



THE COMPUTERWORLD HONORS PROGRAM

CASE STUDY

ORGANIZATION:

Foundation ASTRON

PROJECT NAME:

LOFAR

Summary

LOFAR is a wide-area, adaptive sensor network integrated with a distributed, hierarchical computational network. Its purpose is to develop the technologies for real-time interpretation and on-line forecasting (with associated uncertainties) of complex natural phenomenon in the environment.

Introductory Overview

Understanding the environment has become a necessity for our survival as a species. In recent decades the techniques for sensing and analyzing the world around us have evolved to include remote sensing from space and interpretation of phenomena using sophisticated numerical models. An important next step will be large scale deployment of real-time sensing and analysis capabilities to provide decision makers with real-time, on-line interpretation and prediction of complex, rapidly evolving situations. The LOFAR project provides a research platform for developing techniques for achieving this goal. The project aims specifically to connect some twenty thousand sensors directly to a massive computational facility. The sensors are spread across an area hundreds of kilometers in diameter. The data from the sensors are conditioned locally and then streamed in real time via wide-band glass fibers to a specially developed supercomputer, where they are processed on-the-fly and used to guide ensembles of models running in parallel. From an IT perspective, LOFAR is a wide-area, adaptive, inhomogeneous sensor network integrated with a distributed, hierarchical computational network. In sound-bite terms, LOFAR explores how sensors may be added to the GRID. From an applications perspective, LOFAR can be whatever its sensors make it. Radio frequency antenna sensors are provided for application in radio astronomy, where LOFAR is revolutionizing the design methodology of radio telescopes. But while astronomers look to the heavens, geophysicists have added seismic sensors to the network to develop new techniques for imaging the deep underground. Agricultural scientists have added a diversity of in situ sensors to develop techniques for precision crop management using growth models. And meteorologists and hydrologists are planning to add sensors for high resolution weather forecasting, improved wind energy prediction and the management of water in coastal flood plains. From an innovation perspective, the project aims

LOCATION:

Dwingeloo, Netherlands

YEAR:

2006

STATUS:

Laureate

CATEGORY:

Science

NOMINATING COMPANY:

IBM Corporation



THE COMPUTERWORLD HONORS PROGRAM

CASE STUDY

ORGANIZATION:
Foundation ASTRON

PROJECT NAME:
LOFAR

LOCATION:
Dwingeloo, Netherlands

YEAR:
2006

STATUS:
Laureate

CATEGORY:
Science

NOMINATING COMPANY:
IBM Corporation

both to bring cutting edge technologies to bear on a generic problem and simultaneously to promote cross-fertilization among disciplines, particularly at the level of algorithmic approaches. LOFAR began life in response to a problem in astronomy. Astronomy is remote sensing at its most remote and progress in the field relies heavily on the continuing development and application of new technologies. But the history of machines in general and of research instruments in particular shows that any given technology may be developed only to a certain point, after which longer cost effective and the adoption of another technology is indicated.

So it is with radio telescopes. Current telescope design relies on large mechanical reflector antennas to gather and focus signals from the sky. The relevant technology dates from the 1950's and seems to have reached the limit of its affordability. That is, increased sensitivity is needed to detect and study ever fainter radio emissions to ever greater distances, but such increases can only be achieved by increasing the signal collecting area of the antennas. Progress using this design paradigm therefore requires ever larger (numbers of) antennas, the cost of which scales with collecting area and with the cost of steel, which is not decreasing with time. The cost of significant new capability in radio astronomy using the conventional design paradigm of large mechanical antennas is becoming prohibitive. An alternative is provided by a concept involving large numbers of simple, wide-field antennas (e.g. dipoles) feeding a central signal processing system. In this case the signals from the individual antennas may be combined appropriately to provide the same functionality in software as a parabolic antenna provides in hardware – e.g. focusing and pointing. In addition, many copies of the signals may be made prior to combination, thereby yielding multiple, independently operating software telescopes all sharing the same IT infrastructure and having no moving parts. New high-end information and communications technologies are becoming available to make such a 'shared aperture, multi-telescope' possible in practice. In the process, new functionality may be realized and the observational parameter space available to astronomers is broadened in several directions. The LOFAR project was conceived to develop the first operational radio telescope of this kind. Scientifically, LOFAR-the-telescope opens up a previously unexplored part of the observational parameter space: high angular resolution at frequencies between about 20 MHz and 240 MHz (radio wavelengths of 1.2 to 15 meters). For astronomers these are low frequencies and have led to the name LOFAR: LOw Frequency ARray. It is known that coherent plasma phenomena can produce very strong radio emission in this frequency range, and do so in the outer atmosphere of the Sun and in the magnetosphere of Jupiter. With the two orders of magnitude greater sensitivity over any previous telescope at these frequencies, one expects to discover a wealth of phenomena in the more distant Universe involving this emission mechanism. But one also anticipates discoveries relating to the evolution of black holes in galaxies, to the nature of the first luminous objects to form after the Big Bang creation event, and to the origin of the highest energy cosmic rays (of which even the existence is unexplained). Technologically, LOFAR consists of some fifteen thousand antennas, located in clusters and spread over an area several hundred kilometres in diameter. Each antenna receives signals from the whole sky and generates a data stream of about 2 Giga-bit/sec. Local processing conditions the signals, removing local radio frequency interference and generating up to eight independently pointed measurement beams on the sky. The partly processed data is then transported to a central processor on wide band optical fibers, where the beams are combined interferometrically to generate and calibrate multiple 'telescopes' in software. The computational challenge of calibration is substantial: imaging errors caused by the ionosphere across the entire sky must be measured and corrected every 10 seconds. The final output data is stored briefly before being exported via the academic internet and GRID



