





i-science

**computer
science
research
cluster**



i-science computer science research cluster

production

Hedi Poot, Manoe Mesters

text

This booklet is an abstract of the official programme text written by the I-Science Programme Preparation Committees.

The full text of the I-Science Research Programme can be downloaded at: www.nwo.nl/i-science/programmetext

design

Tegenwind Grafisch Ontwerpburo, Leiden

printing

Artoos, Rijswijk




publication

The Netherlands Organisation for Scientific Research
Research Council for Physical Sciences

The Hague, October 2005

This booklet is intended to provide general information to the public and researchers. No rights can be derived from this text.

contents

	preface	5
	introduction	7
	glance in brief	9
	glance	13
	view in brief	25
	view	29
	stare in brief	37
	stare	41
	organization	47
	abbreviations	51

preface

In the previous years, the Physical Sciences Council of the Netherlands Organization for Scientific Research has worked towards a coherent reinforcement of research in computer science. In the period 2001-2005, their policy was based on the National research agenda for computer science (NOAG-i), which identified seven priority research themes. The **i-science** programmes **glance** and **view** result directly from this research agenda and focus on fundamental research within the NOAG-i themes Parallel and Distributed Computing (PDC) and Modeling, Simulation and Visualization (MSV). The NOAG-i has now been succeeded by a research agenda for the period 2005-2010, the NOAG-ict, in which the importance of the research in **glance** and **view** is still strongly acknowledged, as well as the importance of more cooperation with other science disciplines such as astronomy. The **i-science** programme **stare** is a concrete step in this direction.

The scientific content of the **i-science** programmes was defined by the experts in the field. Their choice for the research questions described in this text was partly determined by their relevance and urgency and partly by the available expertise for solving them in the Netherlands. With its relatively large grants, and the choice for the right research themes, **i-science** will make a significant contribution to the necessary realization of more focus and mass in computer science research in the Netherlands.

Dr. Annejet Meijler

Director NWO Physical Sciences

i-science introduction

The programme cluster **i-science** consists of three research programmes: **glance**, **view** and **stare**. **glance** focuses on large-scale parallel and distributed systems, **view** focuses on generic visualization techniques, and **stare** is a research programme stimulating multidisciplinary research on the borderline between computer science and astronomy. The programme cluster aims for reinforcing the research discipline computer science in the Netherlands. Another aim of **i-science** is to contribute to preservation and improvement of the prominent position of the Netherlands in the international developments.

An eye-catching recent international development in computer science is e-science. *“E-science is about global collaboration in key areas of science, and the next generation of infrastructure that will enable it.”* (John Taylor, Director General of the UK Research Councils). E-science involves many computer science disciplines, such as grid computing, programming, and data and visualization technologies. In the UK and USA large national e-science programmes have recently been deployed. Even though such a programme does not yet exist in the Netherlands, scientists and institutions from the Netherlands already play an important role in e-science developments.

Through the I-science research cluster we make a contribution to e-science, focusing on generic research in the area of large-scale parallel and distributed systems and on visualization. Building on these themes is the programme **stare**, which focuses on challenging Computer Science research triggered by astronomical problems.

This booklet describes the thematic content of the programmes. On the final pages, the organization of the programmes is explained.



glance

glance in brief

In 1969, Leonard Kleinrock, one of the chief scientists of the original ARPA project which started the Internet, wrote:

“As of now, computer networks are still in their infancy, but as they grow up and become more sophisticated, we will probably see the spread of “computer utilities”, which, like present electric and telephone utilities, will service individual homes and offices across the country.”

Despite major advances in both hardware and software, thirty-five years later this may be still as true as it was in 1969, for how far away are we still from computing as a utility? The **glance** programme is designed to foster research of a fundamental nature in large-scale parallel and distributed systems, in order to make another step in this direction. The programme is developed in cooperation with NCF (National Computer Facility).

Parallel and distributed computing deal with the structure and the use of collections of network-connected processors and computers. For parallel computing, such a collection is typically homogeneous, has very fast interconnections, is managed by a single resource manager, and runs applications

with a regular structure - in short, such a collection is a parallel system. For distributed computing, such a collection is typically heterogeneous with both fast and slow (and possibly wired and wireless) interconnections, is managed by the separate resource managers of the constituent parts, and is used for applications with components with different functionalities - in short, such a collection is a distributed system.

Over the last two decades, distributed computer systems have become the rule rather than the exception. With the proliferation of the Internet and the Web, which is one of the prime examples of a distributed system, it is probably very difficult to point at any two computers in the world that are not in some way connected and that cannot in some way communicate. Two more recent developments in distributed systems are grids for compute-intensive and data-intensive scientific and engineering applications, and overlay networks (peer-to-peer systems), which are so far mainly used for file exchange. Relatively homogeneous distributed systems made up of standard components are increasingly used for demanding parallel applications. However, customized parallel machines are still being designed and built to deal with the seemingly insatiable need for high-performance computing in science and engineering. The latter use custom interconnects and sometimes custom processors.

Two important characteristics of distributed systems are autonomy and cooperation. On the one hand, the constituent components of a distributed system should enjoy some level of autonomy in that they manage their own resources, on the other hand they have to cooperate to achieve common goals, to the point that they present themselves to their users and the applications as single entities. Autonomy is taken to its extreme in overlay networks, with all peers cooperating on a voluntary basis. The most important characteristic of parallel systems is efficiency. In particular, with the increasing use of clusters for parallel applications, the *raison d'être* for true parallel systems is to obtain a high performance, which scales with the number of processors used.

