FRÓÐSKAPARSETUR FØROYA

NVD*Rit*

SERRITGERÐ
Thesis

TØKNIFRÁGREIÐING
Technical Report

UNDIRVÍSINGARTILFAR
Teaching Material

UPPRIT
Notes

High-Power Computer Facilities at the Faculty of Natural Sciences, University of the Faroe Islands

Hans Pauli Joensen; Brandur K. Jacobsen; Khanh Duc Nguyen; Bárður A. Niclasen; Knud Simonsen; Dávur Sørensen

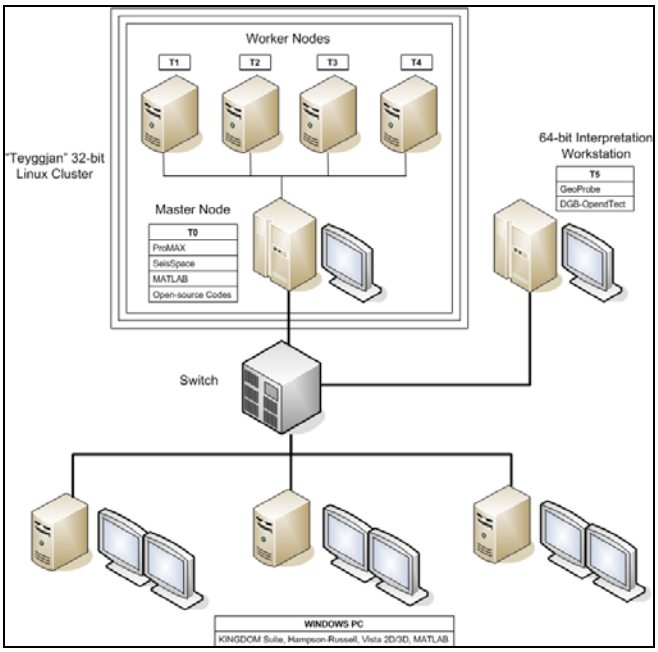

NVD*Rit* 2007:9

NÁTTÚRUVÍSINDADEILDIN
Faculty of Science and Technology

FRÓÐSKAPARSETUR FØROYA
University of the Faroe Islands

Heiti / Title High-Power Computer Facilities at the Faculty of Natural Sciences, University of the Faroe Islands

Høvundar / Authors Hans Pauli Joensen; Brandur K. Jacobsen; Khanh Duc Nguyen; Bárður A. Niclasen; Knud Simonsen; Dávur Sørensen

Ritslag / Report Type *Tøknifrágreiðing/Technical Report*

NVD*Rit* 2007:9

ISSN 1601-9741
Útgevari / Publisher Náttúruvísindadeildin, Fróðskaparsetur Føroya
Bústaður / Address Nóatún 3, FO 100 Tórshavn, Føroyar (Faroe Islands)
Postrúm / P.O. box 2109, FO 165 Argir, Føroyar (Faroe Islands)
✆ • 🖷 • @ +298 352550 • +298 352551 • nvd@setur.fo

Table of contents

Foreword

In natural science research, many disciplines require high-performance computing facilities. This includes research undertaken by the Faculty of Science and Technology (NVD) at the University of the Faroe Islands, mainly within oceanographic modeling and seismic data processing.

Supercomputers don't exist in the Faroe Islands, but NVD has obtained a high-performance computing system through a sponsorship from The Faroes Partnership comprising the oil companies Amarada Hess, British Gas, DONG and Atlantic Petroleum. The sponsorship was part of Enclosure 4 in the license agreements from the first license round in the Faroese territory. The computer system is a Linux cluster configuration with parallel processing facilities. It has the Faroese name Teyggjan.

The report describes the history and motivation behind Teyggjan, and the hardware and software installed on it. Teyggjan fulfills the present need for computer power, but future needs for high-performance computing facilities are discussed as well. The report is organized as presentations from users and administrators.

Hans Pauli Joensen,
Editor

History and motivation behind Teyggjan

Authors: Hans Pauli Joensen and Knud Simonsen; NVD

Teyggjan is the Linux Cluster at the Faculty of Science and Technology (NVD) at the University of the Faroe Islands (UFI). It is a high-performance computer system with parallel computing facilities.

Teyggjan was sponsored in 2003 by The Faroes Partnership comprising the oil companies Amarada Hess, BG, DONG and Atlantic Petroleum. The sponsorship was part of Enclosure 4 in the license agreements from the first license round in the Faroese territory. The original configuration is described below, but the cluster has been upgraded in 2007, as described in the section “Hardware configuration”:

> 4 x Dell PowerEdge 2650: 2 x 2.8 GHz P4 Xeon (Hyperthreading), 2 x 36 GB SCSI, 2MB Ram, 2 x 1G Ethernet, Dual PowerSupply, Redhat 7.3.
>
> 1 x Dell PowerEdge 2600: 2 x 2.8 GHz P4 Xeon (Hyperthreading), 2 x 36 GB SCSI, 6x 146 GB SCSI, 2MB Ram, 2x 1G Ethernet, Dual PowerSupply, Redhat 7.3
>
> 2 x 1G Ethernet switch; Power Switch; Console switch; 42 U Rack, w/ LCD monitor + Keyboard. 2 x 3000VA UPS. Tape station.

Teyggjan was hosted by Føroya Dátusavn until mid-2006, when the computer department at UFI, KT-tænastan, was able to take over and moved Teyggjan to NVD. In this process, Teyggjan became an integral part of the UFI computer network, called Traðarnetið. The new hosting makes data transfer more effective and less costly for the typical user. External users have access to Teyggjan through the Internet.

Teyggjan has introduced a new era in the history of Faroese computer power. In the 1980's and 1990's, NVD conducted some minor numerical simulation projects to investigate possible impacts of planned constructions on the circulation in some Faroese fjords. The simulations were made on local workstations. NVD wanted to encourage numerical modeling of the ocean circulation, but at that time the computational capacity to run state-of-the-art 3D ocean models was not available in the Faroe Islands.

In 1999 a position in oceanography with emphasis on numerical modeling was established at NVD. Through a lateral agreement between the Nansen Environmental and Remote Sensing Center (NERSC) in Bergen, Fiskirannsóknarstovan and UFI, the position at UFI also held an adjoint position at NERSC. This agreement provided access to the supercomputer facilities at the University of Bergen, which included a Cray supercomputer with some 100 nodes.

Detailed numerical ocean models generate relatively large amounts of data, and consequently high requirements on the transmission capacity, which become a severe bottleneck within the cooperation across the Norwegian Sea. Generally the time required to transmit the data from a simulation on the computers in Bergen to the Faroe Islands was longer than the time needed for the simulation itself. In particular in the early phase

of a project, when adjustments and several test runs are required, this situation was not acceptable. The volume of the transmission generated by this activity was comparable or even greater than the need of the entire remaining research community in the Faroes at that time, which in turn also led to financial implications through the cost of the transmission.

In the building up of education programs and research related to the oil industry, the demand for high performance computational facilities was growing in the late 1990's and early 2000's. Analysis of seismic data requires, e.g., large computational capacity.

The experience from the new research activities and inclusion of numerical simulation in the training in several fields led to the conclusion that it would be beneficial to have the super computer facilities locally. This was partly driven by the problems in transmitting large amounts of data to and from centers abroad, but also by a need to have a certain control over the available resources and by the expectation of increased synergy between various research fields through exploitation of a common computational platform.

The primary users of Teyggjan have been scientists at NVD, but other potential users are welcome to use the computing facilities. As of today, the most likely external users will be scientists at Jarðfeingi (the Faroese Geological Survey) and Fiskirannsóknarstovan (the Faroese Fisheries Laboratory).

Hardware configuration

Authors: Dávur Sørensen[1] and Khanh Duc Nguyen[2]; [1]KT-tænastan, [2]NVD

Teyggjan consists of five dual-processor servers: one master node (HEAD) and four computing nodes – these are named t[0-4]; Figs 1 and 2. HEAD is the 't0'. In cluster-speak, this setup is a 5-node cluster, assuming that each physical server, regardless of the number of CPUs, is = one node[1].

The HEAD server is typically to be considered a task dispatcher, which distributes tasks to the computing nodes but on Teyggjan, because of the relatively low number of available nodes, HEAD is also in being used for computing node tasks, though it is not recommended.

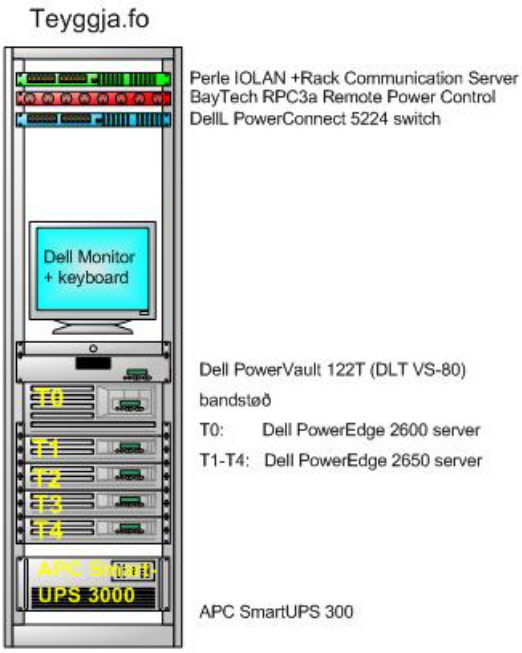

Outlook of Teyggjan.

HEAD is a Dell 2600 2.6-GHz Xeon server. Each of the slaves is a Dell 2650 2.6-GHz Xeon server.

In our upgrade of January 2007, all servers received 4 additional Gigabytes of RAM each, totalling now to 6 GB pr. server; all disks were replaced, going from 10000 rpm[2] to 15000 rpm. The main storage capacity, located on t0, is 146 GB x 6.

The operating system on all servers is RedHat Enterprise Linux 3, update 8; and ScaliManage, the cluster management software, is now in version 5.

There is also a sixth machine, the 't5', attached to the cluster, a powerful workstation which is not to be used for parallel computing but for heavy computing tasks that aren't fit for parallelization. Such tasks are mostly regular serial number-crunching, and especially graphics generation from data sets from parallel runs on t[0-4].

The 't5' is a Xeon dualcore 2.0-GHz 64-bit 'Dell Precision 690 Workstation' with a 512-MB nVidia Quadro FX4500 DUAL graphics card, 8 GB of RAM and a160-GB harddisk. The operating system is RedHat Enterprise Linux 4 (64bit version).