THE COMPUTERWORLD HONORS PROGRAM

CASE STUDY

ORGANIZATION:
Foundation ASTRON

PROJECT NAME:
LOFAR

LOCATION:
Dwingeloo, Netherlands

YEAR:
2006

STATUS:
Laureate

CATEGORY:
Science

NOMINATING COMPANY:
IBM Corporation

environment to users scattered around the world. Although the algorithmic problem is quite different, the data storage and computational requirements are very similar to those of the Large Hadron Collider at CERN. Economically, the finance for LOFAR derives in substantial part from regional development programs rather than from the usual research financing system. The project therefore also includes structures for the involvement of the commercial sector and for business creation. A legal limited partnership has been set up that is operated by a separate foundation to ensure that the LOFAR project is able to interact with the commercial sector to mutual benefit. And a concerted effort is being made to have LOFAR assume the role of launching customer for start-up companies in relevant areas of technology. Early in the project it was realized that the high-level system architecture – large scale, wide area, sensor network being coupled to massive computational capability – could be applied to a wide variety of important problems outside of astronomy. While much R&D is taking place around the world on integrating sensors into our environment – in automobiles, household appliances, etc – the coupling of many kinds of sensors in real-time with dynamic, often non-linear, numerical models of physical phenomena is likely to become an increasingly fertile generic technology. One can imagine, for example, by combining data from satellites with in situ measurements, weather may be forecast on sub-kilometre scales, providing services for applications as varied as traffic control, precision agriculture, wind energy management, flood control, outdoor public events etc. The first community to propose to add additional sensors to the LOFAR data transport and computational network was geophysics. By adding geophones and hydrophones under the radio antennas, this community aims to image the underground to depths of 30 km. These sensors are spectacularly sensitive and will be deployed at depths of 50 meters to suppress signals from footsteps overhead and nearby trees blowing in the wind. The added value with respect to existing seismic arrays is twofold: First, the sensors are permanently installed, thereby permitting the temporal monitoring of structural changes during gas and oil extraction (an important application in the north of the Netherlands) on timescales of years. Second, the use of natural seismic signals becomes possible, eliminating in principle the need for manmade explosions and permitting exploration of new algorithms for improving the quality of the imaging. The meteorological research community has added infra-sound sensors (micro-barometers) to the LOFAR network, to image turbulence in the atmosphere and explore long distance acoustic detection of phenomena such as earthquakes and explosions. Agricultural researchers have added weather and spectral sensors to monitor the sensitivity to disease of potato crops during critical growth periods. And plans are being developed for combining sensors in dikes, the near-shore sea and under the ground with water flow models for improving the system of water management in the Netherlands, where this is of course a crucial application. The generic model of computation, common to all these applications, involves streaming data from an array of sensors being processed in real-time and coupled to dynamic, generally non-linear, numerical models running in parallel and guided by the sensor data. The goal in each case is real-time interpretation and future prediction (with associated uncertainties) of one or another complex natural phenomenon. The project has moved through the design and prototyping phase, and is currently (March 2006) in final design review. Procurement activities have begun with a view to beginning operations early in 2007 and achieving full operations in mid-2008. Initial scientific results have been achieved with the prototype system and are noted briefly below.