This is the challenging territory of the projects in **glance**.



glance global computer science

glance aims at strengthening fundamental research on large-scale parallel and distributed systems in the Netherlands. Tremendous advances in computer hardware (processors, networks, disks) have led to parallel and distributed systems of unprecedented scales, as exemplified by IBM's Blue Gene with projected systems with up to 64K nodes, and on the Web with around a billion computers. It is this scale and the ensuing scalability issues that form the focus of the programme **glance**. In addition, as there is already a substantial application-oriented research programme on grids in the Netherlands (e.g., in the Virtual Laboratories for e-Science programme), research in **glance** is of a fundamental nature. Thus, project proposals in **glance** address fundamental research questions in one or more of the following themes in large-scale parallel and distributed systems:

- 1 Management and analysis
- 2 Data acquisition and data processing
- 3 Autonomous configuration, composition, and coordination

The following pages describe each of these themes with some relevant research questions.

theme 1 management and analysis

Managing large-scale parallel and distributed systems has been recognized as a major problem for a long time, but relatively little fundamental research has been devoted to it. Management includes various aspects, such as complexity, resource management, fault tolerance, and performance analysis, as described in more detail below.

reducing the complexity of management

A major problem in global-scale computing is the tremendous complexity of all the systems combined. Management and behavior analysis of both the system resources and the applications have become challenging tasks. Therefore, it is important to simplify these tasks for application developers and users alike. Self-managing systems are now recognized as a promising approach to reducing the complexity of management, control and analysis of large-scale systems. Self-managing systems can be broadly categorized as self-configuring, self-optimizing (performance), and self-healing (fault tolerance). Self-managing systems have to analyze the available resources and configure themselves to execute accordingly in a dynamic environment, enabling applications to run, optimize their performance, and adapt to changing conditions. For this purpose, they need appropriate mechanisms.

The users of large-scale systems need to be provided with adequate mechanisms and interfaces to analyze and control the behavior of their applications and identify and isolate error conditions. For this purpose, access via commodity devices ranging from PC's to PDA's and mobile phones is likely to become necessary, and so the management framework should be flexible and adaptive to the communication means employed by the user. When there are multiple options, this framework needs to set up the best possible connection to the user, selecting appropriate encryption mechanisms, adjusting the volume of management information that is transmitted (which may include compression), presenting the appropriate control mechanisms to the user, etc. Therefore, in addition to the self-management mechanisms of large-scale systems,

dynamically adjusting the configuration of the control framework is an important second component of self-managing systems.

The software managing large-scale systems is increasingly based on open standards designed to support distributed environments, such as web services in combination with XML. These technologies are fairly new and the development of the appropriate management information models, as well as the realization of the protocols to support the framework described above, will be an important step. Part of the challenge is to manage the complexity of protocols that interoperate and have subtle cross-layer inter-dependencies. Especially in the face of adaptive configurations, the implications of the dependencies need to be explored.

resource management

Large-scale parallel and distributed systems potentially consisting of tens of thousands of hardware components and millions of software components have to be managed by resource managers. If a single resource manager would have to deal with this number of resources, it would be very complex, but if there are many different resource managers, their cooperation becomes complex. For instance, a single scheduler of a parallel machine with tens of thousands of nodes must be very sophisticated to make optimal use of the (expensive) system. The schedulers in grids made up of many separate subsystems may employ sophisticated negotiation schemes for balancing the load. But when some system's capacity is sacrificed, mutual peer-to-peer negotiations may be a very simple paradigm for reaching resource management decisions.

In a system composed of independent administrative domains, resource usage policies induce additional complications. Such policies are generally dynamic in nature, and the actual policy and current system state may not be available in full to entities outside the resource's home organization. Resource metrics (based on self-discovery) that differentiate between policy for specific groups must be derived to enable inter-organizational scheduling - not only for compute resources, but for storage and bandwidth as well.

Research questions may be found in areas such as resource discovery (information dissemination, multilevel peer-to-peer systems), resource selection (negotiations, optimization, matching resource and application characteristics), and resource access (prefetching executables and files, adapting the environment to applications).

fault tolerance

Hardware and software faults will occur frequently in large-scale systems, and both resource managers and applications should be able to deal with this. Several techniques have been studied in the past to discover or handle faults, such as monitoring (discovering faults), replication (of files and computations), and checkpointing (storing the state of a distributed program or recording intermediate results). The scalability of many existing techniques, however, is poor. Also, it is unclear yet whether such techniques can be made transparent to the application (as several parallel systems have tried with limited success) or need to be driven by the application.

performance analysis

Performance is of paramount importance and is one of the prime reasons for adopting parallel or distributed systems in the first place. The true challenge is in the scale of the systems. For instance, much previous research in load sharing in distributed systems was performed with simulations involving only tens of processors. In this theme there is ample room for all three classical methods of performance analysis, viz. measurements, simulations, and mathematical techniques.

