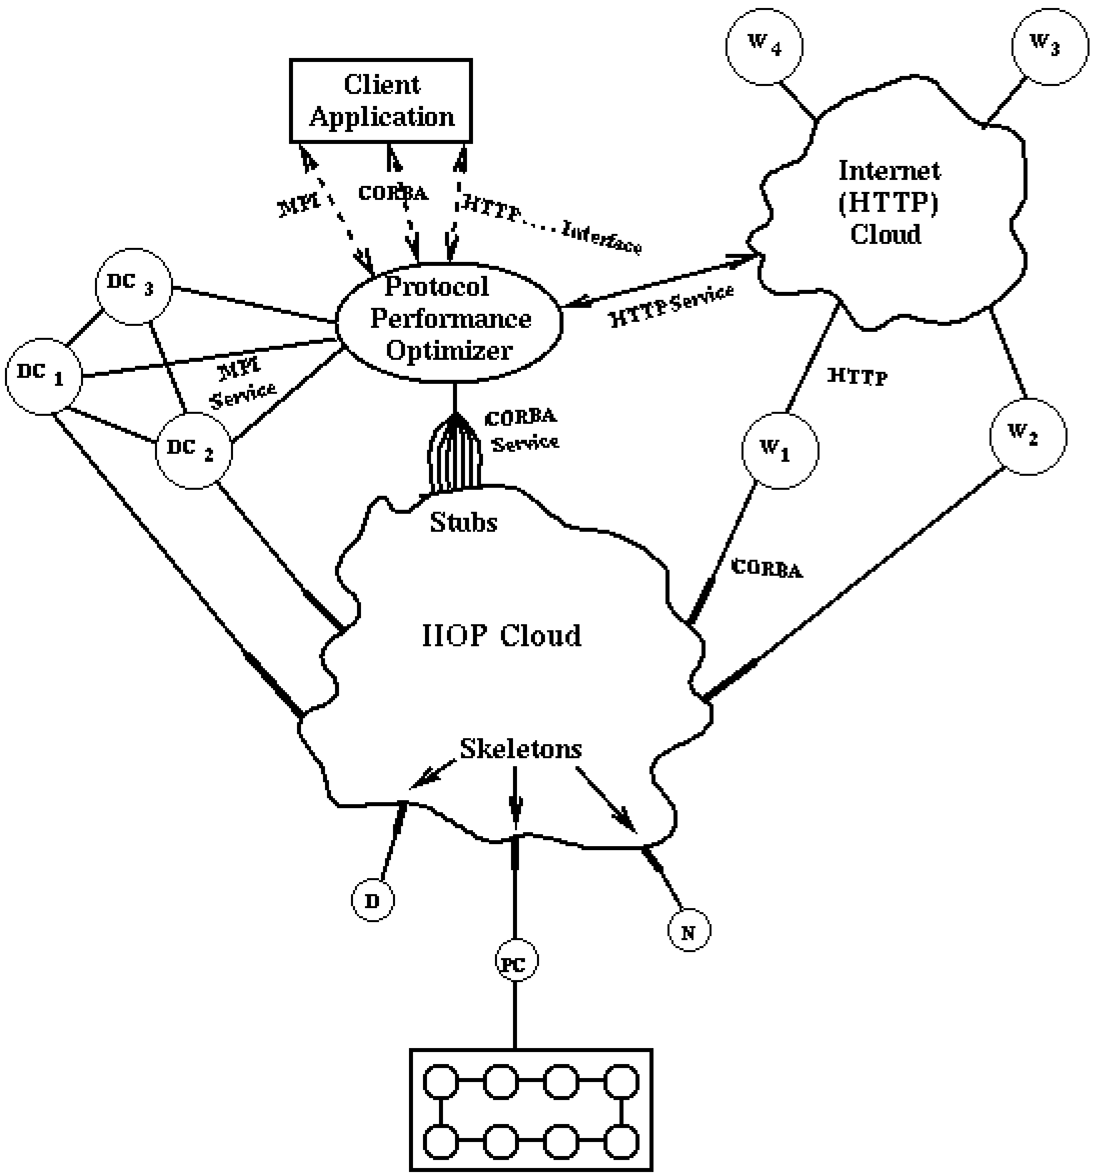

Fig 15: A message optimization bridge allows MPI (or equivalently Nexus, Globus or PVM) and commodity technologies to coexist with a seamless user interface.

