Nationaal Lucht- en Ruimtevaartlaboratorium

National Aerospace Laboratory NLR

NLR TP 97229

Classical HPCN geared to application in industry

W. Loeve, J.Verwer, E. Snijdoodt and A. ten Dam

DOCUMENT CONTROL SHEET

	ORIGINATOR'S REF. NLR TP97299 U		SECURITY CLASS. Unclassified

ORIGINATOR
National Aerospace Laboratory NLR, Amsterdam, The Netherlands

TITLE
Classical HPCN geared to application in industry

PRESENTED AT

AUTHORS	DATE	pp	ref
W. Loeve, J.Verwer, E. Snijdoodt, and A. ten Dam	970421	18	20

DESCRIPTORS

Applications of mathematics	Multiple access
Computational fluid dynamics	Netherlands
Computer networks	Research projects
Computerized simulation	Software tools
Distributed processing	Supercomputers
Environment pollution	Training simuolators

ABSTRACT
In The Netherlands as a result of a national HPCN initiative a Foundation HPCN was established in 1995. The purpose of this Foundation is to stimulate structural and lasting cooperation of universities, technological institutes and industry in economically relevant applications of HPCN. Projects have been selected and since the beginning of 1996 projects are being executed.
In the present paper an overview is given of the eight projects that are being executed. The principles of the HPCN program in The Netherlands is illustrated for flow simulation projects in which the use of the most powerful existing computer servers is essential and made feasible for industry including SMEs. The approach is based on integration of local workstations with remote servers for information management and computing tasks.

Contents

(18 pages in total)

CLASSICAL HPCN GEARED TO APPLICATION IN INDUSTRY

W. Loeve[1], J. Verwer[2], E. Snijdoodt[3], A. ten Dam[4]

Invited paper presented on behalf of The Netherlands Foundation HPCN at the Conference HPCN Europe'97, Vienna, April 1997

ABSTRACT

In The Netherlands as a result of a national HPCN initiative a Foundation HPCN was established in 1995. The purpose of this Foundation is to stimulate structural and lasting cooperation of universities, technological institutes and industry in economically relevant applications of HPCN. Projects have been selected and since the beginning of 1996 projects are being executed.

In the present paper an overview is given of the eight projects that are being executed. The principles of the HPCN program in The Netherlands is illustrated for flow simulation projects in which the use of the most powerful existing compute servers is essential and made feasible for industry including SMEs. The approach is based on integration of local workstations with remote servers for information management and computing tasks.

1 INTRODUCTION

A national HPCN initiative originated by the government resulted in The Netherlands in a national HPCN program. The program aims at structural cooperation of universities, technological institutes and industries in application of HPCN. The Foundation HPCN was established in 1995 to effectuate the objectives. The Foundation so far selected eight projects. An overview of the projects is presented in appendix A. The Foundation also supports development and use of the HPCN infrastructure. An overview of this is presented in Appendix B.

Classical HPCN in the present paper means the application of computers that give at a certain moment in time the maximum attainable performance. For HPCN to become economically relevant the definition has to be limited by adding that the maximum performance has to be realized for practical problem sizes. The TOP500 list of the most powerful computers installed in the world [Dongarra 1996] learns that the above given definition of HPCN excludes clusters of workstations, at least currently. However so far industry in general prefers the use of clusters of available workstations for high performance computing in engineering. The reason is that workstations are available in industry and engineers are used to working with them.

The most powerful computers in the TOP500 list are expensive and are mostly installed in universities and in technological- and scientific institutes. The accessibility of the most powerful computers technically, organisationally and psychologically is limited for most industries. This is an undesirable situation because several enterprise operations in engineering require high performance computing. The performance of the most powerful computers for many applications is much larger than that of workstation clusters and the price/performance is comparable [Saini 1996].

The Netherlands has four computers on the TOP500 list that are available for general use. These systems all are connected to the 34Mb/s Netherlands scientific network. The systems are given in figure 1 together with the rating in Dongarra's list of november 1996. In the figure also two systems are presented that are installed in a private company (Shell). In figure 1 the most powerful systems are a shared memory vector computer (NEC SX-4) and a distributed memory scalar computer (CRAY T3E). These computers that are available for general use are installed in the aerospace technological institute NLR and in the Technical University Delft respectively.

---

[1]National Aerospace Laboratory NLR (chapter 1,2,3 and editor); E-Mail loeve@nlr.nl

[2]Centre for Mathematics and Computer Science CWI (chapter 4); E-Mail janv@cwi.nl