Measurements are indispensable for understanding and explaining the behavior of large systems. For instance, our understanding of the actual behavior of the Internet and of P2P systems (self-similar behavior, heavy-tailed distributions, flash-crowd effects, power laws) and of the behavior of parallel applications in wide-area distributed systems has only come about from measurements. Many resource management mechanisms and policies are far too complicated to be amenable to mathematical analysis, but an experimental evaluation may not reveal their fundamental behavior due to

“contaminating” details of actual systems. In such circumstances, simulations are called for. Mathematical techniques have proven useful for analyzing phenomena in large-scale systems. Examples are differential equations (for epidemic protocols), queuing theory and fluid models (for P2P systems), and statistics (for workload modeling).

theme 2 **data acquisition and data processing**

In many areas of science, industry, and business, huge data streams are generated and huge data sets are maintained and processed, which creates large requirements for bandwidth, storage, and processing capacity. Although in many applications these data sets have only a limited number of sources, increasingly, (massively) distributed data-acquisition systems are built for all kinds of scientific experiments and monitoring applications. For example, in the LOw Frequency ARray (LOFAR) astronomy project which is currently under development in the Netherlands, a radio telescope will be built consisting of tens of thousands of sensors spread across an area of hundreds of square kilometers. Because this system produces continuous data streams, there are also real-time constraints, and because of the sheer size of the data streams, a supercomputer is needed to process them.

Important problems in this area are data acquisition, reducing the data to manageable sizes, and distributing, replicating and storing the data.

data acquisition

An important problem in large-scale distributed data acquisition systems which may consist of tens of thousands of sensors is fault tolerance and robustness. In addition to “ordinary” hardware and software failures, an aspect that is unique to data-acquisition systems is that the signal from a sensor may not simply fail in an on/off fashion, but rather can

be corrupted by some noise source (e.g., by radio-frequency interference in the case of LOFAR). Although at some level of corruption it is better to discard the signal from the sensor in question altogether, there is a “grey area” in which the signal is still useful, albeit with a lower weight.

Logging the performance of each sensor in order to correct the results later may be impractical because of storage problems, but more importantly, because the type of processing involved makes later correction impossible. Therefore, data processing and management strategies must not only be robust against the complete failure of components, but must also be able to simultaneously deal with processing the performance-level data. Given the hierarchical nature of several (proposed) systems in this field, the performance data must be correctly transferred through all levels of the hierarchy for optimal quality of the final result. Ideally, such strategies for dealing with the performance data should be transparent to the applications, but it is an open question as to how this should be achieved. Although some forms of noise and other corruption problems might be modeled mathematically, simulation studies are probably needed in more general cases.

Another robustness issue in data acquisition systems is the heterogeneity of sensors in a single system. Although some systems are homogeneous regarding the sensor type and type of signal produced, others (e.g., for the global acquisition of meteorological data) have highly heterogeneous sensor types, which leads to problems of configuring and coordinating the system.

data reduction

In large-scale distributed systems in which large data sets are generated and maintained, there is often a trade-off between processing capacity on the one hand and bandwidth and storage capacity on the other by reducing or compressing the data. For instance, the raw data streams from sensors may first be processed and combined (at or close to the sensors), leading to a reduction of their overall size, before being transmitted to the storage locations. Similarly, multimedia data may be stored in compressed formats, if they can be recreated for replay within

certain time constraints. There are already various efforts in the area of data grids to tackle distributed storage problems, but few of these include time-constrained data acquisition and reduction from very large numbers of data sources. Although the data reduction and data management strategies will be application dependent, the development of more-or-less generic methods for their design and implementation, and studies of their scalability and performance are of great interest.

The hierarchical architecture of many large-scale systems creates challenges in areas of system architecture and bandwidth. Besides, hierarchical algorithms, in which there are parallelization problems that have not been solved yet, are crucial in this respect.

This applies both to computational aspects and I/O aspects: the available communication bandwidth may dictate the type of algorithms that usefully can be employed. Also, at present parallel filesystems are far from standardized and also need further research to meet the large data management requirements sketched above.

data distribution and data analysis topologies

The hierarchical architecture is not necessarily limited to systems in which data acquisition itself is distributed, and decoupling the analysis topology from the acquisition topology gives the option to share the workload over a larger number of organizations. Distributing the raw data, possibly in real-time, to many different locations also ensures that subsequent access to the data does not suffer from bottlenecks. At the same time it leaves the possibility to store far more raw data and thus incrementally improve on data reduction algorithms, or to evaluate multiple reduction algorithms - possibly by different independent organizations. Large scale distribution or replication of data sets reveals the "missing link" between storage technology and high-bandwidth networking. Although both areas have experienced significant performance improvements, successfully storing large data volumes fast is still an open problem. With network bandwidths of 10 Gbit/s and more, storage systems can no longer cope with the data rate, and reliable novel distributed storage techniques are needed.

theme 3 **autonomous configuration, composition, and coordination**

The increasing power of processing and information resources available today is a trend that is likely to intensify in the foreseeable future. The decreasing costs and higher speeds of processors, now with chip-level concurrency, larger memories, and connections with higher bandwidth, provide more raw processing power and interconnection possibilities for global and ubiquitous computing than ever. This increases the significance of concurrency and its relevance in systems on both micro and macro scales. Modern systems will exploit concurrency at many different levels: at the instruction level and data level to exploit the massive on-chip concurrency of future microprocessors, and at the system level through process or component coordination.