Benefits

Stakeholders in LOFAR include research communities in the fields of astrophysics, geophysics



THE COMPUTERWORLD HONORS PROGRAM

CASE STUDY

ORGANIZATION:
Foundation ASTRON

PROJECT NAME:
LOFAR

LOCATION:
Dwingeloo, Netherlands

YEAR:
2006

STATUS:
Laureate

CATEGORY:
Science

NOMINATING COMPANY:
IBM Corporation

and agriculture, and both the research and commercial communities in the high-end IT-sector. More than 20 organizations are formal members of the LOFAR consortium and many more participate in other capacities. The benefits to these communities may be summarized as follows.

In astrophysics, dramatic new capability is being created. LOFAR will be a hundred times more sensitive than previous instruments at these frequencies. It will be the first radio telescope having instantaneous all-sky sensitivity and play-back capability, making possible the detection of transient phenomena in a monitoring mode followed by (re)forming of a measurement beam at an earlier time in the direction of the transient, to allow its study in detail. Other astrophysical goals include the first deep survey of the sky at these frequencies, discovery and study of the first luminous objects that formed just after the Big Bang creation event, improved study of the solar wind and the solar magnetic storms that result in damage to satellites and disturbance of radio communications, and tomographic imaging of neutron star magnetospheres.

LOFAR's multi-telescope feature is expected to have exciting sociological consequences for the way astronomers do their research. Currently, a limiting factor to innovation in research is the peer review process to obtain access to very expensive telescope time: every minute of integration time is granted based detailed written proposals that are carefully screened by a review committee. This procedure is believed to maximize the scientific output of current facilities, but it strongly favours conservative research and discriminates against high-risk, high-payoff exploration. We are acutely aware that major discoveries – the first extra-solar planets, neutron stars, the cosmic background radiation etc – were not made using instruments at front-line public facilities, but at private and second tier instruments to which access was easy and abundant to obtain. Because LOFAR is opening new scientific possibilities and providing researchers with eight independent telescopes simultaneously, the limiting factor is expected not to be access time but the ability of researchers to digest the acquired information. In this case, we expect there to be abundant observing time available for high-risk, potentially high-payoff programs, even for young researchers to 'play' with the instrument to explore new approaches to observing. We hope that LOFAR will initiate a transition in astronomy similar to the revolution in computing when mini-computers challenged the reigning main-frame machines.

In geophysics, arrays of geophones and hydrophones are being installed with both algorithmic and applications goals. Algorithmically it should be possible to combine natural seismic signals continuously in such a way that a high quality three dimensional images of the underground are obtained. The LOFAR facility will be the first expressly implemented to allow this approach to seismic imaging in an optimum way. If it succeeds it should be possible for example to study the ground subsidence occurring in the northern Netherlands as a result of natural gas and oil extraction. This subsidence is expected but is occurring ten times faster than predicted, with the result that the Dutch government will have to invest in dikes and additional pumping capacity at an accelerated rate. The hope is that LOFAR seismological research will lead to understanding of the physical effects involved and therefore possibly also to potential counter measures.

In agriculture, the goal is to develop techniques to complement Earth observation from space, weather forecasting and farmer's intuition, by adding arrays of inexpensive sensors (temperature, humidity, spectral sensors) across fields of crops. Data communication is currently by wireless LAN with development in the direction of RFID technologies. Assimilation and fusion of available information with on-line growth models are expected to enable farmers to optimize irrigation, application of pesticides and harvesting in both space and time. Estimates are that



THE COMPUTERWORLD HONORS PROGRAM

CASE STUDY

ORGANIZATION:
Foundation ASTRON

PROJECT NAME:
LOFAR

LOCATION:
Dwingeloo, Netherlands

YEAR:
2006

STATUS:
Laureate

CATEGORY:
Science

NOMINATING COMPANY:
IBM Corporation

cost savings between 10% and 20% can be achieved while minimizing environmentally undesirable effects. While the astrophysical and geophysical applications are well under way, producing usable results even with the prototype hardware and software, the agricultural application of LOFAR has a longer timescale such that there are as yet no benefits to report.

Finally, among our stakeholders are the IT research community and commercial IT-sector. Our successes in this area are to be noted in three specific cases. First, we have worked together with IBM Research to specify and validate a version of the BlueGene/L machine optimized for streaming applications. This was the second machine delivered outside of IBM, for a time was the largest supercomputer in Europe, and we may safely claim that the ensuing publicity helped IBM to attract additional customers. Together we discovered various undesirable 'features' and worked together to fix them. Second, we have been a launching customer for a number of small, regional companies in the sector. For example, the 'health maintenance' software environment for LOFAR has been provided by a local start-up, and enabled the company subsequently to win a contract from the European Space Agency to study a similar system for a possible future manned mission to Mars. We have helped another to become familiar with the SCADA monitoring and control system used for the LHC at CERN, adapting it for LOFAR and developing adaptive capabilities.