In principle, each service can be accessed by any applicable protocol. For instance a Web Server or database can be accessed by HTTP or Corba; a network server or distributed computing resource supports HTTP Corba or MPI. Of course one can substitute equivalent commodity (RMI, COM) or HPCC technologies (PVM, PVMPI, Nexus, Globus) in the above discussion. Note that there are two ways of linking MPI and Corba. Firstly there is the MPI function call that actually calls a Corba stub; secondly a Corba invocation can be trapped and replaced by an optimized MPI implementation. We are investigating this at NPAC in the context of a Java MPI linkage which raises already questions about extending MPI to handle more general object data types. One could for instance extend the MPI communicator field to indicate a preferred protocol implementation. This preference can be set using the mechanism of sec. 5. As stressed by Gannon, there are important research issues in efficient object serialization needed for a high performance implementation of fig. 15.

7.2: HPJava -- Java as a High Performance Computing Language

We have in fact discussed many critical uses of Java in both client interfaces and tier-2 servers to high performance systems. Here we discuss its direct use as a scientific and engineering programming language. We first discuss Java's possible role as the basic programming language for science and engineering - taking the role now played by Fortran 77, Fortran 90, and C++. In our 3-tier architecture, this is the use of Java in tier-3 engineering and science applications or in the CORBA HPCC vertical facility. We now discuss this controversial area. One of Java's important advantages over other languages is that it will be learnt and used by a broad group of users. Java is already being adopted in many entry-level college programming courses and will surely be attractive for teaching in middle or high schools. Java is a very social language as one naturally gets Web pages from one's introductory Java exercises that can be shared with one's peers. We have found this as a helpful feature for introductory courses. Of course, the Web is the only real exposure to computers for many children, and the only languages, to which they are typically exposed today, are Java, JavaScript, and Perl. We find it difficult to believe that entering college students, fresh from their Java classes, will be willing to accept Fortran, which will appear quite primitive in contrast. C++ as a more complicated systems-building language may well be a natural progression, but although quite heavily used, C++ has limitations as a language for simulation. In particular, it is hard for C++ to achieve good performance on even sequential and parallel code, and we expect Java not to have these problems.

In fact, let us now discuss performance, which is a key issue for Java. As discussed in the two workshops on Java for Science and Engineering computation (Ref. 27), there seems little reason why native Java compilers, as opposed to current portable JavaVM interpreters or Just in Time compilers (JIT), cannot obtain comparable performance to C or Fortran compilers. Difficulties in compiling Java include its rich exception framework that could restrict compiler optimizations. Users would need to avoid complex exception handlers in performance critical portions of a code. Other important issues with Java include the lack of any operator overloading, which could allow efficient elegant handling of Fortran constructs like COMPLEX. There is much debate on Java's rule (philosophy?) that code not only run everywhere but give the same value on all machines. This inhibits optimization on machines such as the Intel Pentium that include multiply add instructions with intermediate results stored to higher precision than final values of individual floating point operations.

An important feature of Java is the lack of pointers and their absence, of course, allows much more optimization for both sequential and parallel codes. Optimistically, we can say that Java shares the object-oriented features of C++ and the performance features of Fortran. One interesting area is the expected performance of Java interpreters (using just in time techniques) and compilers on the Java bytecodes (Virtual Machine). Here, we find today perhaps a factor of 3-10 lower performance from a PC JIT compiler compared to C compiled code and this can be expected to improve to become "only" a factor of 2 performance hit. As described above, with some restrictions on programming style, we expect Java language or VM compilers to be competitive with the best Fortran and C compilers. Note that we can also expect a set of high performance ``native class'' libraries to be produced that can be downloaded and accessed by applets to improve performance in the usual areas one builds scientific libraries.

In order to discuss parallel Java, we consider four forms of parallelism seen in applications.