Although the forms and the technical details of concurrency at micro and macro levels are very different, they now both share a common context of utilization in modern computing environments: massively concurrent, heterogeneous, distributed, dynamic systems. This gives rise to a number of non-trivial conceptual and technical issues other than what has already been addressed in the classical studies of concurrency. Before demanding new emerging applications in e-science and e-business can put the raw power of modern computing and network facilities to good practical use, challenging issues in the areas of coordination, composition, deployment, and distributed intelligence to support them must be resolved.

coordination models and languages

In contrast to middleware software such as MPI and CORBA, coordination models and languages are meant to offer a systematic means to close the conceptual gap between the cooperation model of an application and the lower-level communication model used in its implementation. The inability to deal with the cooperation model of a concurrent application in an explicit form contributes to the difficulty of developing working concurrent applications that contain large numbers

of active entities with non-trivial cooperation protocols. A complex protocol is typically not recognized as a “commodity” in its own right, because it is only implicit in the behavior of the rest of the concurrent software. This makes maintenance and modification of the cooperation protocols difficult, and their reuse next to impossible. Coordination languages are the linguistic counterpart of the ad hoc platforms which offer middleware support for software composition.

The popular shared data space coordination models offer a metaphor that directly matches only a particular class of applications, where entities share and exchange data through a common, public medium with one another. Catering for other classes of applications often involves implementing non-trivial communication and coordination protocols on top of shared data space models. Furthermore, realizations of such shared data space models, at least in their current implementations, do not scale well over large distributed systems. Challenging and open problems in this area include non-shared-data-space models and languages, especially to accommodate point-to-point and multi-party interactions; coordination by third parties (exogenous coordination); support for mobility and dynamic reconfiguration; models and languages that espouse both synchronous and asynchronous primitives as well as their combinations to construct coordination protocol; and models and implementations that are scalable over large heterogeneous distributed platforms.

composition

Component-based software engineering advocates the construction of software systems through the composition of coordinated autonomous components. Its benefits include software reuse, simpler and faster construction, enhanced reliability, and a dramatic reduction in the complexity of construction of provably correct critical systems, many of which involve real-time concerns. Effective and flexible component composition is still a challenge today, and real-time constraints may make it even more demanding.

In service-oriented computing, a web service can represent a unit of business that an organization exposes to other organizations on the World Wide Web. The recent Web Services standard governs how one defines, advertises, and uses web services. Composition of primitive web services into complex ones is a next challenge. Existing proposals for languages for service composition (also called choreography of web services) often lack foundations in theoretical computer science and possibilities to address composition and coordination from a more general perspective than business process and workflow applications only. Composition of software components or web services to build complex functionality or services is a research area that was recently recognized to be of great importance to the Internet (Grid and Web) research community.

Current models and standards for components and web services incorporate quite inflexible mechanisms for their composition, where each entity imposes severe restrictions on how it can be composed with others. A flexible composition paradigm wherein the same set of components or services can be composed in different ways to yield systems with different emergent behavior constitutes a prominent challenge in this area.

software deployment

Large distributed applications are increasingly constructed as complex systems by integrating heterogeneous collections of software components (programs, databases, web services, agents, etc.), which may have been constructed by different developers. Moreover, such components can change rapidly and independently, and may be added to a system while it remains in operation. As a consequence, it may be difficult to deploy and manage the whole system. Therefore, software deployment becomes a crucial step in the software life cycle, leading to new challenges involving languages, models, and systems for its support and automation. Software deployment involves all activities related to the release, the installation, the activation, the update, the deactivation, and the removal of components, as well as the configuration of the resources to satisfy the performance requirements of the whole system. In order to support deployment, concurrency must become the

primary typing information, which must be quantifiable and must be available dynamically as some kind of meta-data to the entities that aggregate, coordinate, monitor, and reconfigure future distributed applications and systems. For the deployment of such distributed applications, (new versions of) software components and web services must be uploaded to (remote) resources, must be (automatically) installed, and must be configured for their tasks.

distributed intelligence

Along with the advantages of distributed components such as efficient construction of systems out of reusable third-party and off-the-shelf software, come challenges such as coping with changes mandated by the independent evolutions of a system's constituent components and protocols that happen at the discretion of their respective autonomous providers and custodians. Composing distributed components and services, therefore requires advanced automated processing to give semantic meaning to software components before they can be used as building blocks to realize new complex functionality. Distributed software components, web services, and resources are described by, or annotated with, metadata to give a well-defined meaning to each entity. This distributed metadata allows software agents, users, and programs to interoperate, dynamically discover and use resources, and extract knowledge to solve complex problems.

Open systems as envisioned by the Internet community need metadata descriptions of the entities in the system, and need to share a common formalism and rules for inferring new metadata and knowledge through ontologies. Services are needed for reasoning about and querying metadata and ontologies. The challenge is to provide the management functionality for knowledge discovery and applications and systems management; semantic modeling of application tasks and requirements, services, data sources, and computing platforms; and self-configuration and autonomous management where entities manage themselves according to their observations and interactions with other entities. The required scalability and the complexity of systems calls for this management functionality to be autonomous and intelligent.



view

view in brief

In 1965 Ivan Sutherland wrote in a visionary article, entitled "The Ultimate Display":

"One must look at a display screen as a window through which one beholds a virtual world. The challenge to computer graphics is to make the picture in the window look real, sound real and the objects act real."