These benefits are concrete, but represent only a small subset of what could be achieved with further development and utilization of the LOFAR research infrastructure. We are looking at applications as diverse as optimizing wind-energy farms, on-line access for schools, and using commercial FM-radio stations for passive radar applications.

The Importance of Technology

LOFAR is in fact the application of cutting edge information technology to a suite of important problems. The driving application is radio astronomy, but the generic capabilities being developed will find application in many areas.

The areas in which information technology is vital to the project may be summarized as follows:

1. Streaming architecture. Our total data generation rate is 37 Tera-bit/sec. The IBM BlueGene/L that we have adopted as central processor is the first commercially produced supercomputer having both the internal architecture and I/O capacity necessary for the radio astronomical application. Our system is capable of accepting over 700 Giga-bit/sec, and together with the conventional LINUX cluster in which it is embedded has a total processing capacity of about 43 Tera-FLOP/sec. Our data transport back-bone from the sensor array centre to the central processor (a distance of some 70 km) has a bandwidth of 800 Giga-bit/sec. Our hierarchical, distributed, embedded processing system (which is spread over 77 sites across the Northern Netherlands) achieves a continuous processing capacity of about 116 Tera-ops/sec. We have temporary storage at the central processor of half a Peta-byte, with data being forwarded as soon as practical (but in any case within a day or two) to users at universities in the Netherlands and distributed across Europe. These specifications are ambitious but all subsystems have been validated. The design process has made use of a formal design methodology that allowed final component selection to be made at the last minute; the risk taken with this approach is balanced by the advantage of the availability of the latest, high performance hardware.

2. Algorithmic research. The LOFAR research consortium is exploring new algorithms in a



THE COMPUTERWORLD HONORS PROGRAM

CASE STUDY

ORGANIZATION:
Foundation ASTRON

PROJECT NAME:
LOFAR

LOCATION:
Dwingeloo, Netherlands

YEAR:
2006

STATUS:
Laureate

CATEGORY:
Science

NOMINATING COMPANY:
IBM Corporation

number of areas and studying the demands they place on operating systems and computing architectures. Examples include the following: The concept of streaming databases is being developed and the fusion of inhomogeneous sensor data with ensembles of numerical models is becoming a major research area. Making a basic SCADA monitoring and control system adaptive, and optimizing health maintenance algorithms, are being studied. Several approaches to suppressing radio frequency interference, including advanced algorithms for spatio-spectral nulling, are being tried out. And new imaging algorithms for longitudinal seismic imaging studies are being researched.

3. On-line access. It is clear that our several research communities must be able to operate LOFAR remotely. There are no moving parts, so LOFAR may be considered a massive computational facility having an unusual data flow architecture but otherwise fitting into the general GRID concept. The project is adopting the conventional GRID environment as developed at by the high-energy community for the LHC at CERN. For specifically the LOFAR research communities (excepting perhaps the physicists studying the highest energy cosmic rays) this environment is unfamiliar and we are expecting its adoption will lead to new ways of working in these communities and likely also to some evolution in GRID tools.

Originality

We believe that LOFAR is a unique research infrastructure, both in its multi-disciplinary organization and in its technological realization. We have presented the project to a broad range of audiences and always meet with a similar reaction: The ingredients of the facility are recognizable but the combination is innovative and elicits spontaneous discussion of potential new application areas. For example, most audiences are aware of the streaming video accelerators used by computer games but not of the potential for extending the technology to the processing of wide-area sensor data. Most have heard of military phased array radars or of software radios, but the concept of a giant, shared aperture, software multi-telescope is a revelation. That weather stations and satellites exist and provide input to weather prediction models is well known, but that dense arrays of in situ sensors coupled in real-time to numerical models may lead to forecasts on 100-meter scales that in turn will provide new services and markets ranging from water management to precision agriculture, from traffic control to more efficient wind energy generation, is new.

In the astronomical application LOFAR is definitely unique and will open a largely unexplored spectral window on the distant Universe. No other radio telescope has all-sky sensitivity, no other has designed-in interference suppression, no other has the sensitivity to detect the first objects forming in the early Universe, and no other is sensitive both to electromagnetic and to astroparticle radiation.

Regarding its detailed implementation LOFAR also breaks new ground. Early on we realized the importance of the IBM BlueGene architecture and thanks to an excellent research cooperation with IBM Research were able to take delivery of the first /L version outside the U.S. defense establishment. As the LOFAR central processing engine it is not only one of the most powerful supercomputers in the world but its internal architecture and exceptional I/O bandwidth allow it to perform effectively in the required high-data volume streaming mode.