1.Data Parallelism

Here we refer to natural large-scale parallelism found from parallel updates of grid-points, particles and other basic components in scientific computations. Such parallelism is supported in Fortran by either high-level data parallel HPF or at a lower level Fortran plus message passing (MPI). Java does not have any built in parallelism of this type, but at least the lack of pointers means that natural parallelism is less likely to get obscured. There seems no reason why Java cannot be extended to high level data parallel form (HPJava) in a similar way to Fortran (HPF) or C++ (HPC++). This can be done using threads on shared memory machines as pioneered at Indiana while in distributed memory machines, one must use message passing. Here the situation is clearly satisfactory for Java as the language naturally supports inter-program communication, and the standard capabilities of high-performance message passing are being implemented for Java. In section 7.1, we pointed out that there an interesting relationship between the Java-CORBA and Java MPI linkage. At NPAC, we are focussing on SPMD rather than the HPF style approach to data parallel Java. This corresponds to providing Java with a powerful runtime to support data parallel applications.

2.Modest Grain Size Functional Parallelism

Here we are thinking of the type of parallelism used when computation and I/O operation are overlapped as exploited extensively by web browsers. This parallelism is built into the Java language with threads but has to be added explicitly with libraries for Fortran and C++.

3.Object Parallelism

This is quite natural for C++ or Java where the latter can use the applet mechanism to portably represent objects. We have already discussed this in sections 6.1 and 6.2 for shared objects and the large class of military simulations that use large scale object based models.

4.Metaproblems

This is the parallelism in applications that are made up of several different sub-problems which themselves may be sequential or parallel. We have already discussed in the earlier sections, the power of Java in this case for overall coarse grain software integration. This is use of Java in CORBA and web servers in tier 2 and is explicitly discussed in sec. 5.1 to link application components together.

In summary, Java directly addresses three of the four forms of parallelism described above. In these areas, it seems superior to other languages. Java needs to be augmented to fully support data parallelism but so do Fortran and C++!

Interpreted Environments

Java and Web technology suggest new programming environments that integrate compiled and interpreted or scripting languages. In fig. 16, we show a system built at NPAC (ref. [37]) which uses an interpreted Web client interacting dynamically with compiled code through a typical tier 2 server. This happens to be designed for HPF back-end but the architecture is independent of the back-end language. The Java or JavaScript front-end holds proxy objects produced by an HPF front-end operating on the back-end code. These proxy objects can be manipulated with interpreted Java or JavaScript commands to request additional processing, visualization and other interactive computational steering and analysis. Note that for compiled (parallel) Java, use of objects (as opposed to simple types in the language) has probably unacceptable overhead. However they are well used in interpreted front-ends where object references are translated into efficient compiled code. We see such hybrid architectures as quite attractive and warranting further research.