[3]Siemens (chapter 5); E-Mail ewoud.snijdoodt@gw.siemens.nl

[4]National Aerospace Laboratory NLR (chapter 5); E-Mail tendam@nlr.nl

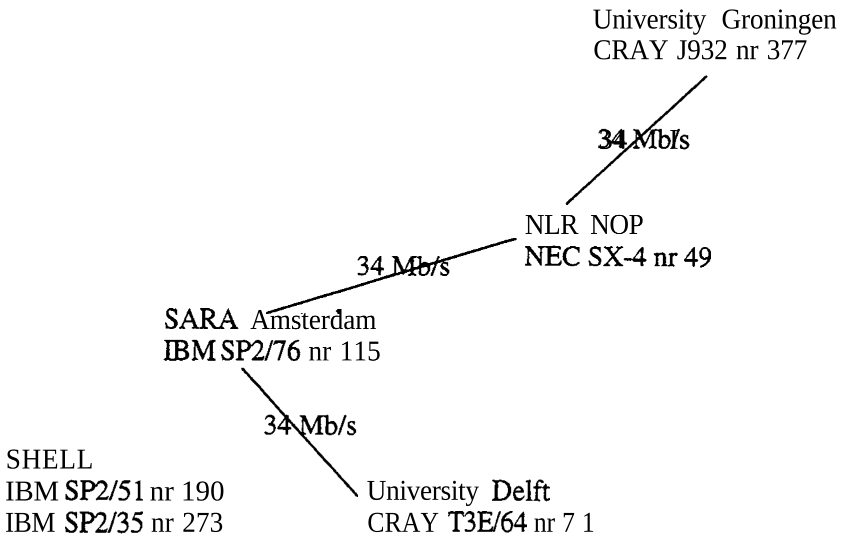

Figure 1 Computers in The Netherlands on the TOP500 list, November 1996

In the European Union it is clear that stimulation of application of classical HPCN in industry especially in Small and Medium Enterprises (SMEs), is very difficult. As a result application of clusters of workstations is stimulated instead in the HPCN programs. In The Netherlands, part of the HPCN program is devoted to classical HPCN. For the present presentation it was considered interesting to describe the relevance of classical HPCN for application in engineering and for support of policy making by government with respect to environmental pollution in particular in the atmosphere and in surface water. This is the topic of chapter 2 of this presentation. The relevance of HPCN for many activities justifies lowering the various barriers for application of HPCN. How this is done in The Netherlands HPCN program is described in chapter 3.

2 THE NEED TO LOWER THE BARRIERS FOR APPLICATION OF CLASSICAL HPCN

HPCN in industry is applied in support of engineering by simulation. Simulation is based on mathematical models of physics such as of flows and materials. The models are solved by means of numerical mathematics. In all modelling steps errors are introduced that mostly can not be removed as a result of shortcomings in the knowledge of the physical reality or as a result of limitations in the computer power that make simplifications necessary. Finally software has to be produced. The software has to be constructed in such a way that it is possible to define and execute verification tests to enable finding and correcting possible software errors. The software also has to be constructed in such a way that errors can be removed with minimum change to introduce new errors. Simulation software in many cases is produced by scientists and engineers with little software engineering know-how and interests. This causes bad testability and maintainability of simulation software. As a result the application of simulation software in industry is costly in many cases. This situation has to be improved by cooperation of simulation experts and suitable informatics experts in development and production of simulation software [Loeve 1992]. Engineering of Computation Fluid Dynamics (CFD) software for testability and maintainability is the topic of [Vogels 1997].

The applicability of the mathematical models and the computations has to be validated by comparison of results with information about the concerning phenomena from other sources. In figure 2 the various sources of errors in simulation are presented in the context of the validation process. Validated software systems for simulation, the supporting documents, results obtained by simulation and documented engineering procedures represent much engineering know-how.

Enterprises are confronted with the challenge to lower cost of product engineering. Lowering costs of engineering nowadays very often means that engineers are only hired for specific jobs. The result is that engineers enter and leave enterprises more frequently than in the past. Proper conservation of engineering know-how for use by successive generations of engineers is essential for the efficiency of enterprise operations. IT can enhance enterprise operation by giving engineers that work in or for the enterprise, easy access to the engineering know-how. This requires central measures for control and management of the know-how [Loeve 1996-2].

Computational Fluid Dynamics (CFD) for simulation of flows in many practical engineering applications is one of the simulations for which always computer performance has been the limiting factor. Advanced industrial CFD software is based on the Reynolds-Averaged Navier-Stokes equations [Rubbert, 1994]. One not too detailed computation for a full aircraft configuration requires about 5 hours in case of a sustained computer performance of 500 MegaFLOP/s. One of the applications of CFD for the near future is support of aerodynamic optimisation [Borland, 19943. Optimization requires more computer performance than analysis, because for optimization for instance about 40 times solution of the flow equations in combination with an optimisation algorithm is required. For NLR, aerodynamic optimisation is the most demanding engineering application [Loeve 1996-1]. Similar computer limitations apply to the field of environmental pollution simulation [Peters 1995].