Success



THE COMPUTERWORLD HONORS PROGRAM

CASE STUDY

ORGANIZATION:
Foundation ASTRON

PROJECT NAME:
LOFAR

LOCATION:
Dwingeloo, Netherlands

YEAR:
2006

STATUS:
Laureate

CATEGORY:
Science

NOMINATING COMPANY:
IBM Corporation

We measure our success in three ways: (i) interest from research communities in participating in the development and exploitation of the LOFAR infrastructure; (ii) research results published in important international journals; and (iii) the influence of the LOFAR 'idea' on policy planning by relevant public and private organizations, when that demonstrably occurs following contacts with the LOFAR team.

(i) Interest. As is evident from the text above there has been considerable interest from several research communities in participating in LOFAR. This interest is more than academic, because major project finance has derived from subsidy programs in the Netherlands in which the participants must pay half the costs themselves. As LOFAR in total will cost nearly € 150M (half of which is contributions by participants) this interest is demonstrated in practice by a very serious willingness to invest, both by universities and commercial companies. Interest restricted to the operations phase has expanded during the last year to include additional scientific institutes in Germany, France, Italy, Poland, Sweden and the UK. Many of these institutes are planning to invest in antenna fields nearby and in fiber connections to transport the data to the BlueGene/L for processing in real-time.

(ii) Research. Because the project is currently in final design review, scientific publications are limited to descriptions of the system, articles on the IT research undertaken during the design phase, and applications results acquired during the prototype testing phase, which is now completed. The main scientific success has been the demonstration that ultra-high energy elementary particles from celestial sources may readily be detected at these radio frequencies as they interact with the Earth's atmosphere, and the subsequent unravelling of the physics of this mechanism (Falcke et al, Nature 435, 313, 2005). Videos with nanosecond time resolution show the direction and intensity of the cosmic ray events, and the technique is now being studied for implementation at the world's largest astroparticle detector array at the international Pierre Auger Observatory in Argentina.

(iii) Influence. We have presented the LOFAR concept of wide area, real-time sensing and interpretation in many forums. We believe we have instrumental in having the concept adopted by the European Union's e-Infrastructure Reflection Group's recommendations on investment priorities for the European Commission in the 2007-2013 period (see for example section 18 of the e-IRG White Paper at URL: <http://www.e-irg.org/publ/2005-Luxembourg-eIRG-whitepaper.pdf>). More locally in the Netherlands, our national Applied Physics Organization (TNO) has prepared a Road Map for sensor technology development that includes LOFAR as a potentially important economic investment area. And finally, the EADS Space company (a member of the European Aerospace and Defense Systems group, makers of Airbus, helicopters, satellites etc) has proposed a major initiative to establish a research infrastructure on the back side of the Moon, and has adopted the LOFAR architecture as the most cost effective possibility for preparatory studies.

We conclude that, even at this early stage of our project, we can claim a satisfactory level of success.

Difficulty

The most important obstacles to success have been two.

First, as do many multi-disciplinary projects, we have experienced considerable difficulty with



THE COMPUTERWORLD HONORS PROGRAM

CASE STUDY

ORGANIZATION:
Foundation ASTRON

PROJECT NAME:
LOFAR

LOCATION:
Dwingeloo, Netherlands

YEAR:
2006

STATUS:
Laureate

CATEGORY:
Science

NOMINATING COMPANY:
IBM Corporation

the peer review process leading up to financing. Our main problem was that the reviewers simply did not believe that a research institute with a mission in astronomy would really be willing and able in practice to execute successfully a complex project involving several research communities. In the end, we were championed from an unexpected quarter. Following an extensive evaluation process our proposal was not recommended for financing. But the final decision was to be made by the Cabinet (Council of Ministers) of the Dutch government. Normally, of course, in such a situation the package of recommended and rejected projects would be agreed in a regular Cabinet meeting with no discussion. In our case, however, Minister Maria van der Hoeven (Education, Culture and Science) personally intervened and convinced her colleagues to take a chance on such an unusual project. Her vision and political courage in that one meeting have made LOFAR possible.

Second, during execution our main difficulty has indeed been sociological. We have had excellent cooperation from the research and commercial IT communities, but our astronomical, geophysical and agricultural research colleagues are having great difficulty cooperating across disciplines and realizing cross-fertilization, even when there is clear potential for mutual benefit. We conclude that it will take several years of joint operations to surmount the barriers of jargon and to build the necessary levels of trust and cooperation.