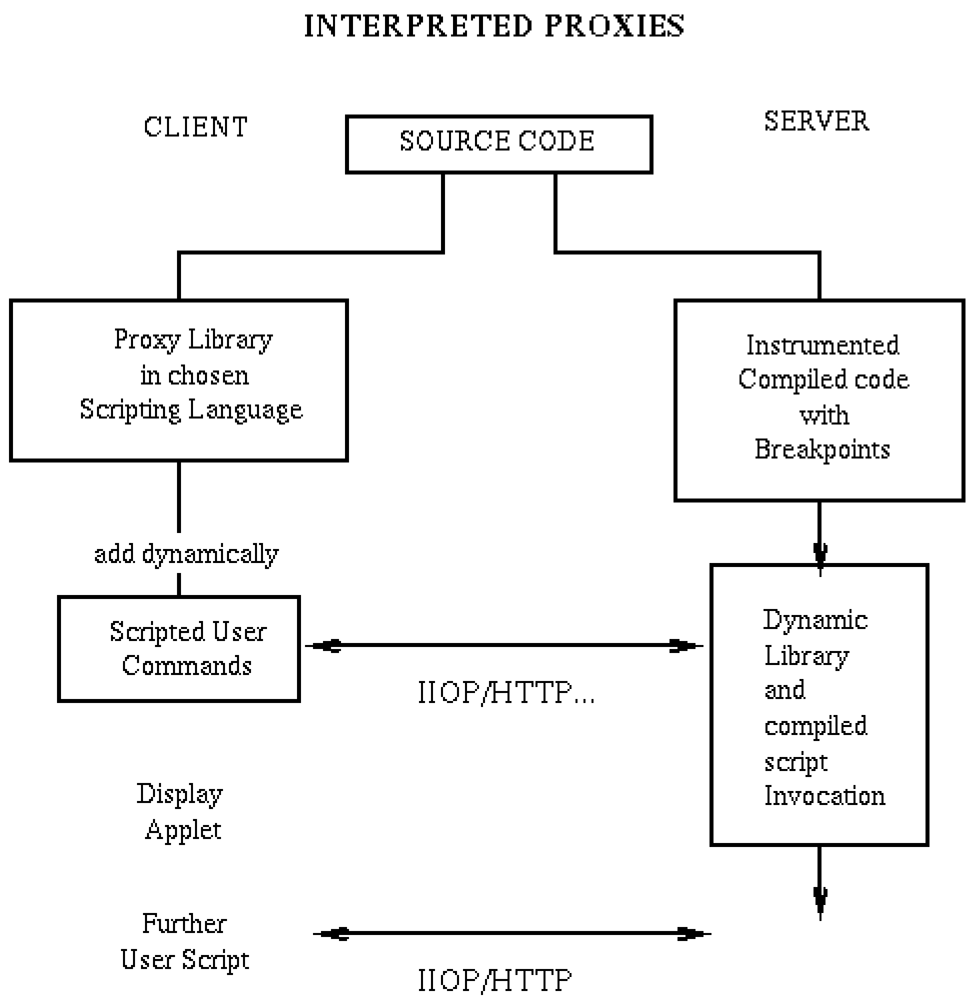

Fig. 16: An architecture for an interpreted Java front-end communicating with a tier 2 server controlling dynamically an HPCC backend.

In summary, we see that Java has no obvious major disadvantages and some clear advantages compared to C++ and especially Fortran as a basic language for large-scale simulation and modeling. Obviously, we should not and cannot port all our codes to Java. Rather, we can start using Java or more precisely Javabeans for wrappers and user interfaces. As compilers get better, we expect users will find it more and more attractive to use Java for new applications. Thus, we can expect to see a growing adoption by computational scientists of Web technology in all aspects of their work.

8: Related Work and References

The base material on CORBA can be found on the OMG Web site [1]. This includes OMG Formal Documentation [2], Recently Adopted Specifications and The Technical Committee Work in Progress [3], which offers up-to-date on-line information on the individual RFPs and their adoption process. One such RFP of a particular relevance for this Chapter, CORBA Components [4], has been recently posted by OMG in response to the Position Paper [5] by IBM, Netscape, Oracle and SunSoft, with mid November '97 as the first submission deadline. It is expected that this activity will result in JavaBeans based ComponentWare model for CORBA. Primary source of information on JavaBeans is the JavaSoft Web site [6]. See also the recent O'Reilly book by Robert Englander [7]. A good recent reference on Microsoft COM (Component Object Model) is Microsoft's Press book by Dale Rogerson [8]. CORBA/COM integration model is specified in the Core CORBA 2.1 document [9], Chapters 14 (Interworking), 15 (COM/CORBA mapping) and 16 (OLE Automation/CORBA mapping). A good overview of CORBA/Java integration and the Object Web concepts can be found in the recent book Robert Orfali and Dan Harkey [10]. The two currently most popular commercial Java ORBs are: OrbixWeb by IONA [11] and VisiBroker for Java by Visigenic [12]. The first public domain ORBs became recently available such as JacORB [13] by University of Berlin, omniBroker by Olivetti and Oracle Labs [14] or Electra by Olsen & Associates [15]. These public domain ORB's facilitate several ongoing research projects on using CORBA for reliable distributed or/and high performance computing which we list below.

Nile, a National Challenge Computing Project [16] develops distributed computing solution for the CLEO High Energy Physics experiment using a self-managing, fault-tolerant, heterogeneous system of hundreds of commodity workstations, with access to a distributed database in excess of about 100 terabytes. These resources are spread across the United States and Canada at 24 collaborating institutions. NILE is CORBA based and it uses the Electra ORB.