1 I prefer one processor = one node

2 rpm: revolutions per minute

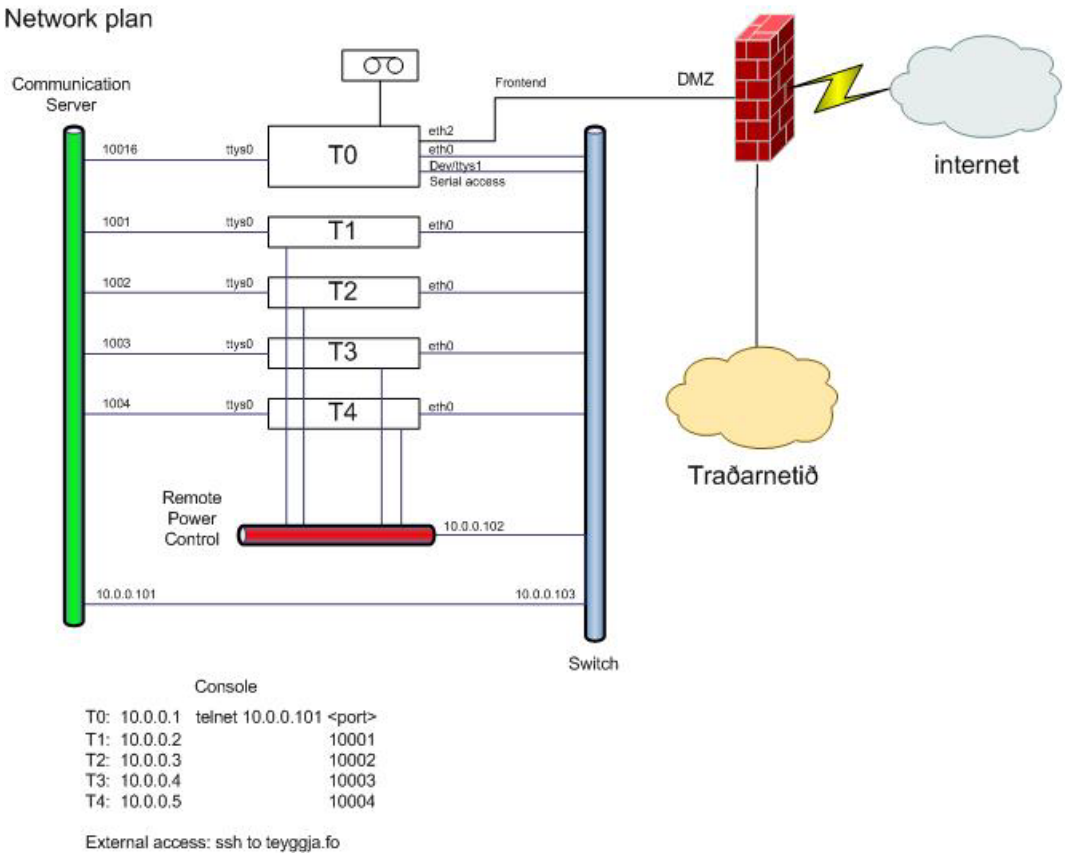

Figure 1 Network configuration of Teyggjan.

How does it work?

Teyggjan is a high-performance, massively parallel computer that perform similarly to a supercomputer for a fraction of the price. It uses a cluster of Linux PCs to form a parallel, virtual supercomputer for parallel computation. Teyggjan currently consists of one master node, and four worker nodes connected together via Ethernet (see Figure 1). Each node is a dual-processor computer with dual-core[3] processors (i.e. computer with essentially four CPUs). Appendix 1 shows the detailed configuration for these nodes.

The master node controls the whole cluster and serve tasks to the worker nodes. It is the cluster's console and gateway to the outside world. The master node is also the storage node, and exports the file systems via network file system (NFS). Worker nodes are configured and controlled by the master node, and do only what they are told to do. In fact, Teyggjan behaves more like a single machine rather than a group of many computers. However, Teyggjan is a distributed-memory machine that differs from a NUMA[4] machine, which also looks like a single large SMP machine and consists of many shared-memory machines connected together.

Because there is no need for worker nodes to access machines outside the cluster, nor for machines out the cluster to access worker nodes directly, the worker nodes (t1-t4) use the private IP addresses: 10.0.0.2/5 (RFC 1918 http://www.alternic.net/rfcs/1900/rfc1918.txt.html).

[3] A "core" is a CPU that may coexist with other CPUs within a single microprocessor
[4] Non Uniform Memory Access

The master node also has this private IP address (10.0.0.1) but uses a second network card to connect to the outside world (IP address: 192.168.150.71 for LAN or hostname: www.teyggja.fo). The most common way of using the system is to access the server's console directly, or either telnet or remote login (ssh, XDMCP or rsh, etc.) to the server node from a personal workstation. Once on the server node, users can start their applications, edit and compile their code, and also spawn jobs on all nodes in the cluster.

Teyggjan, like many other clusters, uses OpenPBS[5] as a job scheduler or queue. OpenPBS (http://www.openpbs.org) was originally developed for NASA in the mid-1990s. Batch scheduling usually means that a script file is created (UNIX shell script), which will be run when the job scheduler has enough resources for the job. The environment indicates which nodes were allocated so that the user may launch programs on them. Commands to the scheduler can either be given on the command line to the *qsub* program, or as comments in the batch file.

In order to achieve the parallelism and take advantage of cluster processing in Teyggjan, we need an application programming interface (API) standard, such as Message Passing Interface (MPI) or Parallel Virtual Machine (PVM). Several MPI implementations exist, both commercial and free. Almost all of the hardware vendors have a MPI implementation available. SGI for instance has a Message Passing Toolkit (MPT). Scali has created a commercially available MPI implementation (ScaMPI). The most used free implementations of MPI include MPICH (http://www-unix.mcs.anl.gov/mpi/mpich/) and LAM (http://www.lam-mpi.org/). ScaMPI and MPICH are installed in Teyggjan, while several of ProMAX's processes are capable of parallel processing using PVM (http://www.csm.ornl.gov/pvm/).

[5] Open Portable Batch System. Teyggjan uses TORQUE which is built upon OpenPBS

Users and the future

Author: Dávur Sørensen; KT-tænastan

Users

The number of registered users is 35 but this number is misleading in terms of usage – quite a few of them have never been active to any noticeable extent.

Active users in the latter half of 2007 are mainly staff and PhD students at NVD, but staff at Jarðfeingi have expressed great interest in getting CPU-time on Teyggjan.

Potential users are also found at Fiskirannsóknarstovan, who have mentioned a few projects where HPCC[6] would come in handy, though no concrete plans have been brought forward yet.

Planning for the future

The current setup has entered the final phase of its optimal life expectancy. Apart from that which was affected by the upgrade of January 2007 (hard disks and RAM), most of the hardware is already outdated. And actually, as a rule of thumb, any computer system passing the age of three years is a candidate for hardware-related interruptions and breakdowns – it is no different for a cluster.

Long-term plans are vitally important. It is necessary to establish the following:

- What is the estimated need for HPCC in the Faroese scientific environment?
- What kind of setup do we want that will meet this estimated need?
- According to the estimated need, is there sufficient technical expertise and funding available to operate, develop and maintain the kind of setup that we want?

Regardless of the outcome of such an analysis, there are some general observations worth taking into consideration already at this early stage.

For starters, talks and mail correspondance with Dell consultant Carsten Hansen, who specialises in HPCC design, indicate that the highly redundant hardware configuration of Teyggjan may most likely not be appropriate as the model for a new design.

Redundancy may rightfully be considered a good thing in terms of ensuring uptime for both the single server and therefore also the cluster as a whole, especially when the cluster consists only of a few nodes. However, it seriously increases the TCA (Total Cost of Acquisition) when using highly redundant servers (i.e. servers that have 2 system disks, 2 power supplies and 2 or more network interfaces).

6 HPCC: high performance computing and communications

For other purposes, KT-Tænastan has purchased two highly redundant Dell 2950 servers within the past year. This is the current generation of the Dell 2650 model used in Teyggjan, and it is known to cost approximately 30,000 DKK. In comparison, we have received recommendations to use non-redundant Dell 1435 servers that cost approximately 18,000 DKK.

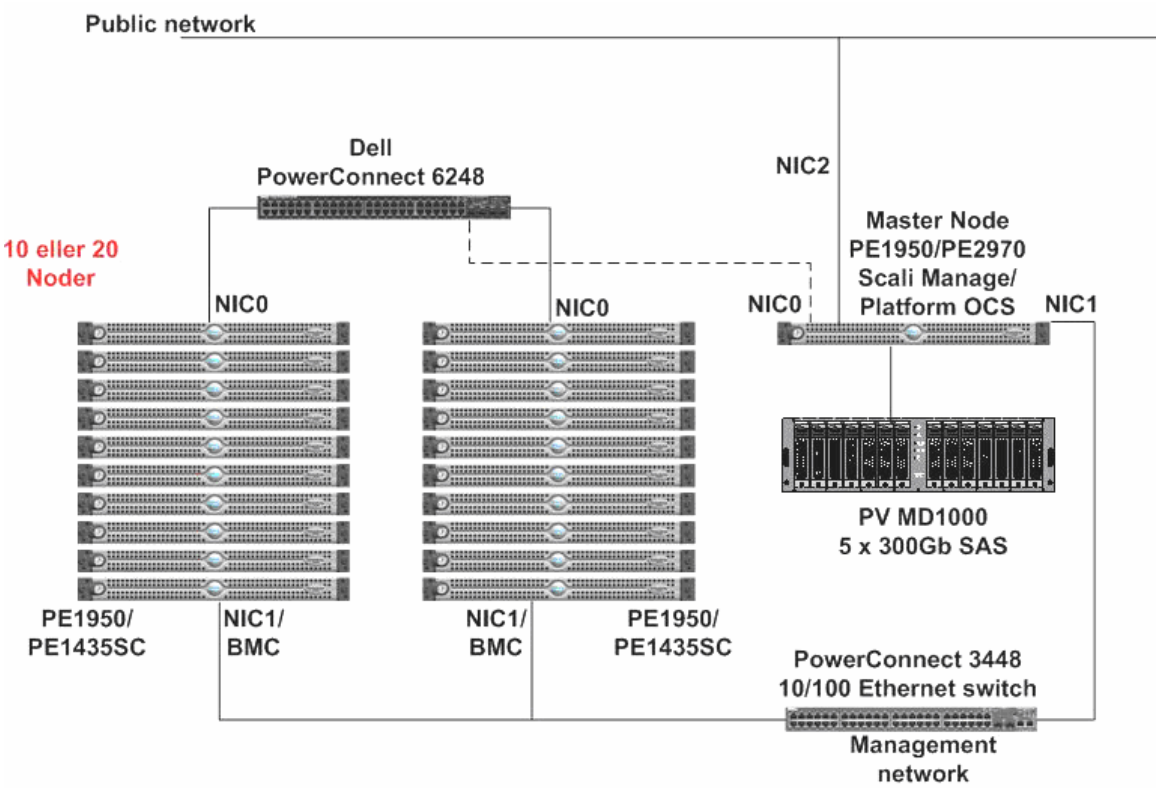

Figure 2. High Performance Computing and Communications (HPCC) cluster.

If HPCC is to continue to be available to the Faroese scientific community, it will be necessary to build a new cluster in the near future. In order to narrow in on the current prices, we asked Dell together with Senta (formerly Farodane) to put together a full design proposal, showing both 10 and 20 servers (Figure 2; and Appendix 3 - in Danish). The costs with 10 servers and 20 servers amount to maximum 196,000 DKK and 358,000 DKK, respectively. The prices do not include installation, which is roughly estimated at 100,000 DDK.

Software on Teyggjan. Current installations and future updates.

Author: Brandur K. Jacobsen; consultant

Teyggjan relies heavily on Open Source and GPL software from volunteers all around the world. Software like the Linux kernel, SSH (Secure Shell), and NFS (Networked File System) are nuts and bolts in Teyggjan. Without these elements and their like, the cost of setting up a computer cluster like Teyggjan would be much higher.