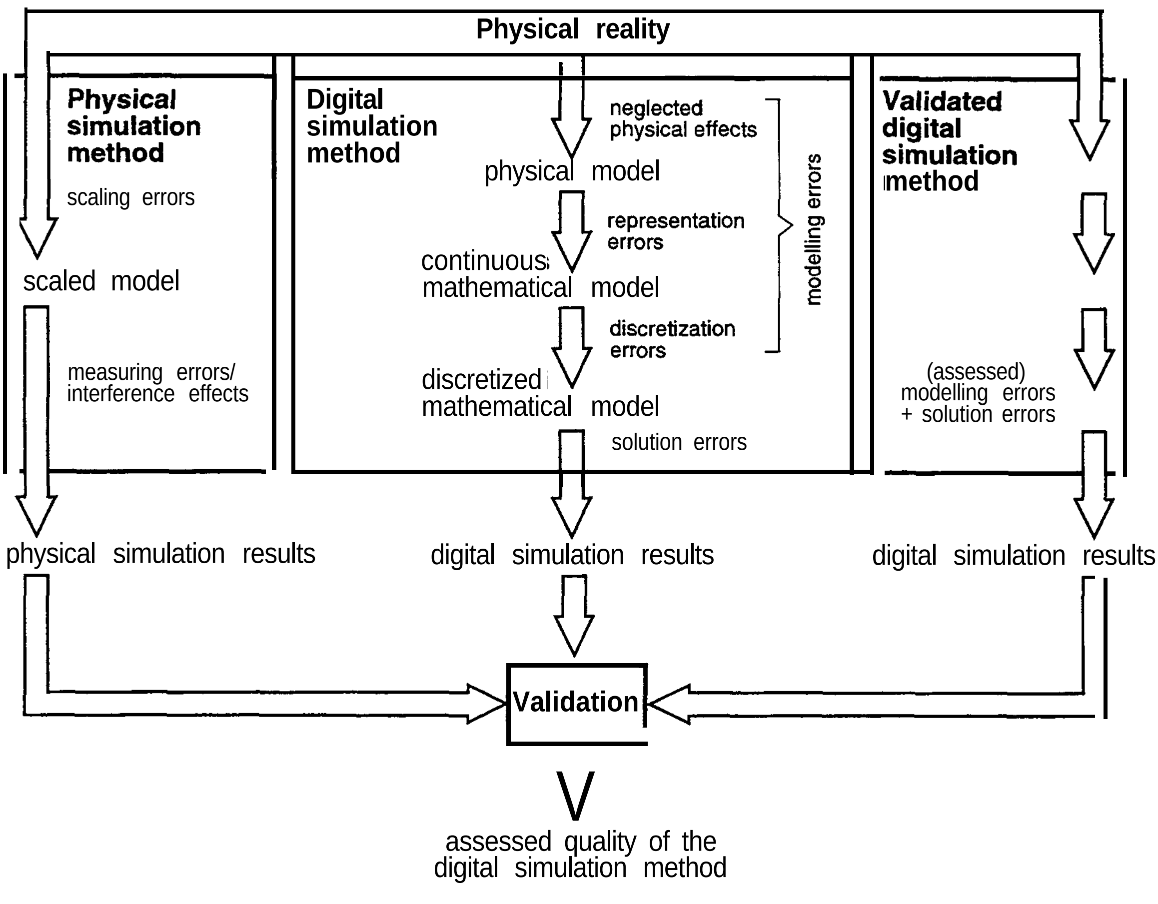

Figure 2 Validation of simulation and the sources of errors

With a NEC SX-4/16 (a vector computer with 16 processors and shared memory) it is easy to realize for CFD in engineering a computational performance on 14 processors of 8,500 MegaFLOP/s [Ven van der, 1997]. With knowledge of only a few techniques on vectorization and coarse grained parallelization, this is realized with relatively little effort (only DO-loop parallelization). One aerodynamic optimization of 40 analysis runs then takes about 10 hours. With some more effort Hatay on NLRs SX-4 realized that 15 processors worked concurrently 96% of the time, each with 1.2 GigaFLOP/s sustained performance yielding overall performance of 18.3 GigaFLOP/s [Hatay 1996]. This range of performance is required for CFD in meteorology, air quality research (see chapter 4) and for computational chemistry. As a consequence it is desirable that the available most powerful HPCN facilities can be accessed for the mentioned applications of simulation. In each country only a few HPCN facilities are installed mostly in universities and institutes. This means that it is required to lower the barriers for application of remote HPCN facilities in combination with easy access to enterprise know-how.

For The Netherlands the question has to be answered if it is worthwhile to make available the installed most powerful computers for practical CFD applications. To answer this question it is possible to make use of information in [Saini 1996]:

- the most powerful servers with scalar processors, the DEC Alpha server (437Mhz) and the SGI Origin2000 (195MHz), are about five times less powerful than one NEC SX-4 processor for a set of benchmarks that in many aspects is representative for flow simulation.
- the performance of one CRAY T3E processor is about equal to the performance of one of the processors mentioned above.
- the price/performance of the three most powerful supercomputers in The Netherlands, the SX with vector processors and the massively parallel CRAY and IBM SP2 with scalar processors, is comparable for the CFD benchmarks.
- the performance of the CRAY T3E with the maximum number of 256 processors in [Saini 1996] is about equal for CFD to the performance of the NEC SX-4 with the maximum number of 32 processors.

It can be added to the above information that for CFD the shared memory NEC SX systems require significantly less implementation and development effort than the scalar massively parallel processors of the CRAY T3E systems. This was the reason for NLR to stick to NEC SX vector computers [Loeve 1996-1]. This notwithstanding the MPP hype. For The Netherlands it therefore may be expected that for CFD, certainly for the near future, the SX-4 vector computer is the best choice if high performance is required. In combination with this for cases that low performance is sufficient, the SGI Origin2000 systems are the best choice, The reason for this is that the price/performance of the SGI Origin2000 is about two times better than that of the NEC, CRAY and IBM supercomputers and of the DEC Alpha Serve?. For the sake of completeness it can be remarked that the price/performance information is from November 1996 and is based on list prices. From experience it is clear that actual prices paid for supercomputers can be more than 50% less than list prices.

Lowering barriers for application of classical HPCN can be done by a combination of distributed and centralised computing, a suitable network strategy and IT tools for management, control and application of the conserved enterprise know-how. This topic is addressed in The Netherlands HPCN project NICE of which the result best can be described with “do-it-yourself flow solutions for engineers and scientists”.

3 THE NICE PROJECT

3.1 Description of NICE

NICE stands for Netherlands Initiative for CFD in Engineering with HPCN. The project covers a wide range of industrially relevant CFD applications in engineering. These comprise waterworks, ships, aircraft, chemistry including flows in ovens, engines and burners. The partners in the project are five institutes and three universities. The institutes are (in alphabetical order): the Centre for Mathematics and Computer Science CWI, the Maritime Research Institute MARIN, the National Aerospace Laboratory NLR, TNO Institute of Applied Physics TNO-TPD and Delft Hydraulics WL. The three universities are, the University Groningen RUG, the Technical University Delft TUD and the University Twente UT. The universities are united in the research school J.M. Burgerscentre. Together the NICE partners have the expertise and the facilities to support a wide range of CFD applications in engineering in industry and especially in SMEs.

---

[5] The price/performance of the SGI Origin2000 workstation is comparable to the price/performance of the Fujitsu vector supercomputer of which the performance is in the same range as that of the NEC vectorcomputer

Engineering processes in many cases are characterised by a large number of activities executed by engineers from various disciplines sometimes located at different geographical locations. The engineers use a variety of simulation software tools. Engineers need workstations for graphical interaction with their tools. Sharing of know-how between specialists has to be safeguarded which means that information servers and intranet servers are required in addition to compute servers and workstations. The network has to give access, for all involved persons and groups, to the computers as well as to the software, data and electronic documents that are stored on the computers. Access to the CFD facilities that are available in The Netherlands is realized in the NICE project in the Centre for High Performance Flow Simulation HFS.

