Nationaal Lucht- en Ruimtevaartlaboratorium
National Aerospace Laboratory NLR

NLR TP 97601

SPINE: A Practical and Holistic Approach to Metacomputing

E.H. Baalbergen

DOCUMENT CONTROL SHEET

	ORIGINATOR'S REF. TP 97601 U		SECURITY CLASS. Unclassified

ORIGINATOR
National Aerospace Laboratory NLR, Amsterdam, The Netherlands

TITLE
SPINE: A Practical and Holistic Approach to Metacomputing

PRESENTED AT
HPCN Europe '98

AUTHORS	DATE	pp	ref
E.H. Baalbergen	971128	15	7

DESCRIPTORS

Architecture (computers)	Distributed processing
Computer networks	Multiprocessing (computers)
Computer programs	Parallel computers
Data bases	Software engineering
Data base management systems	Supercomputers

ABSTRACT
Enterprises constantly invest in their computing infrastructure in order to lower the cost of product engineering. On the one hand, they increase the available computational power, while on the other hand they burden the engineers with the need to constantly familiarize with the changing computing infrastructure. The need to familiarize diverts the engineers' attention from the actual problems to be solved. This paper discusses the major characteristics and issues involved with usage of today's computer networks. It describes efforts in the field of metacomputing that try to hide the networking details from the end user. SPINE is presented as a practical and complete approach to metacomputing. SPINE provides the middleware and tools that enable the recources available from a network to be presented to the users in a form of a single, virtual computer. Engineers operate the virtual computer through a user-orientated graphical user interface which may be customized for a specific application area.

Summary

Enterprises constantly invest in their computing infrastructure in order to lower the cost of product engineering. On the one hand, they increase the available computational power, while on the other hand they burden the engineers with the need to constantly familiarize with the changing computing infrastructure. The need to familiarize diverts the engineers' attention from the actual problems to be solved. This paper discusses the major characteristics and issues involved with usage of today's computer networks. It describes efforts in the field of metacomputing that try to hide the networking details for the end user. SPINE is presented as a practical and complete approach to metacomputing. SPINE provides the middleware and tools that enable the resources available from a network to be presented to the users in the form of a single, virtual computer. Engineers operate the virtual computer through a user-oriented graphical user interface which may be customized for a specific application area.

This page is intentionally left blank.

1 Introduction

In order to be more efficient, and to lower the cost of product engineering and know-how management, enterprises constantly invest in their computing infrastructure. New computers are added to the network, present computers are frequently replaced by more powerful systems, and the network resources are constantly replaced by more advanced and faster ones. However, although the computing power increases through these investments, a growing percentage of the potential power remains unused. On the one hand, an enterprise increases its computational power, while on the other hand it burdens its engineers with the need to constantly familiarize with the changing computing infrastructure.

The efficiency loss is due to the growing gap between the engineer and the computing infrastructure. The effective and efficient use of the potentials of today’s computing power requires thorough knowledge of the resources involved. Consequently, an engineer’s attention may be diverted from solving a problem proper to rephrasing the problem in order to solve it using the available computing infrastructure. Put in other words, the engineer usually must adapt himself or herself to the computing infrastructure rather than vice versa. The constantly evolving and changing infrastructure makes the situation even worse, since engineers must continue to remain familiarized in order to exploit the potentials. A typical problem for the engineer is the need to know how to operate the individual computers connected to the network, as well as how to organize and manage the information in the network, in order to use the available resources efficiently.

Metacomputing is generally recognized - both by software vendors and researchers - as solution to support the engineer in the end use of a network of computers. The key idea is to provide the end user of a network with the potentials (e.g., computing power and storage facilities) of the network, in the form of a single, virtual computer, a so-called “metacomputer”. Details concerning the networking aspects involved are hidden from the user of a metacomputer. Most developments thus far, however, have resulted in partial solutions for metacomputing. Solutions such as networked file systems, systems for scheduling jobs in a network, and tools for remote execution, are generally available, and considered useful, but cover only part of metacomputing. They still leave many technical details to the end user.