Here is a list of currently installed software in Teyggjan that's available for researchers and developers:

- ProMAX 2003.12 (2D, 3D, 4D): ProMAX uses a customized bundled version of PVM (Parallel Virtual Machine) for some processes.
- Torque 2.1.8: Torque is an Open Source PBS (Portable Batch System) for queuing jobs and managing cluster resources. ProMAX makes use of this.
- Mpich2 1.0.5: An Open Source tools and library implementation of the MPI-2 (Message Passing Interface) standard for distributed memory machines.
- Mpiexec 0.82: An alternative to the mpiexec that comes with mpich2, that enables Torque to enforce its queue policy in submitted jobs otherwise not possible with the mpich2 version, which uses external daemons (mpd) outside Torque's control.

The following is a list of planned installations and updates:

- ProMAX 2003.19.1 (2D, 3D, 4D): This update has new features and bug-fixes, and is in the final planning stage.
- Seisspace 2003.19.1: Is a ProMAX add-on flow editor with a modern interface compared to the legacy ProMAX interface.
- Geoprobe 2003.19.1 (64 bit): Geoprobe is in the words of the vendor: "The industry’s leading 3D multi-volume interpretation and visualization solution". It is to be installed on the much expected 64 bit t5 node of the cluster.

There are many different Open Source implementations of MPI (Message Passing Interface) that could be installed on Teyggjan, but for prudential reasons it has been decided to wait until the researchers/developers make specific request for these installations, so as not to clutter the system unnecessarily.

Running seismic applications on Teyggjan

Author: Khanh Duc Nguyen; NVD

Introduction

The recent upgrade of the Linux cluster "Teyggjan" and the establishment of the PUFFINS[7] Project at the Faculty of Science and Technology (NVD) have provided a capable array of computing hardware and software for the purpose of research and education in petroleum seismology and general scientific investigation (Figure 3).

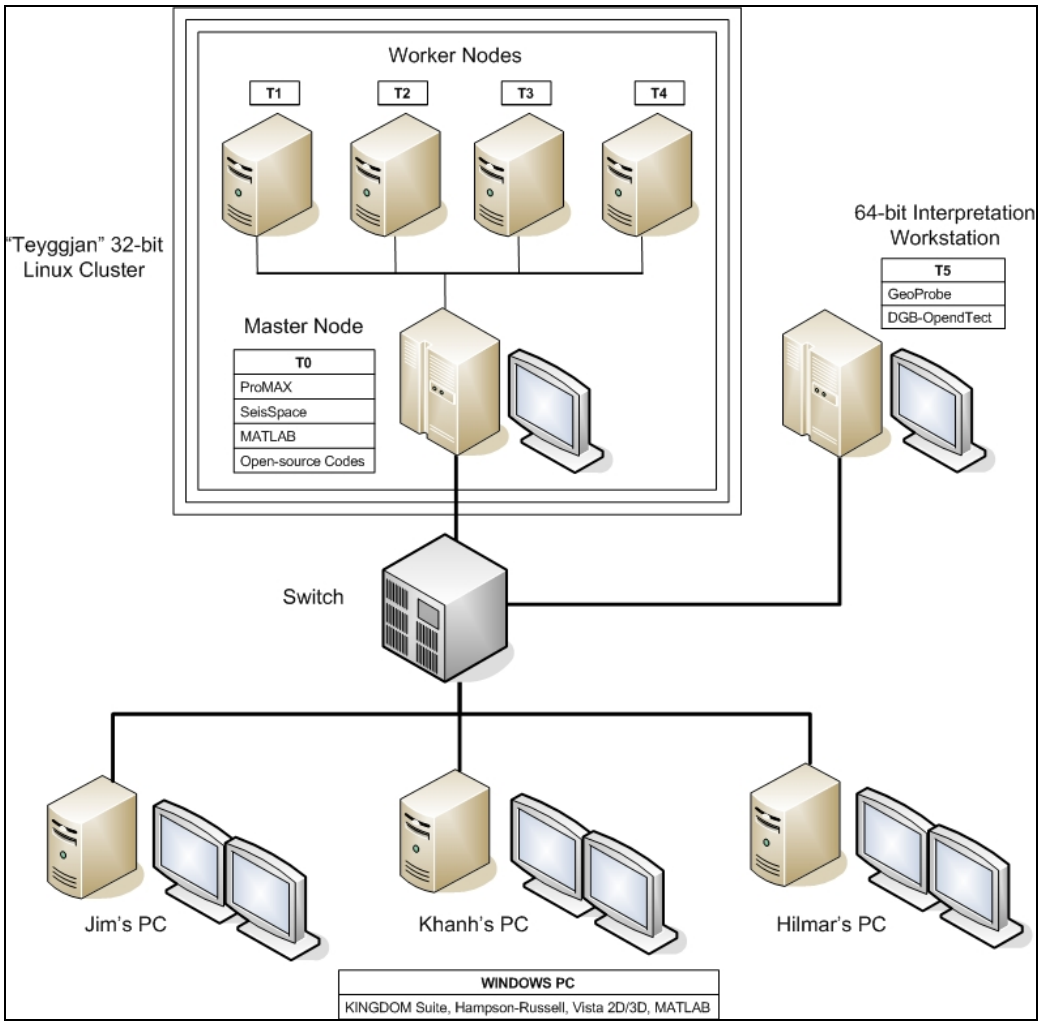

Figure 3. Computer network in the Faculty (NVD) available for research and education in Petroleum Seismology.

This chapter presents an overview of the seismic applications that are installed in Teyggjan as well as the advantage of parallel processing. As Teyggjan's usage will be increased in the future, the need to upgrade and expand the computation facilities, memory and storage capabilities of the cluster is also mentioned.

[7] Probing Underbasalt Formations of the Faroe Islands with New Science

Seismic applications and parallel processing

Through the University's Grant Program and the sponsorship to the PUFFINS project, we have access to several commercial software packages running on both Linux and Windows platforms (see Figure 3). They can be categorized as follows:

- Seismic processing and analysis: ProMAX and SeisSpace (Landmark Graphics), Vista 2D/3D (Gedco), Hampson-Russell
- Seismic interpretation and visualization: GeoProbe (Landmark Graphics), The Kingdom Suites (Seismic Micro Technology), OpendTect[8] (dGB)
- Programming Language: MATLAB (MathWorks)

Several open-source seismic packages now exist and are widely used (SU, SEPlib, FreeUSP, RSF). They are rapidly gaining acceptance for use in both academic and industry due to their versatility, large body of code developed, and low cost. In recent year, the open-source codes have received another boost from the spreading of the Linux operating system and clusters.

In this section, I will briefly introduce only those applications that can be installed in Teyggjan and take advantage of the cluster processing.

ProMAX

ProMAX is one of the most popular software for seismic data processing and analysis. Originally developed by Advanced Geophysical to run on single UNIX workstations (e.g. SunSPARC, IBM, SGI), ProMAX has been ported to Linux cluster and take advantage of its powerful and low-cost computation.

As mentioned above, several ProMAX's processes are capable of parallel processing using PVM (Figure 4a). These processes rely on PVM daemons for their inter-process and inter-node communication. When a flow is executed, one or several PVM process are spawned and displayed in the monitor window where they can be controlled by the user. Each process communicates with the interface through the PVM connection to transfer data, report results or errors. In this case, PVM enables Teyggjan to appear as a single concurrent computational resource.

Another way ProMAX achieves parallelism is to use Parallel Socket Tools (Figure 4b). This allows the user to process data within a given flow on several CPUs in different nodes within Teyggjan. Parallel Executive starts multiple copies of the socket tool on the specified machines and hooks up sockets with each socket tool (node). As a result of multiple socket tools running at once, one node can process several ensembles of seismic data in the time it takes another ensemble to process one.

In Teyggjan, ProMAX uses the TORQUE OpenPBS queue system to distribute jobs automatically over the nodes. Queue is also used to obtain the benefits of sequential

[8] There's an ongoing project at dGB to integrate the open-source processing package "Madagascar" into OpendTect.

release of serially dependent jobs, optimized system performance by controlling resource allocation and centralized management of system workload.

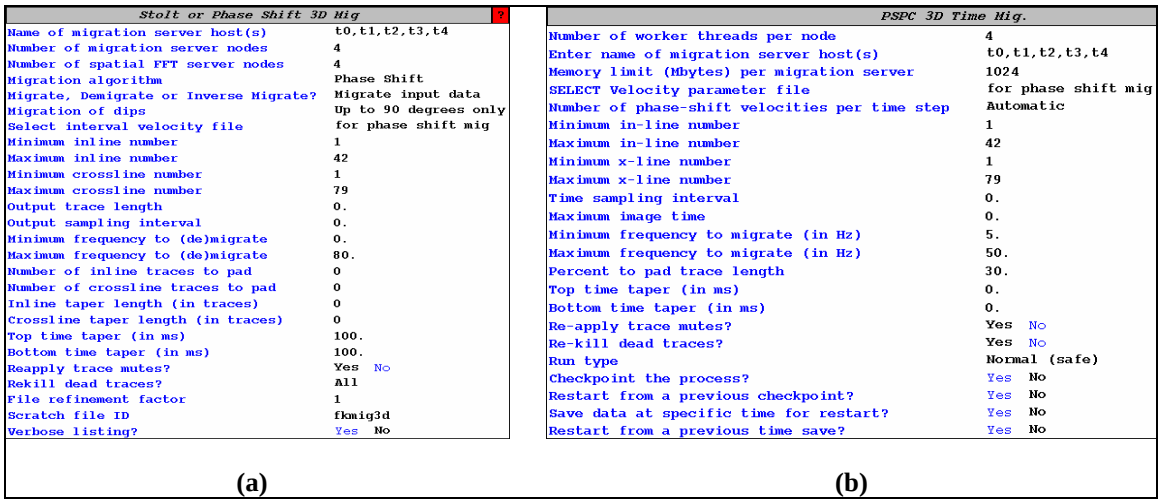

Figure 4. Example of parallel processing in ProMAX using: **(a)** PVM and **(b)** Parallel Socket Tool.

SeisSpace

SeisSpace was originally designed for multiple CPUs and multiple workstations, or cluster computing environments, for large scale, 3D marine data processing applications. It can be considered as an operation add-on for ProMAX users. Major components of SeisSpace include a user interface that conveniently consolidates several ProMAX tasks and workflows into one interface and the ability to increase user's speed and ease of use of ProMAX in a Linux cluster environment. Figures 5 and 6 below show examples of SeisSpace user interface.