Forty years later an enormous amount of progress has been made. On each desktop PC realistic animations of scenes with hundreds of thousands of polygons can be shown, thanks to the developments in software and especially hardware.

The driving force for this success is simply that vision is our most important sense. The use of images and animations enables us to communicate with the computer via a highly sophisticated, parallel high-bandwidth connection, it pays off to exploit this route. The range of applications is large. The main categories are:

- Entertainment. Computer games are a mass market, computer generated imagery is used routinely for special effects and complete movies.

- Data visualization. Visualization is indispensable in science and engineering to understand the large amounts of measured and simulated data.
- Training and simulation. Simulators are a cost-effective means for training (for instance pilots), virtual reality simulations help to remedy phobia.
- Design. High quality renderings of new products are generated without making a prototype; the quality of new buildings can be judged with virtual walkthroughs.

Visualization concerns the interactive synthesis of images and animations in order to provide insight and experiences. It might seem that the aim of Sutherland has already been reached. However, there is a need for more fundamental insight.

The **i-science** programme **view** asks for *generic* methods and techniques. New paradigms are requested that integrate and unify so far unrelated ad hoc methods, that solve open problems, and open up new application domains. Solid foundations on theory, probably from disciplines such as mathematics or psychophysics, are needed, which are used to derive insights and algorithms that are fruitful for many applications. Typically, a model based approach is used. A model serves two purposes here. It is a compact description that provides insight into a class of problems, and also it provides a base for the automatic generation of solutions.

view explicitly asks for *effective* methods. Specifically, it has to be shown that the application of new concepts solves real-world problems, and leads to significant improvements in performance. A variety of measures can be employed here, such as simply the speed of image generation and the size of geometric models that can be handled, but also for instance the quality of the generated experience or the ease with which users can complete their tasks. Measuring such aspects is a difficult problem on its own, the development of methodologies for evaluating and comparing the effectiveness of existing and new methods is needed.

view challenges the research community to come up with proposals to shed more light on this.



view visual interactive effective worlds

The aim of **view** is to stimulate research on visualization, including modeling and simulation to support this. The current base in the Netherlands is narrow. Five relatively small research groups in the Netherlands focus explicitly on visualization and/or computer graphics, about fifteen other groups carry out research in related fields. **view** aims at strengthening the visualization community in the Netherlands, starting from the current interests and strengths, and setting out new challenges. These challenges are characterized by the keywords *generic* and *effective*, as main themes *Interactive Data Visualization* and *Interactive Virtual Worlds* are adopted.

Research projects in **view** address one of the two themes: *Interactive Data Visualization* and *Interactive Virtual Worlds*. These are elaborated in more detail in the next sections, where a description is given and a number of challenges are raised. These challenges should be viewed as inspiration, rather than as an exclusive list.

The type of problems addressed differs per theme, but there is also a considerable overlap. As expressed in their titles, *interaction* is a vital element in both themes, for instance for navigation through space, selection of objects and settings of attributes. There is a strong need for generic models and insight in the effectiveness of solutions for user interaction. *Realism* is also an issue in both, although a more controversial one. What contributes most to realism? How can we measure it? And also, for what applications do we need it, and when is a non-realistic presentation more effective?

view focuses mainly on visual communication, but also audio and haptic feedback can be employed. Research proposals using these modalities are acceptable, provided they address one of the themes and that these modalities are used in combination with visual communication.

theme 1 **interactive data visualization**

Data visualization concerns the use of interactive computer graphics to obtain insight in data. In many occupations professionals are confronted with enormous amounts of data. These data come from a variety of sources. Modern measurement equipment, such as used in medical imaging, satellite imaging, high energy physics, microscopy provide large streams of 2D and 3D time dependent data-sets. The same goes for super-computer simulations, for instance for climate and biochemical research. As an example, a grand-challenge climate modeling project created 8 Terabytes of time dependent 3D data. Furthermore, in a field like genomics large existing distributed databases are used. Standard 2D visualization methods for the interpretation of these data do not scale properly.

The field of data visualization can be subdivided in several ways. A common subdivision is in *scientific visualization* and *information visualization*, where the former is concerned with spatial, continuous data, and the latter with abstract, discrete data. The names of these fields are somewhat misleading, but

they have become standard. Other subdivisions are according to the type of data concerned (e.g. vector field visualization and graph visualization), or the application area (e.g. medical visualization and flow visualization).

The central problem of visualization is *how to visualize large data sets*. The following challenges can be identified.

large data sets

Standard algorithms, such as for instance for iso-surfacing, fall short when really large data sets have to be handled. Both the extraction and also the display of millions of triangles are too time consuming. Much work has already been done in those areas, but more efficient methods are still welcome.

automated analysis and visualization

Statistics and also other fields provide a wide variety of methods and techniques to analyze and reduce large data sets, such as clustering methods and principal component analysis. Most of these methods make assumptions on the underlying data, whereas in visualization it is up to the user to identify patterns and trends. How can we take advantage of such methods, for instance by simultaneous visualization of the raw data and condensed results?

feature based visualization

One solution to the central problem is simply showing only the interesting features. This is easier said than done. For many application domains it is unclear how to define what the features of interest are, and also, how to define in such a precise way that they can be detected automatically. Once those features are found, the next question is how to show them and their attributes such that the viewer can understand them effortlessly.