3.2 Centre for High Performance Flow Simulation HFS

The NICE project aims at a virtual centre for flow simulation. This centre makes use of the compute servers that are available in the contributing organisations [Ven 1996]. The compute servers are in each organisation integrated in intranets. The intranets are connected via the scientific network Surfnet in The Netherlands. Many problems occur in practice when engineers try to use a network:

- Remote access to computers is difficult to organize.
- Programs lack uniformity in user interface and data interface. Incidental use is practically impossible because of time required to familiarize.
- Problems with programs on different computers can be:
  - finding out on which computer a specific program may run,
  - exchange of files between computers,
  - incompatible data formats due to different byte ordering and representation of real numbers in different computers.
- Users have to organize, manage, and find a way through more on-line information than that is accessible on a single computer.

Removal of the technical and organizational shortcomings mentioned above is required to realize the target of the NICE project. Know-how information (data, documents and software) shall be made available on-line. Information shall be managed and controlled in such a way that reconstruction of old information is possible and information can be exchanged between engineers in the same way as if they all make use of one computer. Engineers have to be assisted in organizing, re-using and integrating a continuously growing amount of software, data and supporting help information in the network. Only by making high performance computing a natural part of desktop workstation applications, the use of high performance computing in engineering will spread according to the need for competitive enterprise operation.

To avoid the above mentioned problems use is made in the NICE project, of a software facility SPINE that makes it possible to integrate a local workstation with remote compute servers in such a way that the combination presents itself to the user as one single virtual computer. SPINE is a product of NEC and NLR. SPINE is described in [Loeve 1997-2]. SPINE, by being based on standard UNIX and TCP/IP software, has proved to be applicable for realization of environments that span both local-area and wide-area networks.

The working environments constructed using SPINE are highly portable and easy to customize. Moreover, the working environments are open and extendible, in that any UNIX and TCP/IP based computing system can easily be integrated, and that existing software and other forms of electronically available information can easily be integrated. SPINE makes it possible to incorporate security regulations in the user organization.

The development of simulation software is supported by providing via SPINE a software version management tool, and a toolpack sx4dev which integrates several development tools on the NEC SX-4. A framework for the creation of a regression test suite is available within the working environments. For support of development control, electronic forms are used to manage error reports.

For both users and developers project management systems can be set up. Such a project management system consists of an electronic project archive, a search engine on this archive, a drag-and-drop mechanism to submit documents to the archive, and a database for electronic forms.

The above building blocks of the working environment supporting both use and development of CFD software are shown schematically in Figure 3. As is clear from the figure, development and use interact and overlap.

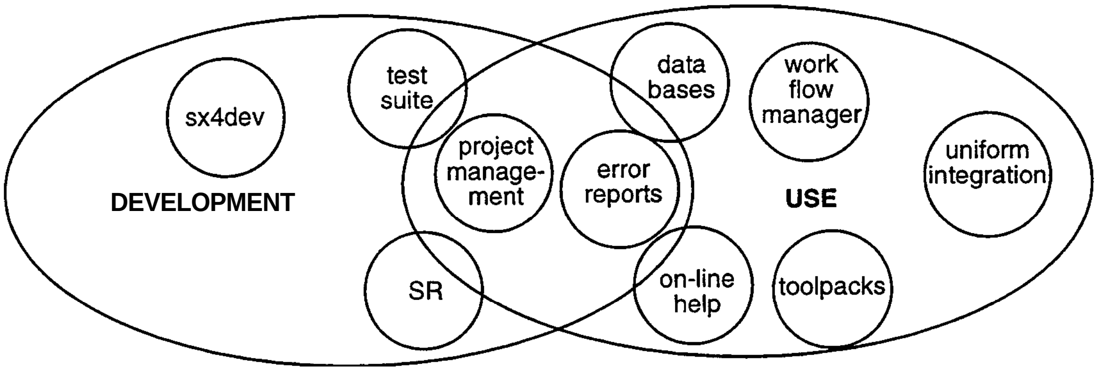

Figure 3 Building blocks for use and development of engineering software.

The development of a working environment is dominated by two principles:

- an integrated tool has added value to the user,
- the working environment shall adapt to the user and not vice versa.

This also implies that integrated tools do not stand alone in the working environment: they are surrounded by tools that either supply input or use the output of the specific tool. This integration of tools greatly facilitates their use.

Since the earliest developments of SPINE it has been applied to ISNaS, a working environment for flow simulations [Vogels 1989]. At first the ISNaS CFD working environment was aimed to support the use of flow simulation packages across the NLR network. The network consists of workstations and terminals that give access to mainframes acting as compute-, intranet-, intemet- and data management servers. A super-computer is integrated in the network since 1987. Because of its success, more functionality was and is added continuously to the network and to the network middleware of SPINE.

ISNaS now supports both use and development of simulation software. The use of the simulation software is supported by integrating the entire pipeline of simulation analysis: geometry modellers, grid generators, flow solvers, and postprocessors. File transfer and remote logins are hidden from the user. Exchange of different file formats between the various tools is made transparent to the user using the SPINE facility for implicit file conversion. Exchange of documents and data is made possible through easy-interface databases. Feedback to the developers of both the working environment and the simulation software is made easy by the use of electronic forms.

The success of the present working environment is probably best exemplified by the following. A CFD trainee was asked to perform an analysis of air heater flow under contract of one of the SMEs in The Netherlands. The trainee had a thorough knowledge of flow physics, and numerical mathematics. But he was less familiar with supercomputers, networks, UNIX and postprocessing. Using the working environment he was able to perform and analyze a specific flow configuration within one week. The analysis consisted of the entire pipeline from preprocessing up to visualization.

In order to facilitate the use of remote resources NICE develops a nationwide simulation centre HFS based on SPINE. HFS consists of two parts: a common part to all NICE partners, and a part specifically tuned to the needs and wishes of a NICE partner. The network on which HFS operates consists of the local network of the partner and the NEC SX-4 at NLR. I-IFS is a specialized 'ISNaS' at each partner, extended with a common set of tools for both project management and development on and use of the NEC SX-4.