Douglas Schmidt, Washington University, conducts research on high performance implementations of CORBA [17], geared towards real-time image processing and telemedicine applications on workstation clusters over ATM. His high performance ORB - TAO [18] - based on optimized version of public domain IIOP implementation from SunSoft outperforms commercial ORB's by factor 2-3. Steve Vinoski, IONA and Douglas Schmidt address current R&D topics on the use of CORBA for distributed computing in their C++ Report column [19]. Richard Muntz, UCLA, explores the use of CORBA for building large-scale object based data mining systems. His OASIS (Open Architecture Scientific Information System)[20] environment for scientific data analysis allows to store, retrieve, analyze and interpret selected datasets from a large collection of scientific information scattered across heterogeneous computational environments of Earth Science projects such as EOSDIS.

NPAC is developing a public domain Java based IIOP and HTTP server, WORB [21], with the alpha release planned by the end of '97 [21]. New anticipated developments in CORBA based distributed computing include emergent CORBA facilities in specialized areas such as Workflow [22] or Distributed Simulations [23].

An elementary overview of "classic" Web Technologies can be found in reference [24]. Reference [25] is one of a set of recent interesting Byte articles on COM, CORBA and Javabeans in the commercial 3-tier model. Reference [26] is the beginning of a compilation of input for defining interfaces for a horizontal CORBA HPcc facility or more precisely it is aimed at one aspect – a seamless interface for users to computing systems. Reference [32] is a group activity aimed at this.

Reference [27] is a compilation of resources relevant to use of Java in Computational Science and Engineering and in particular has proceedings of two workshops in this area. References [28-31,33-37] are from NPAC and describe some aspects of HPcc activities there.

1. Object Management Group, http://www.omg.org

2. OMG Formal Documentation, http://www.omg.org/library/specindx.htm

3. OMG TC Work in Progress, http://www.omg.org/library/schedule.htm

4. "CORBA Component Model RFP",

http://www.omg.org/library/schedule/CORBA_Component_Model_RFP.htm

5. "*CORBA Component Imperatives*" - a position paper by IBM, Netscape,

Oracle and SunSoft, http://www.omg.org/news/610pos.htm

6. JavaBeans, http://www.javasoft.com/beans/

7. "*Developing JavaBeans*" by Robert Englander, O'Reilly & Associates, June

'97, ISBN: 1-56592-289-1.

8. "*Inside COM - Microsoft's Component Object Model*" by Dale Rogerson,

Microsoft Press, 1997, ISBN: 1-57231-349-8.

9. CORBA 2.0/IIOP Specification, http://www.omg.org/corba/c2indx.htm

10. "*Client/Server Programming with Java and CORBA*" by Robert Orfali and

Dan Harkey, Wiley, Feb'97, ISBN: 0-471-16351-1

11. OrbixWeb for Java from IONA, http://www.iona.com

12. VisiBroker for Java from Visigenic, http://www.visigenic.com

13. JacORB by Freie Universität Berlin, http://www.inf.fu-berlin.de/~brose/jacorb/

14. omniORB2 by Olivetti and Oracle Research Laboratory

http://www.orl.co.uk/omniORB/omniORB.html

15. The Electra Object Request Broker, http://www.olsen.ch/~maffeis/electra.html

16. Nile: *National Challenge Computing Project* http://www.nile.utexas.edu/

17. "*Research on High Performance and Real-Time CORBA*" by Douglas Schmidt,

http://www.cs.wustl.edu/~schmidt/corba-research-overview.html

18. "Real-time CORBA with TAO (The ACE ORB)" by Douglas Schmidt,

http://www.cs.wustl.edu/~schmidt/TAO.html

19. "*Object Interconnections*" by Steve Vinoski, column in C++ Report,

http://www.iona.com/hyplan/vinoski/

20. E. Mesrobian, R. Muntz, E. Shek, S. Nittel, M. LaRouche, and M. Krieger, "*OASIS: An Open Architecture Scientific Information System*," 6th International Workshop on Research Issues in Data Engineering, New Orleans, La., February, 1996. See also

http://techinfo.jpl.nasa.gov/JPLTRS/SISN/ISSUE36/MUNTZ.htm

21. WORB - Web Object Request Broker, http://osprey7.npac.syr.edu:1998/iwt98/projects/worb

22. Workflow Mangement Coalition, http://www.aiai.ed.ac.uk/project/wfmc/

23. *High Level Architecture and Run-Time Infrastructure* by DoD Modeling and Simulation Office (DMSO), http://www.dmso.mil/hla