This paper motivates and concisely describes SPINE, which is a system that aims at providing a complete, holistic solution for metacomputing. SPINE provides the “middleware” and the tools that support the realization of application-area specific working environments. A working environment provides the end user with an easy-to-use graphical desk-top system. It facilitates operation of the computer network in a user-friendly and application-oriented fashion, through point-and-click and drag-and-drop operations on windows and icons. A working environment exhibits the single-computer - metacomputer - model. Details concerning the network, remote execution of programs, and file transfers between computers are hidden from the end user.

Section 2 discusses general characteristics of contemporary computer networks, together with the observations that have led to the development of SPINE. It concludes with a concise overview of current research and products in the field of metacomputing that have been motivated by the use of today's computer networks. The SPINE model and the key ideas behind SPINE are presented in section 3. Section 4 describes the implementation of SPINE. Conclusions are contained in section 5.

2 Computer networks

Computer networks found nowadays in industry and research enterprises typically comprise personal desk-top systems as well as centralized, generally available systems. The desk-top systems include PCs (usually running a PC operating system such as MS Windows 95 or NT), graphics workstations (usually running a UNIX-based system such as SGI's IRIX, SUN's SunOS, and HP's HP-UX), and graphics ("X") terminals connected to some UNIX host system. The centralized systems typically include general-purpose as well as dedicated compute, file, and input/output servers, such as mainframes, supercomputers, parallel computers, plotters, and printers. These systems are usually connected with each other, and with other (national and international) networks, through high-speed networks.

An example computer network is that of NLR [5]. The network serves the two establishments of NLR in the Netherlands, and includes:

- a supercomputer for high-performance vector processing: an NEC SX-4 (16 processors, 32 GFLOP/s peak performance, shared memory);
- several servers for scalar processing and software development;
- parallel compute servers for research purposes;
- access possibilities up to 34 Mb/s IP and ATM;
- more than a thousand workstations, PCs and other terminals;
- an integrated software engineering environment, including tools and building blocks for rapid prototyping up to and including full-scale development, and tools for management and control of data, documents and software.

NLR's network is used for design and development of products - including software products - applied in aerospace research and industry [3].

Networks have become more powerful through the years. In addition, the availability of high-capacity wide-area networking facilities has emerged. Internet, ISDN, ATM, and fire wall techniques have made it feasible for enterprises to interconnect their networks, and hence to couple and share their computing power (e.g., [2]).

An general important observation with respect to the existing computer networks is the visibility of networking details to the end users. The integration of interconnected computers is usually accomplished at operating-system level - usually based on TCP/IP and TCP/IP-based software - rather than at end-user level. This means that access to, and usage of other computers is technically supported, but requires special arrangements and the use of special tools by the end user. The network in particular often confronts engineers with the following issues:

- Finding out on which computer in the network a specific program may or must run on.
- Remote login and remote execution of commands, facing the user with the use of special "remote" commands (e.g., *telnet*, *rsh*), and with aspects such as authentication and accounting.
- Fire walls: an enterprise, or a division within an enterprise, may decide to protect its network using a fire wall, which carefully controls access to the network from outside the network.
- The need to explicitly transfer files between computers, or to store and organize data on particular file systems shared between the computers involved.
- System failures, such as hardware and operating systems crashes, network link failures, and capacity problems of systems involved in a computation (full disks and full process tables).
- Hardware heterogeneity: the representation of data values in memory and hence on (binary) files may be incompatible on different computers, due to different strategies with respect to byte ordering, representation of floating-point numbers, and alignments. If binary data needs to be exchanged between heterogeneous computers, conversions are required.
- Operating-system heterogeneity: the engineer usually needs to operate the individual computers in a network, which may confront him or her with different operating systems, or variants of the same operating system.

In addition, use of today's computer networks by engineers in practice raises the following issues:

- The engineer - usually not being a computer expert - is confronted with UNIX and its peculiarities. Most central systems run a UNIX-based operating system, and require the user to operate the system in command-line mode, and to manipulate the information in terms of files and programs.
- The engineer is typically confronted with a variety of heterogeneous software products. Programs come from different sources, such as commercial, developed in-house, or public domain, and hence usually lack uniformity with respect to user interface, input/output data formats, and documentation.
- To make optimum use of the available computing and storage resources, the engineer may have to put much efforts in reorganizing his or her software and data. For example, full exploitation of the computing power of a vector supercomputer or a distributed-memory parallel computer, usually requires the code or even the problem-solving algorithm to be reorganized, or at least a number of compiler options and directives to be used.
- The engineer typically is faced with organizing, managing, and finding a way through an amount of on-line information (software, documentation, data) that is much larger than that possibly available on a single computer. Information drowning is a potential threat in today's computer networks.