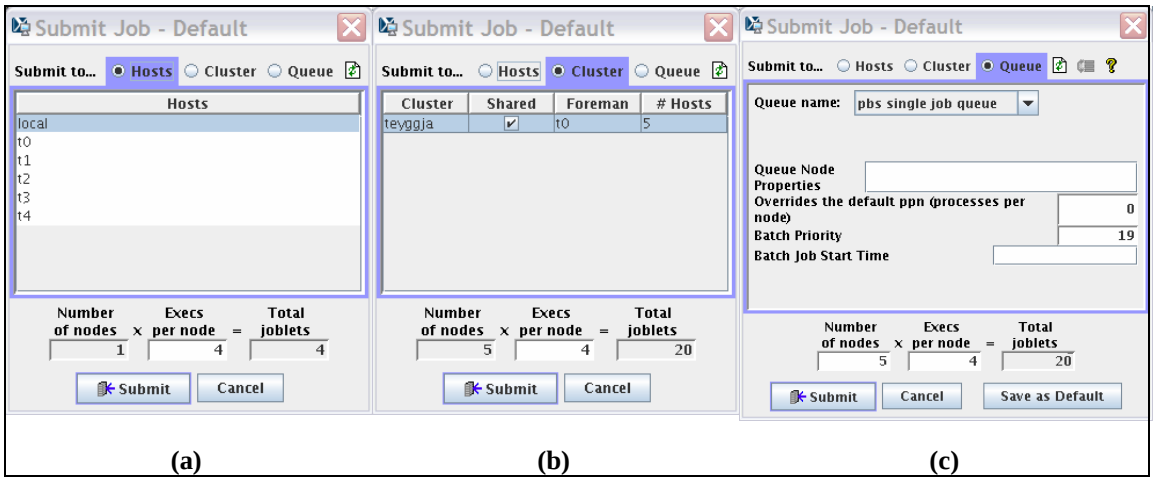

Figure 5. Job submission in SeisSpace. A job can be submitted to **(a)** a single or multiple hosts, **(b)** cluster and **(c)** Torque PBS queue. The user can specify the number of nodes and the number of CPUs per node.

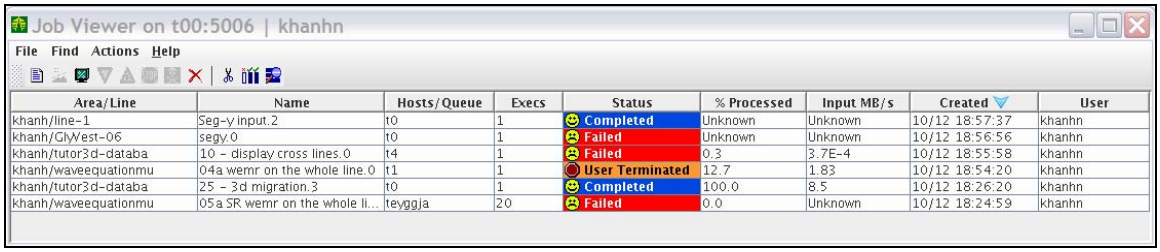

Area/Line	Name	Hosts/Queue	Execs	Status	% Processed	Input MB/s	Created	User
khanh/line-1	Seg-y input.2	t0	1	Completed	Unknown	Unknown	10/12 18:57:37	khanhn
khanh/GlyWest-06	segy.0	t0	1	Failed	Unknown	Unknown	10/12 18:56:56	khanhn
khanh/tutor3d-databa	10 - display cross lines.0	t4	1	Failed	0.3	3.7E-4	10/12 18:55:58	khanhn
khanh/waveequationmu	04a wemr on the whole line.0	t1	1	User Terminated	12.7	1.83	10/12 18:54:20	khanhn
khanh/tutor3d-databa	25 - 3d migration.3	t0	1	Completed	100.0	8.5	10/12 18:26:20	khanhn
khanh/waveequationmu	05a SR wemr on the whole li...	teyggja	20	Failed	0.0	Unknown	10/12 18:24:59	khanhn

Figure 6. The job viewer shows the status of submitted jobs in SeisSpace. The user can monitor, control (pause, resume, kill, and stop) or view detailed information regarding both active and inactive jobs.

Open-source codes

In order to use open-source codes we need a compiler. Several compiler packages are available for the Intel i686 architecture of Teyggjan. The most common ones include the GNU Compiler Collection g^{++} (GCC)[9], which is open source and free, and the Intel C/C^{++} Compiler (ICC)[10]. Portland Group (http://www.pgroup.com/) also has a compiler package available for the i686 architecture. Intel's compiler is free for non-commercial use. When it comes to FORTRAN, the GNU Compiler Collection only includes a Fortran 77 compiler, but Intel and Portland Group have compilers for both Fortran 90 and 95.

The parallel 3D elastic/viscoelastic finite difference seismic modeling (FDMPI)[11] is a parallel implementation code. Communication is performed by using the MPI standard (LAM or MPICH). FDMPI is written in ANSI C and can be compiled using both *mpicc* with MPICH or *hcc* with LAM.

The parallel 3D Prestack Kirchhoff Time Migration (PKTM)[12] is an approach that implements a parallel processing of PKTM using MPI standard (MPICH) to handle the communication and data exchange between nodes. The code is written in C and can be compiled using *mpicc* with MPICH (Seismic Unix must be installed in the cluster)

MATLAB

MATLAB is a numerical computation and visualization environment which is ideally suited to seismic data processing and analysis because of its superior built-in matrix operators and a large library of support functions. The vector syntax of MALAB once mastered, leads to more concise code than most other languages and one can perform sophisticated operations on seismic data without traditional programming. In addition to its computational functions, MATLAB also has a number of graphics and visualization functions which are useful for data analysis. For researchers whose forte is geophysics and not computational science, the high-level programming language of MATLAB is

[9] http://gcc.gnu.org/
[10] http://www.intel.com/cd/software/products/asmo-na/eng/compilers/
[11] http://www.geophysik.uni-kiel.de/bohlen/fdmpi/
[12] http://www.iamg.org/CGEditor/index.htm

able to free them from technical details to allow time to concentrate on developing their algorithms.

There are two ways to achieve parallelism using MATLAB codes:

1. Using Star-P[13] which is an interactive parallel computing platform that extends existing desktop codes for simple, user-friendly parallel computing on spectrum of computing architectures: SMP servers, multi-core servers, and clusters. The Star-P platform allows MATLAB's programmers to re-use existing codes and take advantage of parallel computing power without re-coding their algorithm in C, Fortran and MPI (Figure 7). I've tested the use of Star-P on Teyggjan but unfortunately it does not support the i686 architecture of our cluster.
2. Use MATLAB Distributed Computing Toolbox and Engine (Figure 8). The toolbox provides high-level constructs, parallel algorithms, and MPI-base functions (MPICH). Integration with the Engine enables users to develop cluster-based applications that use any scheduler (queue) or any number of worker nodes. Currently, these toolbox and engine have not been installed in Teyggjan yet.

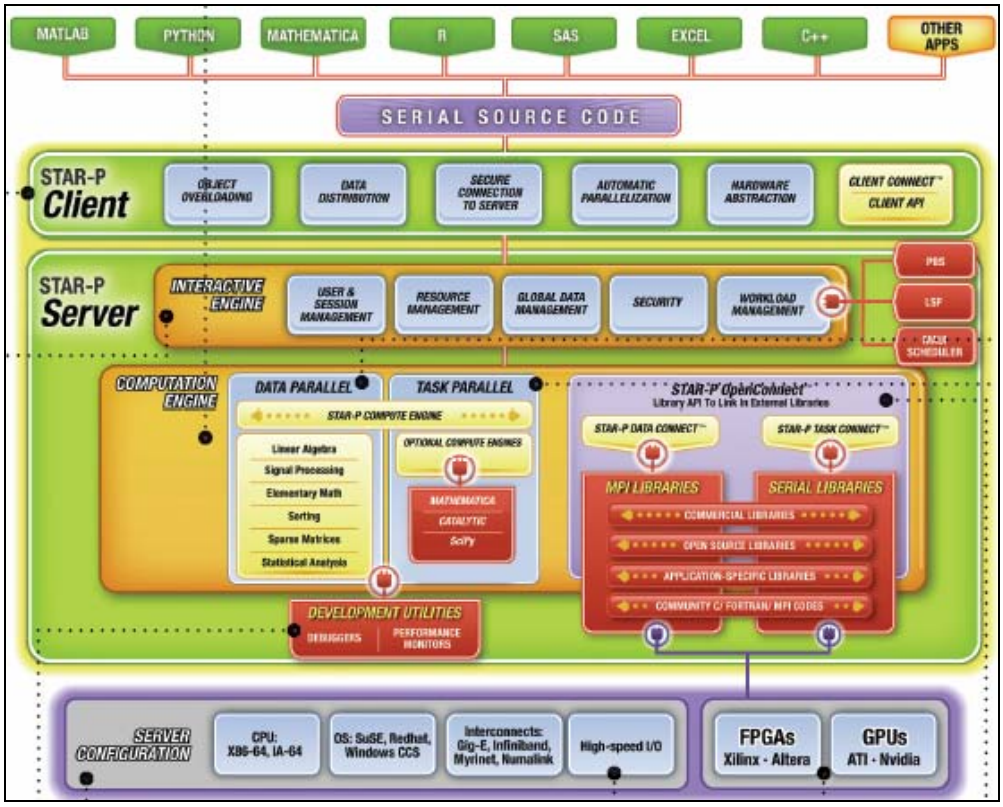

Figure 7. Star-P software is a brigde between popular computing tools such as MATLAB and the grids, servers and clusters.

Although the MATLAB Distributed Computing Toolbox and Engine have not been installed to take advantage of the parallel computing power of Teyggjan, there is still a reason for MATLAB to be installed there. It is that ProMAX – the feature-rich seismic

[13] Originally developed at MIT. Interactive Supercomputing Corporation was launch in 2004 to commercialize Star-P (www.interactivesupercomputing.com)

processing environment, can be used in conjunction with a user-friendly numerical computation environment – MATLAB. This is done through the software[14] that established a gateway between the two environments. The software allows seismic data or database information flow bi-directly between the two environments without any inconvenient processes of import/export. As a result, new seismic processing tools are easily developed without the need for tedious programming in lower-level languages such as FORTRAN and C.

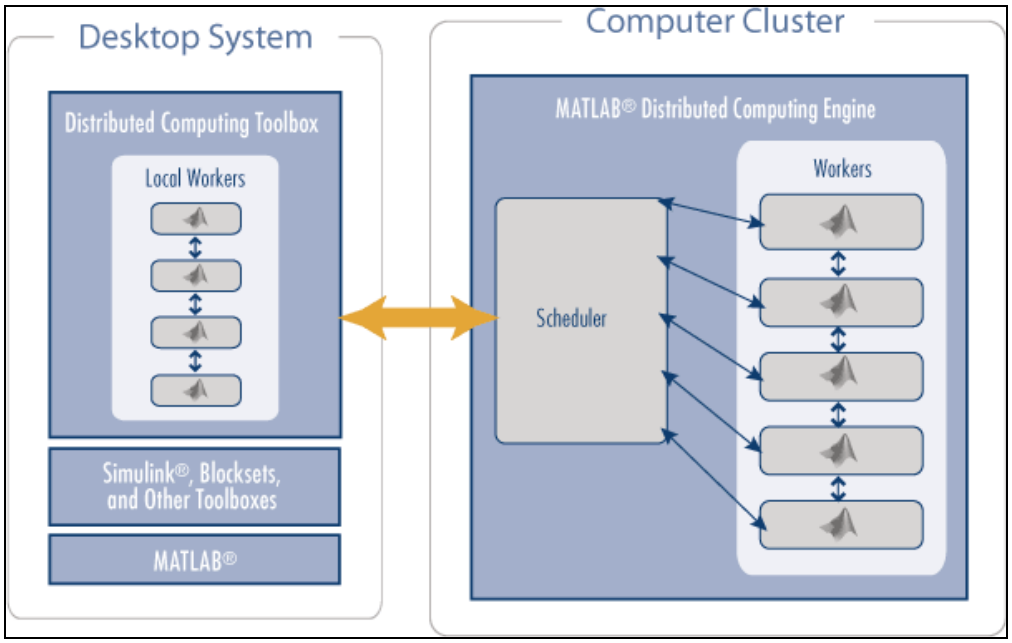

Figure 8. Distributed Computing Toolbox enables application prototyping on the desktop with up to four local workers (left) and, with Distributed Computing Engine (right), applications can be scaled to multiple nodes on a cluster.