24. Geoffrey Fox, *"Introduction to Web Technologies and Their Applications"*, Syracuse report SCCS-790. http://www.npac.syr.edu/techreports/html/0750/abs-0790.html

25. Article in August 1997 Byte on *3-tier commercial computing model*. http://www.byte.com/art/9708/sec5/art1.htm

26. Mark Baker Portsmouth, *Collection of Links relevant to HPcc Horizontal CORBA facility and seamless interfaces to HPCC computers*, http://www.sis.port.ac.uk/~mab/Computing-FrameWork/

27. *Compilation of References to use of Java in Computational Science and Engineering including proceeding of Syracuse (December 96) and Las Vegas (June 97) meetings*. http://www.npac.syr.edu/projects/javaforcse. This includes a summary of discussion of June 21,97 meeting at Las Vegas on *Issues impacting use of Java in Scientific computation*:

http://www.npac.syr.edu/projects/javaforcse/june21summary.html

28. Geoffrey Fox and Wojtek Furmanski, "*Petaops and Exaops: Supercomputing on the Web*", IEEE Internet Computing, 1(2), 38-46 (1997); http://www.npac.syr.edu/users/gcf/petastuff/petaweb

29. Geoffrey Fox and Wojtek Furmanski, "*Java for Parallel Computing and as a General Language for Scientific and Engineering Simulation and Modeling*", Concurrency: Practice and Experience 9(6), 4135-426(1997). Web version in ref. [27].

30. Compilation of *NPAC Activities in Web-based HPcc*, http://www.npac.syr.edu/projects/webspace/webbasedhpcc/

31. Geoffrey Fox and Wojtek Furmanski, "*Use of Commodity Technologies in a Computational Grid*", chapter in book to be published by Morgan-Kaufmann and edited by Carl Kesselman and Ian Foster.

32. Jim Almond at ECMWF in Reading, England "*Resource for seamless computing*", http://www.ecmwf.int/html/seamless/

33. David Bernholdt, Geoffrey Fox and Wojtek Furmanski, *"Towards High Performance Object Web Based FMS"*, White Paper for ARL MSRC PET Program, Sept. 97.

34. D. Bhatia, V.Burzevski, M.Camuseva, G.Fox, W.Furmanski, and G. Premchandran, *"WebFlow -- A Visual Programming Paradigm for Web/Java based coarse grain distributed computing*", Concurrency Practice and Experience 9,555-578(1997).

35. L. Beca, G. Cheng, G. Fox, T. Jurga, K. Olszewski, M. Podgorny, P. Sokolowski and K. Walczak *"Java enabling collaborative education, health care and computing"*, Concurrency Practice and Experience 9,521-534(97). See http://trurl.npac.syr.edu/tango/

36. K.Dincer, G. Fox, *"Using Java and JavaScript in the Virtual Programming Laboratory: a web-based parallel programming environment"*, Concurrency Practice and Experience 9,521-534(97). http://www.npac.syr.edu/users/dincer/papers/vpl/

37. E. Akarsu, T. Haupt, "Integrated Environment for HPF Compiler and Interpreter", http://www.npac.syr.edu/projects/hpfi/

9: Glossary

CORBA (Common Object Request Broker Architecture)

An approach to cross-platform cross-language distributed object developed by a broad industry group, the OMG. CORBA specifies basic services (such as naming, trading, persistence) the protocol IIOP used by communicating ORBS, and is developing higher level facilities which are object architectures for specialized domains such as banking.(fig. 6)

COM (Common Object Model)

Microsoft's windows object model, which is being extended to distributed systems and multi-tiered architectures. ActiveX controls are an important class of COM object, which implement the component model of software. The distributed version of COM used to be called DCOM.