- Most network systems lack, or provide only minimum support for cooperative work. Several “computer supported cooperative work” (CSCW) facilities exist but do support only some aspects of cooperative work. Examples of such facilities are E-mail, video conferencing, Internet (and Internet browsers), and (shared) data bases. Also, concurrent manipulation of information by a team of engineers may threaten the integrity of the information. Access control facilities are required to maintain the consistency of such information.
- Lack of support for quality assurance and quality management. Quality requirements make it essential to apply configuration management and control procedures to product-engineering information.

These issues together make at least occasional use of networks by engineers practically impossible because of the time needed to familiarize.

In summary, today’s computer networks are generally available to the engineers as system-oriented collections of interconnected, individual computers rather than as user-oriented metacomputers that can be customized for particular application areas. The integration of the interconnected computers is accomplished sufficiently at system level, but not at user level. The user needs to adapt to the computing environment, whereas in the ideal situation, the computing environment would be adapted to the user. In addition, the networks lack support for cooperative work. The gap between today’s computing infrastructure and the engineers led to inefficient usage of the available computational power.

The major goal of metacomputing is to make a network of computers look and feel like a single computer, thereby enabling the users to fully exploit the potentials of the network in a user-oriented fashion. A derived, but nevertheless equally important goal is to exploit the computing power to support cooperative work. For an enterprise, metacomputing is of paramount importance for reaching their goal to make their computer infrastructure used efficiently by teams of engineers.

Research and developments into distributed systems and networking have yielded products and concepts that contribute to the realization of metacomputing, such as shared file systems (e.g., NFS, DFS), distributed job management systems (e.g., LSF, CODINE), high-speed and multi-purpose networks, parallel programming support systems (e.g., PVM, MPI), data warehousing and data-mining systems, and CSCW systems such as work-flow management systems. These products, however, cover only part of metacomputing. Research and developments in the field of metacomputing are still continued (e.g., Condor, Polder), but the major goal - the realization of a single, user-oriented, and effective computer, a “metacomputer” - is not yet reached.

3 SPINE model

SPINE is a system that supports the development and operational use of application specific working environments on top of computer networks. A SPINE working environment provides its end user with access to, and efficient use of the resources available from the network, as if these resources are available from one single, virtual computer, a metacomputer. The resources include computing power, data storage capacity, input/output facilities, and information present in the form of tools, software, data, and on-line documentation. The user is confronted with a single, user-oriented and user-friendly computer rather than with a computer network that reveals its peculiarities. Details emerging from the use of multiple systems, networking, and distributed computing are dealt with behind the screen, and remain transparent to the user.

The development of SPINE was motivated by the development of ISNaS in the late 1980s and the early 1990s [7]: the realization of a working environment for numerical flow simulation (CFD) in a computer network - such as NLR's - that is easy to use for the CFD engineer. Instead of developing one specific application-oriented working environment, SPINE was targeted towards a generic solution for the development of application-oriented working environments. The main purpose of SPINE is to support the realization of working environments that:

- are user-oriented, easy to operate by engineers who are neither computer scientists nor computer experts;
- span local-area as well as wide-area, and intra-enterprise as well as inter-enterprise, computer networks;
- make system and network details transparent to the user;
- are easy to tailor for specific application areas;
- are easy to customize to the needs of a user or a group of users, or to the requirements of a project;
- are open and easily extendible, in that existing resources (e.g., tools) can be integrated with minimum effort;
- support quality management, in that tools and facilities are available to support the execution of quality assurance procedures (e.g., [3]), such as software configuration management;
- support cooperative work by multi-disciplinary engineering teams;
- make efficient use of the resources available from a computer network;
- are portable;
- may be used as repository for storing, organizing, sharing, and using product-engineering information (including tools and results).