Management of the project NICE is facilitated through HFS. A project management system is part of the common tools in I-IFS. The system consists of an electronic project archive, electronic error reports, and search engines to search both the project archive and the error report database. The project archive has a technical report repository. Development on and use of the NEC SX-4 is supported by the toolpack sx4dev that integrates tools for compiling, executing, analysing and debugging source. Automatic makefile generation allows for manipulation of the source code. Options of the tools are based on the experience of experts in supercomputing. Using this tool the computational kernel of a flow solver of TNO-TPD was compiled, executed and analyzed within a quarter of an hour. Before that, TNO-TPD had never worked on an NEC supercomputer.

By supporting both development and use of CFD software, HFS is suited for a wide range of customers. From novice users inexperienced with CFD, through experienced users with a need of more powerful resources, up to developers of new algorithms and software.

An example of a customer needing more powerful resources is ESTEC. For a time-accurate simulation using a multi-grid, multi-block structured solver they required more memory than their IBM SP2 could supply. Moreover, throughput times became prohibitive. Even though the code was originally written for scalar machines, a porting effort of roughly three manweeks increased the megaflop rate from 80 MegaFLOP/s on the IBM SP2 to 500 MegaFLOP/s on one single processor of the NEC SX-4.

NEC stimulates use of SPINE to integrate servers and workstations in organisations and to make remote supercomputers easy accessible for users. The Foundation HPCN stimulates actively technology and know-how transfer between the HPCN projects. The first project for which exchange of information and software facilities with NICE was considered to be profitable is the TASC project that also deals with CFD. Another project for which exchange of information with the NICE project is aimed at, is the SIMULTAAN project concerning distributed training systems. TASC and SIMULTAAN are described in chapter 4 an 5 respectively.

In figure 4 an overview is given of the 10 enterprises in which SPINE is installed for integration of local workstations with the remote SX-4. In some cases, SPINE is just as in NLR also used by other organisations for integration of not only user-workstations with the supercomputer but also with other servers.

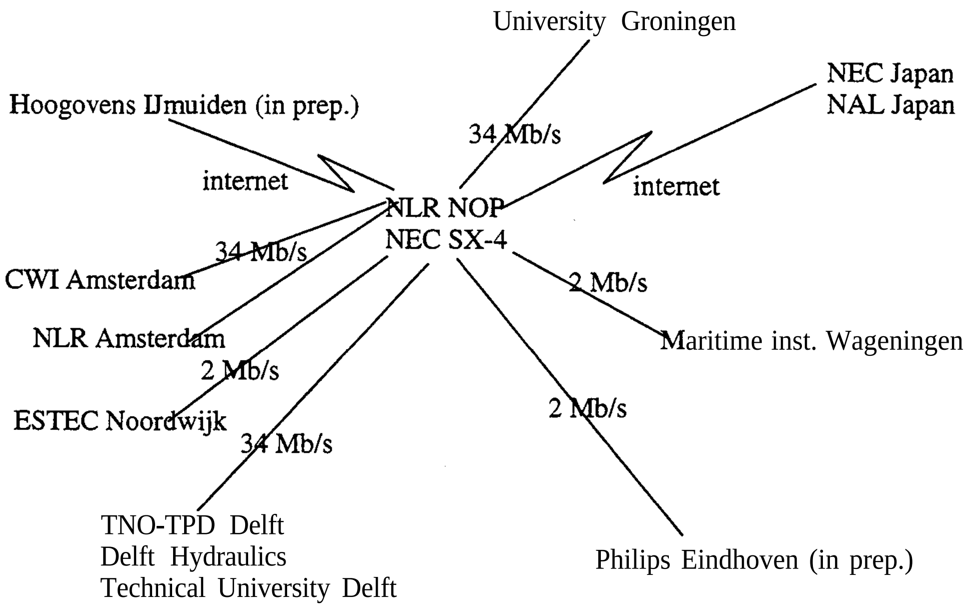

Figure 4 Enterprises in which SPINE is installed april 1997

4 THE TASC PROJECT

4.1 Description of TASC

TASC stands for Transport Applications and Scientific Computing. TASC is a Dutch research consortium involved in the computational modelling of environmental pollution in the atmosphere, surface water or ground water.

Computational modelling of pollution is becoming more and more important in view of the growing awareness of damaging effects which may lead to high economic costs for industry and society. For example, ozone is an hazardous trace gas the levels of which not only depend on natural emissions of different kinds, but also on many anthropogenic emissions (traffic, industry, agriculture). In summer time the ozone concentration can exceed safety standards for humans and animals during short periods (summer smog). Long-term levels can also be too high, however, causing considerable crop damage with high costs for the agro-industry. On the other hand, reduction of anthropogenic emission is in general an expensive procedure too. It is therefore necessary to determine as accurately as possible critical levels and to reduce and control pollution optimally so as to minimize costs. This is a very difficult task requiring complicated optimization algorithms based on comprehensive simulation models.

Four TASC institutes participate in the Dutch HPCN program through the project 'HPCN for Environmental Applications'. This project has started January 1996 and is planned to last 4 years. The institutes are:

- Centre for Mathematics and Computer Science (CWI),
- TNO institute for Environment, Energy and Process Innovation (TNO-MEP),
- Technical University Delft (TUD),
- Delft Hydraulics (WL).

The project focuses on atmospheric air pollution (subject of TNO and CWI) and surface water pollution (subject of WL, TUD and CWI). The primary aims of the project are:

model and algorithm development,
vector and parallel processing,
real life applications.

4.2 **Need for HPCN**

Simulation of long-range transport of pollution is of major importance, as pollution phenomena are almost never limited to areas where dangerous pollutants are emitted. Long-range transport models are based on systems of time-dependent partial differential equations (PDEs), describing advective and subgrid-scale turbulent diffusive transport, chemical and biochemical interactions, emissions, depositions, etc. When many constituents are in the model and the model is 3D, simulation readily becomes highly expensive. To exemplify this for air pollution, we note that in all atmospheric chemistry models the range of time. constants is huge, ranging from milliseconds or shorter (e.g. OH radical) to years (e.g. methane). This means that atmospheric chemistry gives rise to stiff, nonlinear reaction terms which complicates the temporal integration and hence increases CPU time. Each constituent introduces a PDE which is discretized on an Eulerian grid spanning part of, or even the entire globe. Current global models use about 20 constituents, but in [Peters, 1995] it is pointed out that as many as 40 to 100 are necessary for an adequate analysis of perturbations to atmospheric chemistry due to anthropogenic pollution. Because each single PDE already contains many complicated expressions, it should be clear now that CPU times can be enormous. In fact, even with high-speed computers at hand, computer capacity often dictates the grid resolution and the time span that can be covered within feasible ratios of simulation time and CPU time.