Parallel processing and the need to upgrade teyggjan

The ability to perform calculations in parallel opens doors to a whole new way of doing computation. The PC cluster provides a platform for computationally extensive seismic data processing processes such as 3D Finite Difference Modeling or 3D Prestack Migration. Such computational requirements are generally beyond the resources for sequential platforms (single PC or workstation) and even supercomputers with shared-memory configurations. The ratio between the communication time and processing time is a critical indicator for determining the efficiency of the PC cluster. If the processing time is shorter than the communication time, as in the case of I/O process, using more CPUs can not reduce the elapsed time. However, if the processing time is longer than the communication time, as in the case of Prestack Kirchhoff Migration (Figure 9), using more CPUs can efficiently reduce the elapsed time. In this case, the elapsed time is inversely proportional to the number of CPUs.

[14] Developed at CREWES consortium, University of Calgary

The need for more computer power will always be there. Scientists often state that "if you have the computer power, we will use it". Experience of processing seismic data using PC cluster showed that the elapsed time for migrating a 2D seismic line (48805 traces, each trace has 4096 samples with 2ms sample rate) is reduced from 15h using one CPU to 1h using 15 CPUs. The elapsed time for migrating a 3D image (summation of 42 migrated images) is reduced from 630h (26 days and 6h) to 42h using 15 CPUs. Futher reduction can be achieved by using more CPUs. However, an optimal CPU number is expected for an application on PC clusters with large number of CPUs.

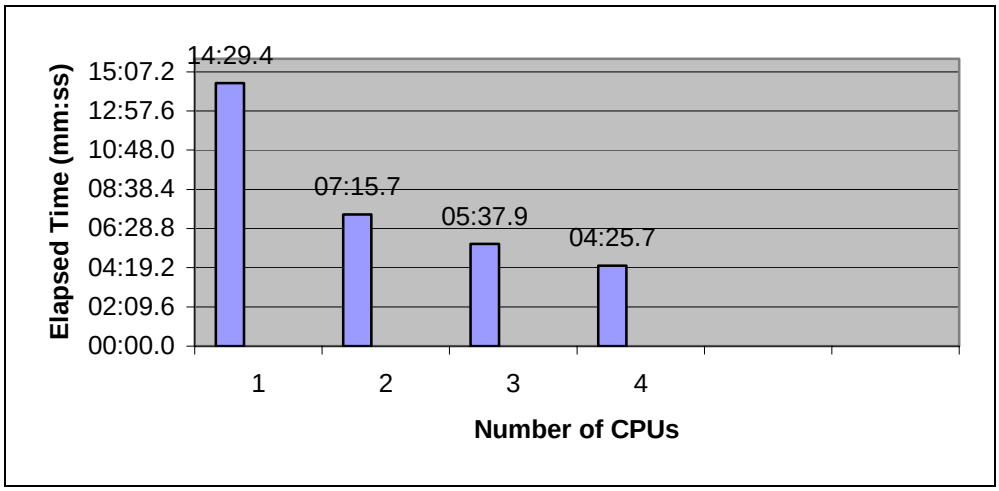

Figure 9. Kirchhoff Migration test of 3000 seismic traces (1001 ms length, 1ms sample rate) on Teyggjan.

We have recently acquired a small 2D seismic survey (GlyVeST project) and processing of these data will occupy most of the Teyggjan's resource in the next few months. Therefore, Teyggjan should be upgraded in the future for processing larger 2D or 3D seismic project. This includes computation power (more nodes), storage (more disk space) and distributed memory (Landmark recommends at least 2Gb of memory per CPU on each node for ProMAX).

Running SWAN on Teyggjan

Author: Bárður A. Niclasen; NVD

The story about SWAN so far

SWAN is an ocean wave model that has been developed at Delft University in Holland. The model was used in a recent PhD project aiming at wave forecasting in the Faroe area, incorporating the influence from tidal currents on the waves. The PhD thesis with the title "An operative wave model for the Faroe Shelf" can be found among the publications on the NVD home page, http://www.nvd.fo/index.php?id=179.

Teyggjan provided the computational power in these investigations, although only serial computations were conducted. Several computations were run on the different nodes, i.e. Teyggjan was used as 5-10 independent powerful PC's and not as a single supercomputer. Using this procedure, up to 10 different runs could be done simultaneously, each of them lasting up to 14 days. So although Teyggjan could not run in parallel mode in this period, it still provided the much needed computational power which was required in order to get the job done within reasonable time. Parallel computations could have been conducted, which would have speeded up the single computations, but the lack of local Linux support made the student postpone this option.

The main complications in running SWAN on Teyggjan are due to the fact that SWAN needs a MPI version Fortran-90 in order run in parallel mode. The factors that have slowed down this process have been the following: 1) The initial Scali software installed on Teyggjan was time limited, 2) Although Scali appears to be able to integrate Fortran-90 compilers, it does in reality only work with Fortran-77 compilers. 3) Support has been purchased from Scali, and for natural reasons they have not been very keen on providing support for rival freeware MPI compilers that are able to integrate Fortran-90 compilers. 4) With exception from shorter periods with Scali support, no funding has been available for support for the users of Teyggjan.

After funding to local Linux support has been granted, some changes have been done and enabled the SWAN model to run in MPI-mode, i.e. Teyggjan can now function as one united supercomputer when running the SWAN code.

Parallel tests

A series of model tests have been conducted with the SWAN code, which give some insight into what gains are obtained when running in different MPI setups. The results from these tests are given below in Figure 10 and Table 1, and some additional details are given in Appendix 2.

The main results from these tests were as follows:

1) Teyggjan has plenty of memory capacity for SWAN runs.

2) Memory distribution between the nodes is almost linear, meaning that if two nodes are used only half the memory is used by each compared to the same simulation using only one node etc.
3) All nodes now compute with the same speed (Table A2 in Appendix 2), this was not the case before the recent upgrade of Teyggjan.
4) Some variation was observed in the computational time of identical parallel runs, but the variation was usually below 10% of the total computation time.
5) The fastest computations with a given number of nodes were achieved by using two treads per note (see Test 1 in Appendix 2) and the fastest results were obtained when using all 5 nodes.
6) Parallel SWAN computations can be conducted in less than 15% of the time required to conduct the serial SWAN computation for the same problem (last row Table 1).
7) The speed-up of using additional nodes is almost constant for runs using the same number of treads per node (one or two), giving almost half computation time each time the number of nodes is doubled (last two columns of Table 1).

There is therefore no clear indication of any saturation limit in the computation time with the respect to the number of nodes used for parallel SWAN runs on Teyggjan. Experienced SWAN users have found that computation times of parallel implementations of SWAN tend to saturate around 30 nodes, i.e. no speedup in computation time is acquired past this number of nodes. As Teyggjan only has 5 nodes, there is still significant room for further improvement with respect to speeding up the parallel SWAN computations.

Future use of SWAN

There is an ongoing project (January 2007 to January 2009) which aims at describing and characterizing the areas on the Faroe Shelf, where tidal currents and bathymetry can effect the incoming waves and potentially generate dangerous seas. Another topic of this project aims at describing where and how much the waves affect light sediment resuspension in the inner areas.

Both topics require high resolution SWAN runs that will demand large parts of Teyggjan's computational power for long time spans. It is estimated that each simulation will run for continuous periods in the order of several days to a few weeks. This makes it necessary to discuss how different computing jobs should be prioritized on Teyggjan.

Table 1. Run times for parallel SWAN runs of Test 2 (see A1) with different number of nodes and treads. The simulation with '*' is the serial reference run.

# notes	# treads	Total time In seconds	Time ratio vs. serial	Time ratio 1 tread	Time ratio 2 treads
1*	1*	2832	1.00	-	-
1	1	3190	1.13	1.00	-
1	2	1955	0.69	-	1.00
2	2	1660	0.59	0.52	-
2	4	1066	0.38	-	0.55
3	3	1151	0.41	0.36	-
3	6	701	0.25	-	0.36
4	4	853	0.30	0.27	-
4	8	533	0.19	-	0.27
5	5	681	0.24	0.21	-
5	10	416	0.15	-	0.21

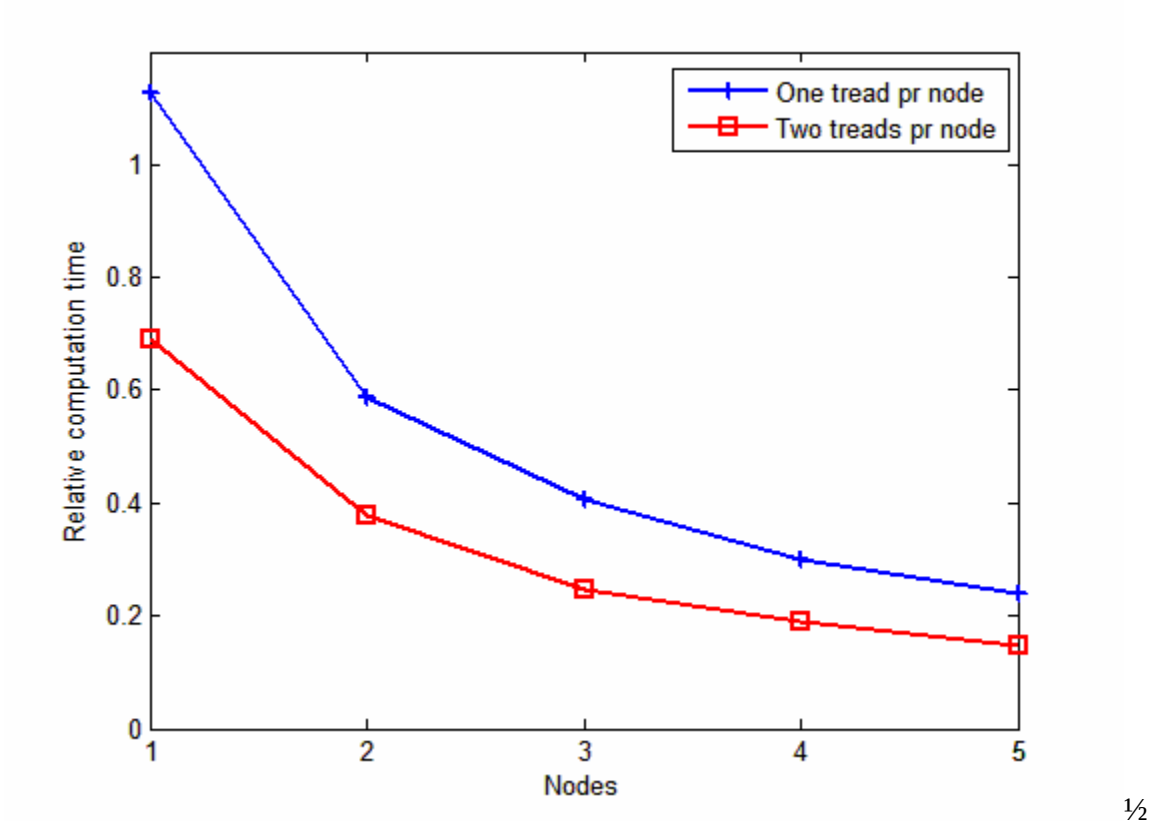

½

Figure 10. Relative computational time of Teyggjan when working on a SWAN problem, using 1-5 nodes and one or two reads pr node. Relative computational time equal to one corresponds to a serial run of the problem.