SPINE provides facilities for the development as well as the operational use of working environments in computer networks. The remainder of this section discusses both aspects.

Development of working environments

Support for the development of working environments comprises a tool for creating an initial minimum working environment, and tools for modifying - tailoring and extending - working environments. The initial working environment consists of a few tools, including a text-file editor, and on-line documentation, organized into a small tree of directories. This working environment may be further developed into an application-specific working environment, by modifying and extending the directory tree through integration of tools, data, documentation, and other electronically available information.

The integration of already available on-line information is simply done by linking or copying the files and directories involved into the directory tree. For example, an interconnected computer's file system, or part thereof, may be linked directly into the directory tree, and is consequently directly accessible for the user of the working environment. The freedom and ease of linking and creating arbitrary files and directories (and trees) into working environments, and the ability to arrange for access permissions to these, allow work spaces to be set up for individual users as well as teams.

The integration of tools and programs available on any host in the network is facilitated by SPINE's Tool Integrator. This facility allows native tools - operating-system specific, commercial, as well as home-brew programs - to be integrated without the need to adapt the native tools proper. It provides a graphical editor that constructs wrappers for native tools, and that supports the combination of several native tools into a single tool. Tool wrappers serve to hide all specific details concerning the activation of native tools, and enables the user to activate tools in a uniform and intuitive way. The Tool Integrator also supports the definition of input parameters to be specified by the caller of the tool (e.g., the end user may be prompted to fill out a form to specify parameters), checks on the number and types of input files to be specified with activation of the tool, and on-line help information about the tool. SPINE, in addition, provides the concept of "tool pack", to support the reuse of tools already integrated. A tool pack is a collection of tools which are readily available for use in a working environment through a simple "import" mechanism.

Operational use of working environments

A SPINE working environment comprises technical solutions for the key goals of metacomputing: a user-oriented desk top, information management facilities, and middleware for managing the network's resources.

The desk top provides the end user of a working environment with a GUI, a graphical shell that facilitates operation of the metacomputer in a user-friendly and application-oriented fashion. Most of the common data management and processing operations are available in the form of point-and-click and drag-and-drop operations on icons in windows. This mode of user interaction originates from the Macintosh, and is a widely accepted solution for user-friendly computing, as can be seen from the popularity of the PC Windows systems. The desk top enables the user to manipulate files and directories, and to activate tools.

A tool - possibly running on different computers, and involving file transfer - can simply be activated by clicking, or dropping input files, on the tool's icon. Additional parameters can be specified by entering values in a form popped up upon tool activation. Help information about specific tools or files can be obtained by dropping the corresponding icons on a help button. The desk top also offers capabilities for data-flow and work-flow driven computing: tools and so-called data containers (placeholders for data files) may be organized into a graph. SPINE's Workflow Editor supports the construction of such graphs, and controls the execution of tools in the graphs. It provides the basic building blocks for definition of work flows.

SPINE supports information management by providing facilities for managing storage of, access to, and modifications of data, documents, software, and related information, and for the exchange of information among users. These facilities include tools for controlling versions of software products, for organizing and manipulating data files, and for definition and usage of electronic forms.

The SPINE middleware serves to "glue" the desk top and information management facilities, and the computing and information resources available from the network, together into one coherent system that may be considered as the operating system of the metacomputer. The major task of the middleware, similar to that of an operating system on a single computer, is to manage the available resources in order to accomplish the single-computer model, and to exploit the potentials of the underlying computer network.

4 SPINE implementation

The implementation of SPINE according to the model as described in the previous section started early 1993 at NLR. SPINE has been used since 1993 for the development of several working environments, initially at NLR, but later also in several national projects (e.g., NICE, [2]) and international projects (e.g., MDO, [6]). Since 1996, the development of SPINE has continued jointly by NLR and NEC. By the end of 1997, the available version of SPINE is version 2.0. SPINE version 2.0 is available for most of today's UNIX variants found in enterprises: HP-UX, SunOS, and IRIX. In addition, it is available for NEC's SX-4 UNIX, SUPER-UX. Although SPINE version 2.0 is targeted to UNIX (TCP/IP-based) networks, it is possible to operate the desk top of working environments from PCs that have NFS and X software installed. NEC has planned to market SPINE version 2.0 early 1998.