Regional and urban atmospheric models are less demanding than global models. However, also in this field excessive computer performance is needed when the aforementioned optimization has to be carried out, like in aerospace design (see chapter 2). For the simulation of transport and chemistry in shallow water and groundwater similar statements can be made.

4.3 **Experiences with HPCN computers**

All TASC partners involved in the HPCN project have access to high-speed computers, including the Cray C98 and the IBM-SP2 at the computing center SARA in Amsterdam, the Cray T3E at TUD and the NEC SX-4 at NLR. Access to the SX-4 has been arranged with support from the HPCN foundation (see Appendix B); access to the other machines with support from NCF (National Computing Facilities Foundation). As yet no practical experience exists with the T3E as this machine has become available only recently.

Similar as for the CFD applications of the NICE project, for the transport-chemistry applications the shared memory machines SX-4 and C98 are more user-friendly than the distributed memory machines. This advantage should not be neglected, since human effort spent on developing and optimizing HPCN codes is increasing with the complexity of the models and can easily form a substantial part of a whole research task, even for those experienced in advanced algorithm design and software development. For HPCN of scientific (and economic) interest is which type of architecture is best in terms of raw speed, now and in the future.

According to the TOP500 list (figure 1), for the current configurations in The Netherlands the most powerful in terms of FLOP rates is the SX-4. In the remainder of this chapter we will briefly argue that for comprehensive transport-chemistry problems, the shared memory, vector/parallel SX-4 will keep the lead in the next years, certainly in the near future.

Vector parallel experience
For a vectorized, global prototype model, [Spee, 1997] reports a very good performance of 500 MegaFLOP/s on one processor of the C98 and promising at expert estimates of parallel speedups (for up to 8 processors). Recall that one processor of the C98 has a theoretical peak performance of about 1 GigaFLOP/s. The experiences in [Spee, 1997] are very well in line with those from [Verwer, 1995] for a related regional prototype model and also with those for a shallow water transport model from [Houwen van der, 1996]. Hence there is enough evidence that on shared memory machines transport-chemistry models do lend themself quite naturally to vectorization and coarse-grained parallelization, something which is confirmed in the literature. Compared to the C98, we expect that on the SX-4, performances roughly can be doubled, since the theoretical peak per processor is about 2 GigaFLOP/s. This is confirmed in [Hatay,1996]. This means that the GigaFLOP/s range, between 1 and 10, say, depending on the number of processors, is within reach. In the separate, forthcoming note [Verwer, 1997], some concrete SX-4 performance results will be reported for a prototype atmospheric transport solver and a prototype shallow-water transport solver.

Massively parallel experience
The solver from [Verwer, 1995] has also been implemented on the Cray T3D, a massively parallel machine with physically distributed logically shared memory, see [Blom, 1996]. The T3D can contain a maximum of 2048 processors, each with a theoretical peak of 150 MegaFLOP/s, so that in the most ideal circumstances 307 GigaFLOP/s can be obtained. However, our experiences with the T3D are less favourable. We measured excellent parallel speedup (almost 100%) on a 64 processor machine, but could obtain only about 10 MegaFLOP/s per processor due to cache restrictions. So this gives even a lower MegaFLOP rate than one processor of the SX-4 and due to the low single processor performance the speedup result may not even be very spectacular. As observed above, as yet no experience exists with the T3E, the successor of the T3D.

For a much more comprehensive model, [Elbem, 1997] reports speedups on an Intel Paragon using up to 128 processors. Using a sophisticated load balancing algorithm, he was able to achieve a speedup of about 90% with 128 processors, which is very good. However, extrapolation of his speedup curves indicates that these drop sharply if much more processors would be used. This drop is related to the key question of massive parallelism: can the speedup be kept at a high level when much more processors will be operational. For an ideal algorithmic-to-communication ratio, i.e. no communication overhead at all, the answer to this question is given by Ahmdahl's law. This law says that the ideal speedup is given by $S = 1/([1-f] + f/p)$, where p is the number of processors and f the parallel algorithmic fraction. For example, this law shows that if $f = 0.99$, i.e. 99% of all the algorithmic work can be carried out in parallel, the actual speedup is bounded from above by 100. So there is no use in having more processors available then.

Chemistry-transport problems are often claimed to be well suited for massively parallel processing, since they can be discretized such that the parallel fraction in the (operator splitting) algorithms is high. However, 99% is already a lot and as yet it is unclear whether for advanced, comprehensive models this level can be achieved. More research is necessary to find out whether the great potential of massive parallelism can be truly effected in transport-chemistry modelling. Crucial for success will be:

- algorithm design: excessively large algorithmic parallel fractions must be realized;
- hardware design: excessively high communication speeds between processors is essential, as well as excessively high single processor performance;
- software design: of major importance for the user community is much easier to use parallel system software and compilers which relieve users from most of the intrincacies of programming parallel computers.

4.4 Acknowledgements

The TASC solvers were developed with support from Cray Research Inc., under Grants 95.02 and 96.03 via the Stichting Nationale Computerfaciliteiten (National Computing Facilities Foundation NCF). For using the

NEC SX-4 at NLR support was obtained from the HPCN program (EZ/ICES) within the framework of the TASC project "HPCN for Environmental Applications".

5 HPCN FOR DISTRIBUTED VEHICLE SIMULATION: SIMULTAAN PROJECT

5.1 Description of SIMULTAAN

The main goal of SIMULTAAN is to strengthen the position of The Netherlands in the national and international simulator market for distributed training systems. The SIMULTAAN partners are
- TNO Physics and Electronics Laboratory;
- National Aerospace Laboratory NLR;
- Siemens Netherlands N.V.;
- Fokker Space B.V.;
- Hydraudyne Systems & Engineering B.V.;
- Delft University of Technology, Faculty of Aerospace Engineering.