perception based visualization

The research on the human visual system is vast, yet visualization researchers rarely study or apply what is known about the visual system when designing visualization techniques. There is much to be gained by integration and translation of the knowledge on perception knowledge into the visualization pipeline.

integration of scivis and infovis

Traditionally, *scientific visualization* (scivis) and *information visualization* (infovis) are seen as disjoint but this is artificial. Discrete datasets can be converted into continuous fields, and vice versa. One can expect this to be beneficial, but there is no methodology or model available for this.

automated visualization

The production of an effective visualization requires often many choices and settings of parameters. For an optimal result, manual tuning gives the best results, but for efficiency reasons models and algorithms for automatic settings would be preferable. Also from a more theoretical point of view this is a challenge, since it requires a precise answer to the question what a good visualization is.

multifield data

Much effort has been spent in visualizing scalar and vector volumetric data. However, nowadays measurement equipment often produce data per voxel in different wavelengths, in chemical simulations the concentrations of many different substances are calculated. And even worse, in radiation simulations per voxel distributions of intensities are generated. It is fully unclear how to get insight in such data sets, except by selection and visualization of these components one by one.

multi-modal visualization

Virtual reality techniques have made it possible to experience abstract information spaces as 3D objects, sounds and forces in virtual worlds. The use of multiple modalities, as well as the use of another modality than the obvious one for a certain phenomenon can lead to more effective visualizations. Psychology can give insights into the characteristics and interaction between human sensory modalities and on how to adapt presentations optimally to our cognitive skills.

a theory of visualization

The field of Visualization consists of a large number of separate ideas, methods, and techniques. Can overall theories be defined? And do these lead to useful results?

theme 2 **interactive virtual worlds**

The simplest form of computer graphics is to make a picture of a single, static object. Interactive Virtual Worlds can be considered as the extreme opposite. The users are immersed in a complex scene with a large number of dynamic objects, which can be explored and interacted with. To achieve full immersion, *virtual reality* techniques such as head mounted displays and head tracking have to be employed, but, as illustrated by the popularity of computer games, also on desktop systems virtual worlds can be highly attractive and compelling. Other applications are design, scientific visualization, training, and virtual communities.

Traditionally, much effort was spent to meet the high frame rates required for immersion and the minimization of lag. For large scenes, these are still relevant issues, but there are many other open issues as well. The following gives a number of examples.

manipulating a virtual world

In virtual worlds users should be enabled to manipulate objects. For example they should be enabled to open doors, pickup objects, deform them, etc. Such manipulation should look realistic and the system should assist the user in these tasks, for example by automatically computing the correct way of holding an object and the corresponding internal motions of the avatar executing the task.

effective navigation in virtual worlds

When inspecting a virtual world the user must navigate through it. Current practice leaves this under full control of the user, using e.g. a joystick or the keyboard. Such navigation is difficult, leads to ugly motions and takes the attention of the user away from the tasks to be performed. New navigation methods are required to assist the user in moving through the environment with minimal effort. This also involves tools to inform the user about where he is and how to get to other places of interest.

level of detail in interaction

Humans, when interacting with each other and with their surroundings, adjust their interaction according to the availability of communication channels and manipulation tools, the characteristics of their interlocutor and the overall goal of the interaction. Present-day virtual reality systems do not allow for this flexibility. It is a challenge to provide a general model of interaction, where communication functions can be realized by different, computer-related and/or human-specific means, allowing different levels of detail of communication.

individual and adaptive ECAs

The individualism of embodied conversational agents (ECAs) has become a hot topic. Particular demonstrators have been around to show individualism in motion characteristics, or verbal or non-verbal style. However, it is still a challenge to develop synthetic characters that are individual along all relevant dimensions (outer look, gesturing, speech, dialog strategy,...) in a consistent way, and which adapt to the characteristics of the human interlocutor.

affective appraisal of virtual worlds

Virtual worlds are often used for inspecting real-life locations without the need of being present. E.g. historical sites, models of new buildings, plans for changing landscapes. It is important that the virtual model gives a similar experience as the real location. The affective appraisal of virtual worlds is heavily influenced by the modeling, lighting, sound, and behavior in the world. The challenge is to create the virtual worlds in such a way that the affective appraisal matches that in the real world.

shared virtual worlds

Given multimodal records of a world from many different perspectives, how can we synthesize and integrate these into a global virtual model to be shared in collaborative settings, spanning both time and space? Application domains vary from the virtualization of archived personal experiences to the effective operation of remote crisis management teams.

realistic motion in virtual worlds

Now that virtual worlds as used, e.g., in games have become visually very realistic, users also expect the entities in these worlds to behave in realistic ways. One basic operation for entities is that they have to move in a natural way between places. This involves questions like: what is a natural route through an environment, how to incorporate constraints (e.g. of cars), how to combine internal body motion (animation) with external navigation, and how groups of entities behave (e.g. crowd modeling).

real-time physical simulation


In virtual worlds physical processes, like falling and bouncing blocks, floating objects in water, car behavior, etc. need to be modeled and computed in real time. Standard physical simulation techniques are too slow for this. The goal though is not accurate behavior but realistically looking behavior (compare the lighting models used in graphics that are not based on the actual physical phenomena). This requires new ways of modeling such processes and efficient algorithms to compute them.

modeling worlds

The definition of accurate and detailed geometric models of real-world scenes and objects remains tedious and costly work, despite all methods that have been developed and tools that are available. Directly sampling the real world, via laser scanners or by image based techniques, is an attractive alternative, but still requires often much manual effort and is by not suitable for imaginary worlds.