Appendix 1: detailed configuration of Teyggjan

```
#
# Host              : t0
#
# model name        : Intel(R) Xeon(TM) CPU 2.80GHz       x      4
# OS                : Linux 2.4.21-47.ELsmp
# MemTotal          : 6167588 kB
# Compiler          : g++ (GCC) 3.2.3 20030502 (Red Hat Linux 3.2.3-56)
#
```

```
#
# Host              : t1
#
# model name        : Intel(R) Xeon(TM) CPU 2.80GHz
                      x      4
# OS                : Linux 2.4.21-47.ELsmp
# MemTotal          : 5907436 kB
# Compiler          : g++ (GCC) 3.2.3 20030502 (Red
                    Hat Linux 3.2.3-56)
#
```

```
#
# Host              : t3
#
# model name        : Intel(R) Xeon(TM) CPU 2.80GHz
                      x      4
# OS                : Linux 2.4.21-47.ELsmp
# MemTotal          : 5907436 kB
# Compiler          : g++ (GCC) 3.2.3 20030502 (Red
                    Hat Linux 3.2.3-56)
#
```

```
#
# Host              : t2
#
# model name        : Intel(R) Xeon(TM) CPU 2.80GHz
                      x      4
# OS                : Linux 2.4.21-47.ELsmp
# MemTotal          : 5907436 kB
# Compiler          : g++ (GCC) 3.2.3 20030502 (Red
                    Hat Linux 3.2.3-56)
#
```

```
#
# Host              : t4
#
# model name        : Intel(R) Xeon(TM) CPU 2.80GHz
                      x      4
# OS                : Linux 2.4.21-47.ELsmp
# MemTotal          : 5907436 kB
# Compiler          : g++ (GCC) 3.2.3 20030502 (Red
                    Hat Linux 3.2.3-56)
#
```

The above information is obtained using the following script:

```
#! /bin/sh
#
# Purpose: check
#         1) Hostname,
#         2) CPU model, speed and number of CPU
#         3) Operating System
#         4) Compiler
#
# Date: Sep-15-2007
#
# Modified:
#
# Author: KDN
#
echo    "# "
echo    "# Host: `hostname -s`"
echo    "# "
echo -n "# `grep -e 'model name' -e 'cpu        ' '/proc/cpuinfo'  | uniq`"
echo    " x `grep 'processor' '/proc/cpuinfo'  | wc -l`"
echo    "# OS: `uname -sr`"
echo    "# `grep 'MemTotal' '/proc/meminfo'`"
echo    "# Compiler: `g++ --version | grep g++`"
echo    "# "

exit 0
```

Appendix 2: Details from SWAN runs

Table A1. Computational examples from two SWAN setups, which are used to test the performance of Teyggjan.

	Test 1	**Test 2**
Number of time steps	12	12
Number of grip points	112 x 125	360 x 330
Total number of variables to compute pr. time step	11.760.000	99.792.000

Table A2. Results from serial (single processor) runs of Test 1, conducted on the different nodes.

Note name	**Total time in seconds**	**CPU load**	**Memory load**
t0	264	25%	1.6%
t1	254	25%	1.7%
t2	260	25%	1.7%
t3	256	25%	1.7%
t4	260	25%	1.7%

Table A3. Results from parallel SWAN runs of Test 1, using different number of nodes and treads. The computational times are given by the SWAN output, and correspond to computational time of one tread. Time ratio to serial the serial run is calculated by assuming 260 second duration in the serial run (Table A2). The communication time corresponds to the total time spent transporting data between the different nodes and treads.

# notes	**# treads**	**Total time in seconds**	**Time ratio to serial run**	**Comp. time in seconds**	**Comm. time in seconds**
1	1	305	1.17	270	30
1	2	182	0.70	142	40
1	3	277	1.07	191	83
1	4	283	1.09	211	69
2	2	164	0.63	123	39
2	3	149	0.57	116	31
2	4	102	0.39	71	29
2	5	174	0.70	55	117
2	6	199	0.77	151	46
3	3	106	0.41	88	17
3	6	68	0.26	45	20
4	4	80	0.31	61	18
4	8	52	0.20	35	16
5	5	66	0.26	48	17
5	10	48	0.18	25	22

Appendix 3: Færøernes Universitet, Løsning fra Dell, 9. November 2007

HPCC Oplæg til Færøernes Universitet

I forbindelse med fremtidigt HPCC cluster har Dell fået lov til på baggrund af oplysninger fra Færøernes Universitet at komme med et løsningsoplæg til dette.

De vigtigste faktorer ved et HPCC beregningscluster er at få så meget performance ud af løsningen som muligt indenfor de givne rammer. Applikationerne er normalt fastlagt på forhånd da det er meget specifikke beregninger som ønskes udført og tit med special udviklet s/w til lige netop disse beregninger. Derfor skal h/w tilpasses så optimalt som muligt til applikationen/erne for at få den mest optimale performance i forbindelse med beregningerne. De vigtigste faktorer i forbindelse med dette er CPU, ram og netværk – det er disse tre faktorer som normalt er de vigtigste i forbindelse med supercomputing da det er disse faktorer som afgør hvor hurtigt beregningerne kan foretages alt afhængig af applikationerne.

Til dette HPCC cluster hos Færøernes Unisversitet ønkes følgende applikationer anvendt:
- ProMAX (version 2003.19.1.1)
- Seisspace fra Landmark.
- SWAN (bølge og havstrøms-modellering) beroende på MPICH2
- Matlab (fra MathWorks)

Færøernes Universitet ønsker et cluster i størelsesordnen 10-20 noder.

I forbindelse med valg af h/w til at afvikle ovenstående applikationer skal der primært tages beslutning omkring valget omkring brugen af Intel og AMD og om der skal benyttes DualCore eller QuadCore CPU´er, igen kan valget hænge sammen med hvordan applikationerne performer i forbindelse med disse to teknologier.
AMD er specielt godt til applikationer som er meget ram intense på grund af den måde AMD´s chipset og busteknologi er bygget op. AMD allokerer ram direkte til hver CPU hvorimod Intel deler ram´en mellem de fysiske CPU´er.

Derudover så er der også mulighed for at vælge blades, rent performancemæssigt er der ingen forskel på om der vælges blades eller traditionelle servere, men i dette setup hvor clustret vil komme til at bestå af 10-20 noder så vil det rent økonomisk ikke være nogen fordel at benytte blades da man man også skal indregne omkostninger til bladechassis´erne, der vil ikke være den store pladsbesparelse – ved 20 noder vil pladsbesparelsen være 6U i et rack, rent strømmæssigt vil der heller ikke være den store besparelse – det vil først være interessant til løsninger på min. 50 noder.

Andre faktorer som spiller ind er mængden af ram samt interconnect (ethernet eller Infiniband f.eks.)

For at kunne tage den rigtige beslutning omkring valg af enten Intel eller AMD så bliver man nød til at kigge på applikationerne og se på hvilken platform de performer bedst på og så vælge udfra dette, umiddelbart vil AMD QuadCore nok være det bedste valg når de indenfor ganske kort tid blever leverbare i PE1435SC serverne.

I øjeblikket kan vi levere QuadCore Intel chipset/CPU´er og i øjeblikket DualCoreAMD chipset/CPU´er (levering af QuadCore AMD er som sagt udsat lidt – løsningen er prissat med DualCore, men QuadCore priserne vil lægge sig meget tæt op af DualCore)

Derfor har jeg valgt at lave et oplæg på vores Intel QuadCore PE1950 samt option på vores PE1435SC server som er decideret designet til HPCC da den er designet på en platform som primært giver mulighed for en masse CPU kraft og masser af ram – hvorimod redundans ikke er mulig på denne. OS på maskinerne vil være Redhat Linux 4 WS x64.

Som interconnect er valgt en PowerConnect 6248 gigabit ethernet, da dette ifølge Færøernes Universitet dækker behovet for performance til interconnect.

Som Masterserver (til indsamling af data og til management) er valgt en PE2950 server med tilhørende diskkabinet til opsamlig af data. OS på Master serveren vil være Redhat Linux 4 ES x64.

Til management/overvågnings netværk er valgt en PowerConnect 3448 10/100 switch.