Individually, each partner has an extensive knowledge of several types of simulators. Collectively, the partners have the potential to form a competitive position in the world-wide simulation market. To expand this position an intensive co-operation and pooling of research, development and marketing is essential.

The market of training simulators is growing fast. Training by means of sophisticated simulators is not restricted to pilot training anymore. For example trainings simulators for, tanks, submarines and weapon systems have become an essential part of training of personel. In the civil market simulators are applied for training of train and tram drivers, and also the car drivers segment shows great promises.

Four main results of the SIMULTAAN project can be distinguished:
1. The SIMULTAAN Simulator Architecture. This will be a generic framework applicable for a wide range of simulators, including manned mock-ups of vehicles, high-fidelity flight simulators and unmanned simulators.
2. Permanent Intellectual Infrastructure. Working relationships between partners will improve. The partners will be more aware of each other's capabilities and available simulator components. These components will form a common repository of resources.
3. Permanent Network Infrastructure. A permanent network infrastructure will be realised to develop new techniques. This will facilitate test and demonstration of capabilities of different simulators and simulator components.
4. Industrial Application. Technology transfer from research institutes to industry will be realised more effectively by joint development of concrete products.

5.2 Vehicle Simulation in a distributed environment.

A SIMULTAAN simulator consists of (multiple) networked simulators and appropriate scenario management facilities. An important feature of a trainings simulator that makes use of the SIMULTAAN Simulation Facility (SSF) is that all components of the simulator work together independent of the physical location of such a component.

Communication between components is supported for a multi-thread system and through networking (e.g., ethernet, ISDN, SCRAMNET). Components can be turned on and off or replaced by similar components. A standardised functional description for each SIMULTAAN component will be made available.

An important part of a training simulator is a realistic dynamics behaviour of the vehicle. To achieve this goal, usually a detailed mock-up of the vehicle, possibly mounted on a motion platform, is complimented by a complex mathematical model of the vehicle.

Within SIMULTAAN, the Dynamic Model component contains a collection of realistic and validated mathematical models of vehicle and aerospace systems suited for real-time simulation in a distributed simulation environment. High-level descriptions of the features of the models are included. This facilitates selection of the appropriate model.

The mathematical model, implemented in code, must run in real-time to provide motion cues. This means that a trade-off between computation-time and model complexity must be made. In the past this trade-off often implied a reduction of model complexity. Nowadays, behavioural and performance requirements are such that this may not be acceptable. To arrive at realistic simulations, the mathematical model must be structured to facilitate high performance. Several options are available to achieve this goal. First, one can tailor the code such that it is suitable to run on parallel processing machines. This may not always be a feasible option: already accepted models that have been around for quite a while are usually more suitable to run on a single processor. Computation power is also important. NLR has available for this the NEC SX-4, and also several high performance workstations. The network of NLR is presented in [Loeve 1997-1].

Another important part of real-time simulation of vehicle dynamics is non-real-time simulation! In order to validate model software, one needs reference models. These reference models, also implemented in code, can be used to establish acceptable approximations of the true model. Equally important, the available models can be used to perform control engineering prior to the actual trainings simulation. For the non-real-time simulation the SPINE facility mentioned in section 3 will be used.

6 CONCLUDING REMARKS

The Foundation HPCN in The Netherlands makes it possible that CFD facilities and expertise are offered to industry in an integrated way by institutes and universities. The NICE partners have the knowledge to cover a wide range of CFD applications in engineering. The TASC partners cover a wide range of applications of CFD in determining pollution distribution in the atmosphere and in water.

The SIMULTAAN partners have the knowledge and capabilities to design and build training simulation systems. SIMULTAAN achieves a unique clustering of this knowledge and experience, which is necessary to compete in the international simulation business. The innovative SIMULTAAN architecture will give the partners a lead in the development of customised simulation systems. Distributed simulation capabilities are essential in this context. Installation of complex mathematical models of vehicles on appropriate high performance computers is one of the applications.

The exchange of software facilities between the HPCN projects, is advantageous for the development of the distributed environment for CFD applications as well as for vehicle simulation..

7 REFERENCES

Blom, J.G., Kessler, Ch., Verwer, J.G. (1996) An evaluation of the Cray T3D programming paradigms in atmospheric chemistry/transport models, CWI Report NM-R9604 (to appear in the Proceedings of the EUROSIM conference, Delft 1996).

Borland, C.J., Benton, J.R., Frank, P.D., Kao, T.J., Mastro, R.A. and Barthelemy, J.F.M. (1994) Multidisciplinary design optimization of a Commercial Aircraft Wing - An Exploratory Study. 5th Symposium on Multidisciplinary Analysis and Optimization, Panama City Beach, Sept. 7-9, 1994. AIAA-94-4305-CP, pp 505-5 19.

Dongarra, J.J., Meuer, H.W. and Strohmaier, E (1996) TOP500 Supercomputer Sites November 1996; retrievable from intemet by anonymous ftp to www.netlib.org/benchmark/top500.ps.

Elbem, H. (1997) Load balancing of a comprehensive air quality model, In: Making its mark - the use of parallel processors in meteorology, Hoffmann, G., Kreitz, N. (eds.), World Scientific.

Hatay, F.F. (1996) Performance of a computational fluid dynamics code on NEC and CRAY supercomputers: beyond 10 gigaflops; email hatay@nas.nasa.gov

Houwen, P.J. van der, Sommeijer, B.P., Kok, J. (1996) The iterative solution of fully implicit discretizations of three-dimensional transport models. CWI Report NM-R9621.

Loeve, W. (1992) Life-cycle-oriented method for development and production of large-scale industrial mathematics software; in ***Computers in Industry, 18*** Elsevier.