rendering

Realistic rendering of large models, with different materials under various lighting conditions is a classic topic in computer graphics. Virtual worlds pose two extra challenges. A real-time performance is required, and also, dynamic scenes are often more critical to render realistically, because short-cuts and simplifications that are fine for still images become apparent when viewed dynamically.



stare


stare in brief

"I have read all the writings of Aristotle several times from beginning to end and I assure you that I have not found anything in them which could be what you are telling me. Go, my son, and calm yourself. I assure you that what you took to be spots on the sun are only flaws in your glasses or in your eyes." Fr. Baseaus (to Christoph Scheiner who had seen spots on the sun in 1611)

One of the biggest challenges in science these days is to let scientific progress keep up with the ever increasing data streams. Many current (international) astronomical projects in which Dutch researchers are involved call for advanced computer science methods in order to handle the enormous amounts of data. Examples are ALLEGRO (ALMA Local Expertise Group, ALMA is Atacama Large Millimeter Array), OmegaCEN (supporting data center for the OmegaCAM wide-field camera), JIVE (Joint Institute for VLBI in Europe), LOFAR (LOW Frequency ARray) and 3D-radiation-hydrodynamics.

The amount of data collected is doubled each year, due to ongoing advances in instrumentation. The raw datasets are now reaching sizes in the order of petabytes. Advanced computer science tools and techniques are required in order to transfer, manage, analyze and visualize this information. To enhance rapid developments, a close cooperation between astronomers and computer scientists is required. The interaction between astronomy and computer science can bring both science disciplines to a higher level. While astronomy can benefit from computer science techniques to get the maximum out of the data, the computer science discipline is challenged by the involved complicated questions from astronomy. The programme **stare** is designed in order to stimulate this interaction.

stare aims at innovative computer science that can enhance astronomy, meaning the research in **stare** will not only contribute to the foundation for e-science, but will actually *be* e-science.



stare star e-science

stare aims at stimulating innovative research at the interface of computer science and astronomy. It does so by facilitating cooperation between computer scientists and astronomers. Research projects are carried out by teams consisting of (at least) one computer scientist and (at least) one astronomer. The programme is explicitly meant to support high quality research with surplus value for *both* research disciplines. This relevance for both disciplines is the essence of the programme, and is a main criterion for judging the programmatic fit of proposals with the programme. Computer science should not be confused with computational science, which is not a focus of this programme. Below some relevant research themes are given.

theme 1 technology development

data handling

The first step after measuring astronomical data is to reduce the usually broad bandwidth data to manageable sizes and formats. This involves for instance calibration, filtering, noise reduction, quality control, the selection between which raw data to store and which to discard or the specific data handling in case of transient phenomena. There is a need to develop techniques for automated parallelization.

data storage

Observational astronomy yields huge data streams which have to be stored in databases for longer periods of time after filtering and calibration. However, storage alone is just the first step. Model- and proximity-based query processing techniques, inspired by the astronomer's vocabulary, should be designed. There is a strong need for the further development of alternative query languages (e.g. array-comprehension-, and functional-based). Another major challenge for both computer science and astronomy is data lineage. Data lineage deals with the problem to properly describe, monitor and manage the derivation process of data. In principle, it should be possible to partly re-calculate derived information if one of the earlier steps in the data lineage path has been changed. More examples are the consistency of distributed datasets, connecting different catalogues by associated entries, database speed, the development and application of smart databases, and the developments of standards for astronomical data.

distributed computing and grids

Astronomy is one of the most important applications to make use of innovative distributed computing and grid technology. This is illustrated by developments such as the Virtual Observatory and distributed telescopes such as JIVE and (in the future) SKA. Challenges lie for example in the area of scalability; How can we set up a networking software/hardware infrastructure that can that efficiently transport the tera-, (or peta)bytes of (streaming) astronomical data. How do we scale existing grid technology to this kind of data? How do

we guarantee Quality of Service for these kind of operations? How does the VO connect to the grid? How does the end user get access by means of VO-like interfaces to processing power to solve his observational-research problem? How do astronomers based in one University connect to and make efficient use of scalable resources at other nodes? Other related questions are whether we can simulate and test large distributed systems such as LOFAR on its full scale using grid technology. Can we make sure the system is robust enough to deal with, for example, failing sensors? How do we correlate data from distributed data streams? And how to solve the I/O problem, which is still the bottleneck when large data streams are transported? Web services can for example be useful for making visualizations of different databases and combining them.

theme 2 **interactive data analysis and interpretation**

The aim of measuring, transporting and storing astronomical data is to enable researchers to interpret what has been measured. Visualization, data mining and the virtual observatory enable us to access and understand the data.

visualization

An important step in the data processing sequence is visualization. The main concern for the visualization of astronomical data is the development of automatic methods that can handle the enormous amounts of data and select the relevant parts of the data to be visualized. At the same time, these automated methods have to leave room for human interaction, since the analytical power of the human brain is still indispensable for a correct interpretation of the data. Challenges lie in the area of automated feature extraction (from noisy data) and pattern recognition, visualizing 3D data cubes or high-dimensional parameter spaces, efficiency with respect to speed and memory of existing algorithms for very

large data sets, integrating filtering and visualization in a problem solving environment, the tension between automation and human interaction, visualization of databases, and the intensive interaction between the database management system and the visualization tool. Also, software parallelization and graphics hardware acceleration are important issues to consider.