SPINE (version 2.0) consists of the following subsystems:

- The *User Shell*, which provides the desk top, the parameter pop-up forms, and the graphical user interfaces for Tool Integrator and Workflow Editor.
- The *Information Management System*, which contains the information management facilities.
- The *Environment Management System*, which provides the facilities for developing working environments, and the SPINE middleware.
- The *Common Tool Set*, which contains tool packs for data-file management (*dfms*), software version control (*sr*), definition of electronic forms (*e-forms*) and search engines (*e-search*), general file and directory manipulation (*common*), and developing software for an SX-4 vector supercomputer (*sx4dev*).

To maintain portability, flexibility, and platform-independence, and to avoid costs for third-party licenses, one of the major constraints applicable to the development of SPINE has been to utilize as many standard, off-the-shelf and public-domain products as possible. For example, the middleware is based on the standard UNIX network facilities *Remote Shell, RSH* (*Secure Shell, SSH*, is supported in order to cope with fire walls), and the *Network File System, NFS*. The graphical user interface software is based on Ousterhout's *Tcl/Tk*. The version-control software is based on Tichy's *Revision Control System, RCS*. In addition, the on-line documentation and electronic forms and search engines are based on HTML as “language”, and *HTTPD* as server.

5 Conclusions

The present paper presents the issues that have motivated the development of SPINE as practical and holistic approach for metacomputing.

SPINE is a practical solution in the sense that SPINE exists, and has proven to be a usable tool for realization of metacomputing in engineering environments. It has been applied successfully for the realization of working environments for numerical flow simulation [2, 4, 5], software engineering, for control engineering [1], and for multi-disciplinary design [6]. In practice, unexperienced engineers were able to utilize working environments within short time.

SPINE is considered to be a holistic approach in the sense that it provides a framework for a complete solution for metacomputing. The framework allows easy incorporation and integration of existing and future partial solutions for metacomputing. Most developments in metacomputing are currently primarily driven by research in computer science. This has resulted in a gap between scientifically well-founded solutions on the one hand, and industrial needs on the other hand. The gap often yielded products that - although being successful and widely accepted - support particular metacomputing aspects only, such as job scheduling and parallel programming, whereas engineers demand complete solutions for metacomputing. The purpose of SPINE is to bridge the gap.

The development of SPINE is currently focused on the provision of extended work-flow management facilities, an enhanced desk top based on objects, full integration of PCs running Windows 95 and Windows NT, and support for integration of CORBA-compliant products.

References

1. M.J.H. Couwenberg, R.J. Cazemier, NLR's CACE working environment ISMuS, in: proc. IEEE int. symp. on computer-aided control system design, Dearborn, MI, USA, 15-18 September 1996.
2. R.J.P. Groothuizen, A NICE HPCN center for flow simulation, extended abstract submitted to HPCN Europe '98.
3. W. Loeve, E.H. Baalbergen, M.E.S. Vogels, H. van der Ven, High performance computing in simulation to fulfil user-driven quality criteria, in: proc. of the DASIA 97 conference on data systems in aerospace, Sevilla, Spain, 26-29 May 1997.
4. W. Loeve, H. van der Ven, M.E.S. Vogels, E.H. Baalbergen, Network middleware illustrated for enterprise enhanced operation, in: proc. of the 6th int. conf. for worldwide industrial research on global enterprise computer applications in production and engineering CAPE '97, Detroit, MI, USA, 5-7 November 1997.
5. W. Loeve, M.E.S. Vogels, H. van der Ven, E.H. Baalbergen, Network middleware for enterprise enhanced operation, in: proc. of the 4th int. conf. on computer integrated manufacturing ICCIM 97, Singapore, 21-24 October 1997.
6. MDO Consortium (1996), http://www.telecall.co.uk/∼srcbae/mdo.
7. M.E.S. Vogels, W. Loeve, Development of ISNaS: an information system for flow simulation in design, in: computer applications in production and engineering (F. Kimural ed.), North Holland, ISBN 0 444 88089 5.