ComponentWare

An approach to software engineering with software modules developed as objects with particular design frameworks (rules for naming and module architecture) and with visual editors both to interface to properties of each module and also to link modules together.

DIS (Distributed Interactive Simulation)

Original framework developed in such projects as SIMNET, to support FMS and IMT applications. HLA and RTI are superceding DIS.

FMS (Forces Modeling and Simulation)

A major class of DoD (Department of Defense) applications, which are often termed "war-games". FMS either involves virtual time event driven simulations or real-time simulations involving people and military vehicles or instruments in the loop. Typically FMS uses a geographically distributed set of computers to simulate a geographically distributed application.

HLA (High Level Architecture)

Current object architecture for FMS and IMT applications. HLA (and RTI) are expected to become a CORBA facility.

HPcc (High Performance commodity computing)

NPAC project to develop a commodity computing based high performance computing software environment. Note that we have dropped "communications" referred to in the classic HPCC acronym. This is not because it is unimportant but rather because a commodity approach to high performance networking is already being adopted. We focus on high level services such as programming, data access and visualization that we abstract to the rather wishy-washy "computing" in the HPcc acronym.

HPCC (High Performance Computing and Communication)

Originally a formal federal initiative but even after this ended in 1996, this term is used to describe the field devoted to solving large-scale problems with powerful computers and networks.

HTTP (Hyper Text Transport Mechanism)

A stateless transport protocol allowing control information and data to be transmitted between web clients and servers.

IMT (Integrated Modeling and Testing)

A major class of DoD applications aimed at detailed simulation of the capabilities of military subsystems such as a new tank. IMT simulations tend to be similar to FMS but include more detail for the component being tested. IMT can involve real systems in the loop, and is of growing importance as DoD cuts back on expensive physical tests. Civilian analogues are car crash simulations used in the automobile industry and computational fluid dynamics simulations substituting for wind tunnels in the aeronautics industry.

IIOP (Internet Inter Orb Protocol)

A stateful protocol allowing CORBA ORB's to communicate with each other, and transfer both the request for a desired service and the returned result.

JDBC (Java Data Base Connection)

A set of interfaces (Java methods and constants) in the Java 1.1 enterprise framework, defining a uniform access to relational databases. JDBC calls from a client or server Java program link to a particular "driver" that converts these universal database access calls (establish a connection, SQL query, etc.) to particular syntax needed to access essentially any significant database.

Javabean

Part of the Java 1.1 enhancements defining design frameworks (particular naming conventions) and inter Javabean communication mechanisms for Java components with standard (Bean box) or customized visual interfaces (property editors). Enterprise Javabeans are Javabeans enhanced for server side operation with capabilities such as multi user support. Javabeans are Java's component technology and in this sense are more analogous to ActiveX than either COM or CORBA. However Javabeans augmented with RMI can be used to build a "pure Java" distributed object model.

Object Web

The evolving systems software middleware infrastructure gotten by merging CORBA with Java. Correspondingly merging CORBA with Javabeans gives Object Web ComponentWare. This is expected to compete with the COM/ActiveX architecture from Microsoft.

OMG (Object Management Group)

OMG is the organization of over 700 companies that is developing CORBA through a process of call for proposals and development of consensus standards.

ORB (Object Request Broker)

Used in both clients and servers in CORBA to enable the remote access to objects. ORB's are available from many vendors and communicate via the IIOP protocol.

RMI (Remote Method Invocation)

A somewhat controversial part of Java 1.1 in the enterprise framework which specifies the remote access to Java objects with a generalization of the UNIX RPC (Remote Procedure Call).

RTI (Run Time Infrastructure)

Run time defined to support HLA compliant simulations including "federated" interoperable simulations.

Web Client

Originally web clients displayed HTML and related pages but now they support Java Applets that can be programmed to give web clients the necessary capabilities to support general enterprise computing. The support of signed applets in recent browsers has removed crude security restrictions, which handicapped previous use of applets.

Web Servers

Originally Web Servers supported HTTP requests for information - basically HTML pages but included the invocation of general server side programs using the very simple but arcane CGI - Common Gateway Interface. A new generation of Java servers have enhanced capabilities including server side Java program enhancements (Servlets) and support of stateful permanent communication channels.