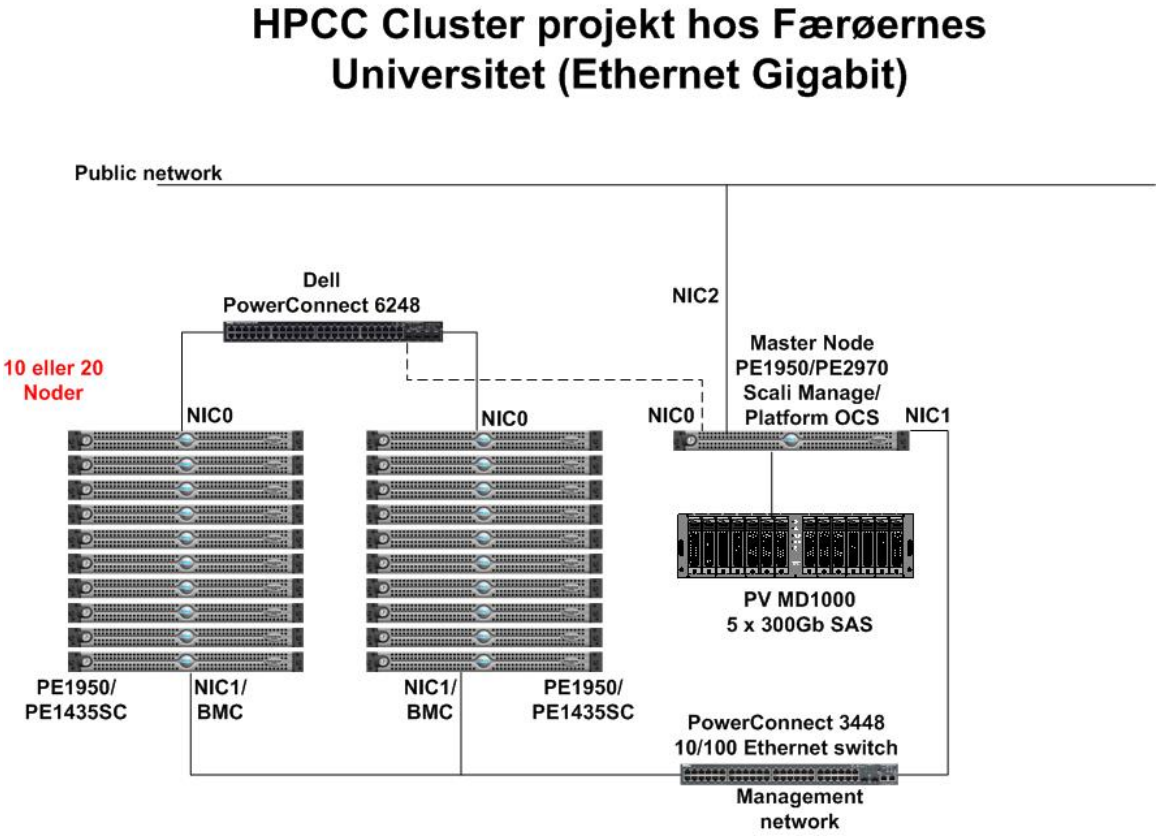

Priser

Products	Total
10 nodes HPCC Cluster	
10 x PowerEdge 1950 (Intel) 2 x 2.33 Ghz QuadCore 2 Gb ram (4x512) 1 x 73Gb 15k SAS disk Non-redundant strømforsyning Inkl. Redhat Linux 4 WS, 3 års Subscription 3 års NBD service	
1 x PowerEdge 1950 (Intel) 1 x 2.33 Ghz QuadCore 2 Gb ram (4x512) 2 x 73Gb 15k SAS disk 1 ekstra RaidController 1 ekstra Dualport netkort Redundant strømforsyning Inkl. Redhat Linux 4 ES, 3 års Subscription 3 års NBDservice	DKK 196.000,00
Løsningen med 20 Noder i stedet for 10	DKK 358.000,00
Eller	
10 x PowerEdge 1435SC (AMD) 2 x 2.8 Ghz 1333FSB DualCore 2 Gb ram (4x512) 1 x 73Gb 15k SAS disk Non redundant strømforsyning Inkl. Redhat Linux 4 WS, 3 års Subscription 3 års NBD service	
1 x PowerEdge 2970 (AMD) 1 x 2.8 Ghz DualCore 2 Gb ram (4x512) 2 x 73Gb 15k SAS disk 1 ekstra RaidController 1 ekstra Dualport netkort Redundant strømforsyning Inkl. Redhat Linux 4 ES, 3 års Subscription	
Fælles for begge løsninger	

1 x PowerVault MD1000 diskhylde 5 x 300Gb 10k SAS 3 års NBDservice	
1 x Dell 42U rack 1 x 16 ports KVM switch + 1 x 8 ports udvidelses moduler 20 x USB kabler til KVM 1 x LCD skærm inkl. tastatur 3 års NBDservice Ingen UPS er inkl. i prisen	
Interconnect (til clusternoderne)	
1 x PowerConnect 6248 Gigabit ethernet Layer 3 switch Ingen RJ-45 kabler inkluderet 3 års NBD service	
Management switch til clusteret	
1 x PowerConnect 3448 10/100 ethernet switch Ingen RJ-45 kabler inkluderet 3 års NBD service	DKK 76.000,00
HPCC Management s/w	
Scali Manage 20 Cluster Noder Ex. installation Inkl. 3 års subscription	DKK 41.800,00
Samme løsning med 10 Noder	DKK 20.900,00
Platform OCS Kan der gives tilbud på hvis dette ønskes ?	

Almindelige betingelser for Dell Danmark

1. Definitioner

Betingelser: disse almindelige betingelser

Custom Factory Integration ("CFI"): en Service, som kombinerer Kunde-specificeret software og hardware med Produkt(er), og som bl.a. kan omfatte indlæsning af image og applikationer og vedligeholdelse heraf samt software- og hardwareintegration og/eller inventarhåndtering.

Dell: den Dell-enhed, som er anført i din Ordrebekræftelse og/eller på din faktura;

Forbruger: en fysisk person, som køber eller aftaler at købe Produkter og/eller Services fra Dell hovedsageligt til privat brug

IM: "Integrationsmateriale", defineret som tredjepartsprodukt(er), der er specificeret eller leveret af Kunden i henhold til CFI

Immaterielle Rettigheder (IPR): patenter, varemærker, registerede design og ansøgninger om registrering heraf samt ophavsrettigheder, designrettigheder, knowhow, firma- og forretningsnavne og lignende beskyttede rettigheder i ethvert land

Kunde: juridisk eller fysisk person, som køber eller aftaler at købe Produkter og/eller Services fra Dell, herunder Forbrugere

Købspris: den samlede pris, som Kunden skal betale Dell for Produkter og/eller Services

Ordrebekræftelse: Dells skriftlige bekræftelse af Kundens ordre

Produkter: en individuel vare (herunder Software) som beskrevet i et gældende dokument, som Dell har udgivet fysisk og/eller på sit websted, eller som beskrevet i en Ordrebekræftelse, og som Kunden køber eller aftaler at købe af Dell, bortset fra komponenter, der tilføjes Dell-hardware via CFI

Services: almindelig service og support leveret af Dell eller Dells underleverandører i henhold til Servicetilbudet, inkl. CFI

Servicetilbud: de services, der udbydes af Dell, som beskrevet i et gældende dokument, som Dell har udgivet fysisk og/eller på sit websted, eller som beskrevet i en Ordrebekræftelse

Software: operativsystemer, middleware/tilpasningsprogrammer, applikationer eller anden software, som fremstilles eller ejes af, eller er givet i licens til Dell

Tredjemandsprodukter: Produkter, som Dell sælger, men som ikke er fremstillet, samlet eller frembragt af Dell

Tredjemandssoftware: operativsystemer, middleware/tilpasningsprogrammer, applikationer eller anden software, som er udgivet eller givet i licens af en tredjemand.

2. **Anvendelse**

2.1 Disse Betingelser finder anvendelse på alle aftaler om Kunders køb af Produkter og/eller Services fra Dell.

2.2 Disse Betingelser træder i stedet for alle andre vilkår og betingelser, medmindre andet er skriftligt aftalt med Dell.

2.3 Alle ordrer på Produkter og/eller Services, der afgives af Kunden, anses som et tilbud om at købe de pågældende Produkter og/eller Services i henhold til disse Betingelser.

2.4 Dell accepterer Kundens tilbud om køb i henhold til disse Betingelser ved at udstede en Ordrebekræftelse til Kunden. Kunden skal gennemgå Ordrebekræftelsen og straks kontakte Dell, hvis Kunden opdager eventuelle fejl eller uoverensstemmelser. I modsat fald vil Dell fremstille og levere Produktet i overensstemmelse med Ordrebekræftelsen, som er bindende.

3.**Tilbud/Ændringer heraf**

3.1 Dells tilbud gælder kun, hvis de foreligger skriftligt, og kun i 10 dage efter afgivelsen af tilbudet, medmindre andet er anført i tilbudet.

3.2 Da Dells politik er konstant at forbedre sine Produkter og Services, og da Kunden muligvis køber Tredjemandsprodukter, forbeholder Dell sig ret til at ændre specifikationerne for Produkter og Services i forhold til det i Ordrebekræftelsen anførte. Dell garanterer, at Produkter og Services altid vil have mindst samme funktionalitet og ydeevne, og Dell vil ikke foretage væsentlige ændringer, uden forudgående aftale herom med Kunden.

4. **Købspris og betaling**

4.1 Den Købspris, som Kunden skal betale, vil fremgå af Dells Ordrebekræftelse og faktura.

4.2 Betaling skal ske før levering eller Service eller, hvis dette er aftalt skriftligt, senest 15 dage fra fakturadato. Dell kan indstille levering, indtil der er sket betaling af det fulde beløb.

4.3 For større ikke-forbruger ordrer, som indebærer levering i fremtiden, har Dell ret til at regulere priserne som følge af ændrede

valutakurser, afgifter og omkostninger til forsikring, fragt og indkøb.

4.4 Ved for sen betaling beregnes rente på 2% pr. måned af det skyldige beløb. Dell kan tilbageholde leverancen i tilfælde af for sen betaling. Hvis Dell er nødsaget til at inddrive sit tilgodehavende og/eller tage Produktet tilbage, afholdes alle omkostninger hertil af Kunden.

5. **Levering**
Leveringsdatoen i Ordrebekræftelsen er omtrentlig. Leveringsstedet er som anført i Ordrebekræftelsen.
Af praktiske årsager kan del-levering af Produkter finde sted (fx ved levering af Tredjemandsprodukter, fremstillet på et andet tidspunkt end et Dell produkt).
Eventuelle manglende, forkerte eller skadede Produkter eller emballage skal anføres i fragtbrevet før kvittering for modtagelse.

6. **Ejendomsret og risiko**
Ejendomsretten til Produkterne overgår først, når Dell har modtaget betaling af det fulde beløb. Dell kan til enhver tid tage leverede Produkter tilbage, før ejendomsretten er overgået, hvis Kunden misligholder Betingelserne. Denne bestemmelse gælder ikke IPR.
6.2 Risikoen overgår til Kunden eller dennes repræsentant ved levering af Produkterne.

7. **Accept og Forbrugeres fortrydelsesret**
Ikke-forbruges accept: Ikke-forbrugere kan kun afvise at acceptere Produkter på grund af væsentlige mangler ved skriftlig reklamation over for Dell senest 7 dage **efter** leveringen; ellers anses Kunden at have accepteret Produkterne.
Fortrydelsesret for Forbrugere: Forbrugere kan fortryde deres ordrer i op til 14 dage efter modtagelsen af Produktet. Dell kan anmode om en skriftlig bekræftelse af udnyttelse af fortrydelsesretten. Forbrugers fortrydelsesret gælder ikke for software og Services, hvis Forbruger tager disse i brug, eller for Produkter, hvis Forbrugeren har gjort disse uegnede til salg. Produkter returneres for Forbrugers egen risiko, i væsentlig uændret stand, samt i originalemballage. Dell vil senest 30 dage efter modtagelsen af Produktet refundere Købsprisen med fradrag af eventuelle rimelige udgifter i forbindelse med tilbagetagelsen, herunder returfragt, og et rimeligt beløb i tilfælde af beskadigede Produkter eller Produkternes hændelige undergang. For at sikre, at Produkter returneres i overensstemmelse med ovennævnte, oplyser Dell en returadresse og et returnummer, som bør fremgå tydeligt af det emballerede Produkt.
Hvis Kunden ved ibrugtagningen vælger ikke at acceptere en eventuel operativsystemlicens, accepterer Dell alene at tage hele Produktet retur og refundere Købsprisen, hvis dette sker inden for de i hhv. pkt 7.1 og 7.2 nævnte frister.
7.4 Produkter, der ønskes returneret, skal kunne afhentes på det tidspunkt og på sted, som Dell med rimelighed anmoder om.