Loeve, W. (1996-1) From parallel CFD to a tool for computer aided engineering; in ***Parallel Computational Fluid Dynamics*** Elsevier 1996.

Loeve, W. (1996-2) Creating capabilities for effective and controllable scientific and engineering computations; RCI European Member Management Symposium IX, June 1996; published in NLR TP 97113.

Loeve, W., Ven van der H., Vogels, M.E.S., Baalbergen, E.H. (1997-1) Network middleware illustrated for enterprise enhanced operation; to be published in ***Proceedings of the 4th International Conference on Computer Integrated Manufacturing ICCIM'97 Singapore.***

Loeve, W., Baalbergen, E.H., Vogels, M.E.S., Ven, van der H. (1997-2) High performance computing in simulation that fulfils user-driven quality criteria; to be published in ***DASIA '97 Conference*** ESA Publications Division.

Peters, L.K., Berkowitz, C.M., Carmichael, G.R., Easter, R.C., Fair-weather, G., Ghan, S.J., Hales, J.M., Leung, L.R., Pennell, W.R., Potra, F.A., Saylor, R.D., Tsang, T.T. (1995) The current state and future direction of Eulerian models in simulating the tropospheric chemistry and transport of trace species: a review, Atmospheric Environment 29, 189-222.

Saini, S., Baily, D.H. (1996) NAS Parallel Benchmark (Version 1.0) Results 11-96; NASA Aimes Research Center, email: saini@nas.nasa.gov.

Spee, E.J., Verwer, J.G., Zeeuw, P.M. de, Blom, J.G. Hundsdorfer, W. (1997) A numerical study for global atmospheric transport-chemistry problems, CWI Report MAS-R9702, 1997.

Rubbert, P.E. (1994) CFD and the Changing World of Airplane Design; AIAA Wright Brothers Lecture, Anaheim California; September 18-23.

Ven van der H. (1996) NICE: do-it-yourself flow solutions for engineers and scientists; International Eurosim Conference HPCN Challenges in telecomp and telecom Delft; published in NLR TP 96424.

Ven, van der H., Vegt, van der J.J.W. (1997) Experience with advanced CFD algorithms on NEC SX-4; to be published in ***VECPAR '96 selected papers*** *from* ***the 2nd International Meeting on Vector and Parallel Processing (Systems and Applications),*** Springer.

Verwer, J.G., Blom, J.G., Hundsdorfer, W. (1996) An implicit-explicit approach for atmospheric transport-chemistry problems, Applied Numerical Mathematics 20, 191 - 209.

Verwer, J.G., Kok, J. (1997) Performances of Two Transport Solvers on the NEC SX-4, CWI Report MAS-R97xx.

Vogels, M.E.S. (1997) Architecture and constructionprinciples for computational fluid dynamics software for engineering in industry; Thesis university of Twente, M.E.S. Vogels NLR Amsterdam.

APPENDIX A

Projects partly financed by the Foundation HPCN

The Foundation HPCN in The Netherlands so far has selected nine projects. These projects are:

- *3DOME*
The purpose of the 3DOME project is the development of a data base with virtual three dimensional objects. 3DOME will become a market place of 3D objects with facilities storing, processing and trading.
Contact person for the project is W. Schuur, E-Mail wim.schuur@rcc.nl

- *ELSIM*
The purpose of this project is to make available software modules for computation in simulations of integrated circuits and printed circuit boards.
Contact person for the project is K. Stam, E-Mail stam@prl.philips.nl

- *HPV*
HPV stands for High Performance Visualisation. The purpose of the project is to develop a generic, user friendly visualisation facility for the application of HPCN.
Contact person for the project is A. Tatters,
E-Mail atatters@inetgate.capvolmac.nl

- *TASC*
TASC stands for Transport Applications and Scientific Computing. The purpose of the TASC project is development of new algorithms and parallel software for simulation of 3D transport of pollution in the atmosphere and in water.
Contact person for the project is J. Verwer, E-Mail janv@cwi.nl

- *NICE*
NICE stands for Netherlands Initiative for Computational Fluid Dynamics with HPCN. The purpose of the NICE project is flow simulation for a wide range of engineering and scientific applications and realization of an integrated center for flow simulation. The center makes use of a nation wide computer network that is presented to the user as one single virtual computer.
Contact person for the project is W. Loeve, E-Mail loeve@nlr.nl

- *IMPACT*
The purpose of the IMPACT project is fast analysing of large quantities of financial data several times per day making use of HPCN technology and presenting data to users.
Contact person for the project is M. Hoevenaars, E-Mail michelh@ing.nl

- *JERA*
The purpose of the JERA project is development of new commercial web-services and development of a distributed fault-tolerant scalable architecture for this.
Contact person for the project is B. Herzberger, E-Mail bob@wins.uva.nl

- *SIMULTAAN*
The purpose of the SIMULTAAN project is development of training simulators for complex control in the sectors transport, defence, industry and space.
Contact person for the project is E. Snijdoodt,
E-Mail ewoud.snijdoodt@gw.siemens.nl

The Foundation HPCN is Technology Transfer Node (TTN) for Preparatory, Support and Transfer (PST) proposals for EU stimulation programs.

Information about the Foundation HPCN is available via the Bureau HPCN, E-Mail bureauhpcn@mxi.nl

APPENDIX B

Contributions to the infrastructure by the Foundation HPCN

The policy of the Foundation HPCN is to contribute to the HPCN infrastructure in The Netherlands if this is required to make advanced HPCN applications possible. The contributions concern so far:

- Use of the SX-4 of NLR for the TASC project.
- Extension of a Parsitec computer for the 3DOME project. This system is operated by the university computer center SARA in Amsterdam.
- Extension of a NCube/Oracle platform for the IMPACT project. The platform is operated by the computer center of the Technical University Delft.
- Contribution to the realization of a 3 x 3 x 3 meter "cave" for 3D stereo-scopic virtual reality projection. The concerning facility CAVE is operated by the university computer center SARA in Amsterdam.
- Use of the SX-4 for another HPCN project as soon as the request is received.