data mining


In order to optimally exploit the huge amounts of data of different frequencies and modalities that are stored in different formats, at distributed locations, very advanced data mining techniques are called for. How can we trace what's relevant from these huge amounts of (visualizable) data? It should be possible to search through different databases efficiently to find possible relations along several dimensions. An intuition based model is often the starting point in astronomy and demands fast and flexible access to experimental data for validation of the model. Large-scale data mining techniques are essential here. A related issue is data mining in numerical models. If something has been calculated before, there is no need to calculate it again, but this is only useful when scientists can find what has been calculated and have access to the results. The storage, query processing and data mining techniques inherently call for a parallel and distributed solution. Both at the level of system software and data exchange, e.g. P2P database technology, and in methods to automate integration of heterogeneous sources. It will become commonplace that information from different sources should seamlessly be combined to reach new scientific insights.

virtual observatory

In the Virtual Observatory (VO), all astronomy data and literature will be organized into a coherent system that can be accessed by anyone, in any form, from anywhere². The VO consists of a collection of data centers, each with unique collections of astronomical data, software systems and processing capabilities. Storage, services and processing are connected in the Virtual Observatory. Through a VO, scientists can get access to computing power and storage capabilities

that exceed by far the possibilities at their local institution. Furthermore, by combining data from different surveys (e.g. with spectra at different frequencies) a deeper and better insight into the observations can be reached. Processing grid technology will in the future be an essential backbone for the virtual observatory. At the moment, the European partners gathered in the EURO-VO project are building the infrastructure for connecting archives of data-centers, allowing astronomers to publish data on VO, which in turn can be searched and datamined by the community using services that are also provided by the data providers. Obviously, to do this for all instruments and archives is an immense task. Challenges for VO-like **stare** projects can be found in new innovative combinations of dedicated web-services, distributed processing, distributed archiving and federated databasing (essentially any of the above technology and interactive data analysis issues) which will help to build the VO and make it useful and accessible for astronomical research with the new generation of instruments.

The legacy problem (can we still use old software and systems for new instrumentation and developments?) is relevant for each of the previous themes.



i-science organization

glance and **view** have their own programme committee. The programme committee is responsible for the programme and its scientific coordination, and consists of experts from different research organizations and universities. The NWO Physical Sciences Council provides the secretary of the programme committee. Tasks of the programme committee are, among others, determining the assessment criteria, reporting on the progress of the programme and stimulating the knowledge transfer.

The proposals in **view** and **glance** are assessed for their scientific quality by an international assessment committee. The proposals in **stare** are assessed by the assessment committees for computer science and astronomy in the Open Competition. The GBE makes the final decision on the allocation of the grants.

For the members of the committees see our website: www.nwo.nl/i-science.

The budget of **i-science** is M€ 10,5, of which M€ 4,5 is allocated to **view**, M€ 4 is allocated to **glance**, and M€ 2 is allocated to **stare**. By this budget, a total of 21 projects can be subsidized through competition rounds organized in 2005 and 2006. A small part of the budget is reserved for programme activities such as the organization of symposia, workshops and other activities stimulating knowledge transfer. This part is administered by the NWO Physical Sciences Council.

i-science is directed at strengthening and creating clusters of expertise that are active within the relevant themes, by means of relatively large grants consisting of about 3-4 researchers. Each grant has a size of k€ 500 and can be used for:

- PhD positions;
- Post-doc positions;
- Assistant/associate professors;
- Technical personnel and programmers;
- Instrumentation;
- Travel costs for congress visits.

Applications can be submitted by researchers working at Dutch universities, NWO- and KNAW institutes. A proposal can be submitted by one researcher (the main applicant) representing one research group or a consortium of research groups. For **stare**, an extra requirement is that applications can only be submitted by teams of researchers consisting of at least one astronomer and one computer scientist.

See the call for proposals for details.

More information on **i-science** can be acquired from dr. ir. Hedi Poot (070-3440915; poot@nwo.nl) at the NWO Physical Sciences Council or on our website, www.nwo.nl/i-science.

abbreviations

BRICKS	Basic Research in Informatics for Creating the Knowledge Society
Bsik	Besluit subsidies investeringen kennisinfrastructuur (Decision on investments in knowledge infrastructure)
CWI	Centrum voor Wiskunde en Informatica (National research institute for Mathematics and Computer Science)
FOCUS	reinforcing computer science
GBE	Physical Sciences Council (Gebiedsbestuur Exacte Wetenschappen)
GLANCE	Global computer science
LOFAR	LOw Frequency ARray
MSV	Modeling, Simulation and Visualization
NCF	Nationale Computer Faciliteit (National Computer Facility)
NOAG-i	Nationale Onderzoeksagenda informatica (National Research agenda for computer science)
NOAG-ict	Nationale Onderzoeksagenda informatie- en communicatie technologie (National Research agenda for information and communication technology)
NWO	Nederlandse Organisatie voor Wetenschappelijk Onderzoek (Netherlands Organisation for Scientific Research)
PDC	Parallel and Distributed Computing
STARE	Star e-science
VIEW	Visual interactive effective worlds
VL-e	Virtual Laboratory e-science