8. **Services**
Materialer og udførelse
8.1 Dell vil fra og med leveringsdatoen reparere eller omlevere mangelfulde Produkter, bortset fra Tredjemandsprodukter, i overensstemmelse med det aftalte serviceniveau og indenfor den i Servicetilbudet aftalte periode. Alle Servicetilbud er tilgængelige på Dells websted og kan ses forud for køb. Forbrugeres ret til reparation eller omlevering af mangelfulde Produkter, herunder Tredjemandsprodukter, følger af ufravigelige retsregler i forbrugerlovgivningen og går forud for, hvad der måtte være anført i Dells Servicetilbud. For Ikke-forbrugere anses fravigelige retsregler i lovgivningen fraveget ved denne bestemmelse.
For reservedele gives en garanti på 90 dage fra leveringsdatoen eller restperioden af det pågældende Servicetilbud, såfremt denne er længere. Dell ejer alle Produkter eller dele heraf, der fjernes og erstattes under reparation. Dell opkræver betaling fra Kunden, hvis Produkter eller dele heraf, der fjernes og erstattes under reparation, ikke returneres efter påkrav herom.
Dell videregiver til Kunden de garantier, som er ydet af Dells producent eller leverandør af Tredjemandsprodukter. Dell yder ingen garanti for IM eller Tredjemandssoftware, som typisk leveres af en licenstager af den pågældende Tredjemandssoftware.
Dell vil i videst muligt omfang bestræbe sig på at opfylde de i Servicetilbuddet anførte svartider. Disse kan variere afhængig af afstanden og tilgængelighed til Produkternes lokation, samt af tilgængeligheden af komponenter. Bortset fra pkt. 11 er Dell ikke ansvarlig for eventuelle direkte eller indirekte tab forårsaget af Dells manglende opfyldelse af relevante svartider, ligesom Dell ikke er ansvarlig for eller forpligtet i forhold til eventuelle mangler i Produkterne eller manglende afhjælpning af mangler inden for en rimelig tid.
Dells forpligtelser er med forbehold for korrekt brug af Produkterne og gælder ikke for eventuelle dele af Produkter, der er blevet ændret eller repareret uden Dells forudgående skriftlige

samtykke. Dells forpligtelser vil ikke være gældende, hvis en mangel skyldes eksterne forhold såsom ulykke, uheld, fugt, elektrisk påvirkning eller andre forhold eller andet anført i det Servicetilbud, der købes sammen med eller er tilknyttet Produkterne. Uanset ovenstående er følgende undtaget fra Service, medmindre andet er anført i Servicetilbuddet: Arbejde uden for den normale lokale arbejdstid, flytning, forebyggende vedligeholdelse, forbrugsartikler, IM; reparation af Produkter, som fungerer i overensstemmelse med branchestandarder, herunder bl.a. defekte pixels på skærme, overførsel af data eller Software samt virus. Der foretages ikke service på dele, som ikke er kritiske for Produktets funktion, herunder bl.a. hængsler, døre, kosmetiske elementer og stel.
Ved reparation benytter Dell komponenter, som er nye eller fremstår som nye i overensstemmelse med branchestandarder og –praksis. Batterier til Notebook leveres med max. ét års garanti, medmindre en kortere garantiperiode er anført i det pågældende Servicetilbud, i hvilket tilfælde denne kortere garanti ikke kan opgraderes.
Egnethed til bestemt brug og krav til ydeevne
8.7 Dell yder ingen garanti for, at Produkterne er egnet til en bestemt brug, eller at Produkterne vil opfylde bestemte krav til ydeevne, medmindre andet er skriftligt aftalt med Dell.
Bortset fra det i disse Betingelser udtrykkeligt anførte, yder eller påtager Dell sig ingen udtrykkelige eller stiltiende garantier vedrørende Produkternes stand, kvalitet, ydeevne, salgbarhed eller holdbarhed, og alle sådanne garantier fraskrives hermed. Kunden er ansvarlig for at fjerne produkter, der ikke er leveret af Dell, og tage sikkerhedskopier af og hemmeligholde alle data i Produktet, herunder før levering af Services.
Service kan leveres pr. telefon eller via Internet, hvor dette måtte være hensigtsmæssigt. Telefonopkald kan blive optaget til uddannelsesformål. Kunden skal være samarbejdsvillig og yde Dell al rimelig bistand og information, som er nødvendig for, at Dell kan levere de pågældende Services. Kunden betaler alle telefon- og forsendelsesudgifter i forbindelse med henvendelser til Dell, medmindre Dell er forpligtet til at betale sådanne udgifter ifølge ufravigelige retsregler i forbrugerlovgivningen.

9. **Custom Factory Integration (CFI)**
CFI kan leveres af Dell i overensstemmelse med Kundens anvisninger , tekniske specifikationer og gældende Dell-krav. Kunden giver specifikationer på eller leverer IM, eller Dell fremskaffer IM efter Kundens anvisninger. Dell oplyser, om IM accepteres og/eller godkendes, og vil herefter integrere IM i Produkterne, hvorved et CFI-Produkt fremkommer. Dell kan installere CFI-Produktet ifølge Kundens anvisninger eller efter Dells egen tekniske rådgivning, hvis dette er aftalt skriftlig. Dell vil ikke udføre CFI-relateret arbejde, hvis dette ikke er teknisk muligt.

10. **Uforudsete begivenheder**
Parterne er ikke ansvarlige for manglende opfyldelse, som skyldes forhold, som den pågældende part ikke har indflydelse på (force majeure), herunder strejker, terroristhandlinger, krig, leverandør-/transport-/produktionsproblemer, valutakursudsving, regerings- eller myndighedshandlinger og naturkatastrofer. En part har ret til en rimelig forlængelse af frister, dog således at hvis de pågældende forhold varer i mere end 60 dage, har parterne ret til med skriftligt varsel at ophæve aftalen uden godtgørelse.

11. **Ansvarsbegrænsning**
Dell accepterer at være ansvarlig for tingsskade, dødsfald eller personskade forårsaget af uagtsomhed udvist af eller forsætlige handlinger foretaget af Dell eller medarbejdere, befuldmægtigede eller underleverandører, der handler på Dells vegne. Dells ansvar kan under ingen omstændigheder, bortset fra ved dødsfald eller personskade (hvor der ikke gælder nogen ansvarsbegrænsning), overstige et beløb svarende til 125% af Købsprisen. Dell er over for Kunder ikke ansvarlig for:
tab hidrørende fra mangler eller defekter i Produkter eller Services, som Dell har afhjulpet inden for en rimelig tid,
indirekte tab eller følgetab, herunder tab af omsætning, overskud, løn, indtægter, data eller forventede besparelser,
tab, som kunne være undgået ved, at Kunden havde fulgt Dells rimelige råd og anvisninger,
tab som følge af Dells brug af materiale eller anvisninger, leveret eller specificeret af Kunden, eller
tab, der udspringer af, at Kunden ikke har foretaget fuldstændige og opdaterede sikkerhedskopier af edb-programmer og data.
12. Immaterielle rettigheder (IPR) og Software

Dell skadesløsholder Kunden for omkostninger og ansvar i forbindelse med krav rejst som følge af at brug af Produktet krænker tredjemands Immaterielle Rettigheder. Dell kan i så fald tilbagekalde og udskifte eller ændre Produktet eller refundere Kunden Købsprisen med fradrag for værdiforringelse. Kunden skadesløsholder Dell for alle krav udsprunget af IM og IPR, der er specificeret eller ejet af Kunden og integreret i Produktet. Dell har tilladelse til at føre retssager

vedrørende krav, herunder forhandlinger og indgåelse af forlig, og Kunden skal bistå Dell i forbindelse hermed, hvis retssagen er direkte forbundet med Kundens Produkt. Al IPR i et Produkt, der tilhører Dell, forbliver Dells. Kunden skal straks underrette Dell om al krænkende eller uautoriseret brug af Produkter eller IPR. Dell skadesløsholder ikke Kunden for: (i) Tredjemandsprodukter og Tredjemandssoftware, (ii) uautoriserede ændringer eller uautoriseret brug eller (iii) krav forårsaget af brug af Produkter sammen med produkter, der ikke er leveret af Dell. Kunden skal overholde licensbetingelserne for al software.

13. **Eksportrestriktioner**
Kunden er opmærksom på, at Produkter, som kan indeholde teknologi og software, er underlagt eksportrestriktioner i EU- og amerikansk lovgivning samt lovgivningen i det land, hvor Produkterne leveres til eller benyttes. I henhold til sådan lovgivning gælder der særlige restriktioner for Produktet vedrørende salg, udlejning og overdragelse til visse slutbrugere, slutbrug og lande. Kunden accepterer at overholde sådan lovgivning.

14. **Databehandling**

Persondata, som Dell har fået fra Kunden, opbevares og behandles i overensstemmelse med gældende lovgivning og Dells retningslinjer for beskyttelse af privatlivets fred. Dell kan videregive sådanne persondata til andre Dell-enheder eller Dells befuldmægtigede eller underleverandører, som udfører tjenesteydelser for Dell. Dell kan ligeledes overføre persondata til søsterselskaber, som kan være hjemmehørende i lande uden for Det Europæiske Økonomiske Samarbejde (EØS) såsom Dell Inc. i USA, i hvilket tilfælde Dell vil sørge for tilstrækkelig beskyttelse af de pågældende persondata. Dells politik til beskyttelse af privatlivets fred kan ses på Dells websted eller fås ved at kontakte Dell Data Protection Officer, DELL Germany, PO box 20 02 34, 60606 Frankfurt am Main. Kunden giver sit samtykke til, at Kundens persondata må behandles som anført ovenfor.

15. Fortrolighed
Begge parter skal behandle alle oplysninger, som modtages fra den anden part, og som ser ud til at være fortrolige, på samme måde, som den pågældende part generelt set ville behandle egne fortrolige oplysninger, men ikke med mindre end en rimelig grad af omhu.

16. Ophævelse
Ophævelse kan ske fra begge parters side, hvis den anden part (i) væsentligt eller vedholdende misligholder nærværende Betingelser og ikke afhjælper misligholdelsen senest 30 dage efter skriftligt påkrav herom eller (ii) bliver insolvent eller ikke længere er i stand til at opfylde sine gældsforpligtelser. Dell kan ophæve med øjeblikkelig virkning ved skriftlig meddelelse til Kunden, hvis Kunden (i) ikke foretager rettidig betaling eller (ii) misligholder lovgivningens eksportrestriktioner eller Dell har en rimelig mistanke om, at Kunden har gjort dette.
Følgende bestemmelser i nærværende Betingelser skal forblive i kraft efter nærværende Betingelsers ophør eller udløb og skal fortsat være bindende for parterne, deres successorer og rettighedserhververe: Pkt. 4.2, 4.4, 6.1, 6.2, 7.4, 8.7, 8.8, 11, 12, 13, 15 og 18.

17. **Forbrugeres retsstilling**
Ufravigelige retsregler i relevant forbrugerlovgivning påvirkes ikke af disse Betingelser.

18. **Lovvalg og værneting**
Disse Betingelser er underlagt dansk ret, og begge parter accepterer, at de danske domstole har enekompetence. Hvis en domstol måtte finde eventuelle bestemmelser i disse Betingelser ugyldige, påvirkes de resterende ikke heraf. Alle meddelelser skal ske skriftligt og sendes til en juridisk medarbejder hos den anden part på den adresse, der fremgår af fakturaen. Wien-konventionen om aftaler om internationale køb finder ikke anvendelse.

19. **Overdragelse**
Dell kan helt eller delvist overdrage og overføre sine rettigheder og forpligtelser til kompetente tredjemænd, herunder lade sine rettigheder og forpligtelser udføre helt eller delvis af underleverandører. Kunden må kun gøre tilsvarende efter Dells skriftlige samtykke.

20. **Diverse**
Kunden kan se alle Dells retningslinjer, nærmere oplysninger om Produkter og Servicetilbud samt meddelelser på www.dell.dk