Computing at the chair Richter

June 21, 2010

Contents

1 Preliminary remarks

1.1 About this document

This document serves as a short introduction to the computer resources of the chair Richter and gives an overview about existing possibilities. It shall help new group members to get acquainted with our computer system. Complimentary to this document, there is the *Computing policy of the chair Richter* which contains some rules of conduct for a smooth operation within the chair. Up-to-date versions of both documents can be found at *Homepage Richter→Internal*. Some terms in this manual are highlighted by a box. They are explained in more detail in the glossary.

1.2 System administrators

If you have questions or problems concerning computing, don't hesitate to contact your system administrators:
Jan Bundesmann, PHY 4.1.17, Tel. 2016
Tobias Dollinger, PHY 4.1.17, Tel. 2016
Timo Hartmann, PHY 4.1.17, Tel. 2016
Viktor Krckl, PHY 7.1.04, Tel. 3194

2 Introduction

2.1 Computing at the chair Richter

Every member of the group is provided with a desktop PC. The default operating system is Linux. Besides these personal computers, there are some shared devices, such as printers, the Linux compute cluster, a multimedia Windows PC and laptops.
As a matter of principle, Windows is only installed on the laptops and the multimedia PCs (PHY 4.1.29 and PHY 5.1.34).

2.2 User name and passwords

Every user is registered in the computer system of the university of Regensburg with a so-called *NDS-short tag*. This short tag consists of three letters and 5 digits, e.g. dau12345. It also serves as user name for Linux.
Every user has one central password:

NDS password: required for Windows and Linux and some Internet-based services of the university. This password has to be changed every 90 days. It can be changed at:
http://www.uni-regensburg.de/e/r/Benutzer/Allgemein/01504_en.phtml
or, in Linux, by the command `passwd`.

Apart from that, there are some other, less frequently used passwords:

Library password: used to borrow books from our local library of the chair Richter. It can be changed at:
Physik → Mitarbeiter(innen) der Fakultät → Modify your own entry

Samba password This is needed to make in Windows a connection to the Linux-home-directory. It can be changed at:
https://www-cgi.uni-regensburg.de/cgi-bin/RZD/smbpwd.phtml (only in german)

2.3 Internet presence of the chair

2.3.1 Homepage Richter

The chair Richter is represented by a website on the Internet. There, you can find information about current research, publications, seminars and lectures. The homepage can be found at

http://www.physik.uni-regensburg.de/forschung/richter/richter/ .

2.3.2 Library database

The chair has a small library (PHY 4.1.13). Its books are available to any group member. They are managed by a database which can be accessed at *Homepage Richter→Internal*. Please enter books that you checked out or returned!

2.3.3 Mailing lists

Every member of the chair is subscribed to a mailing list: `richter-liste` for non-professors, `richter-profs` for the professors. Via these mailing lists you will receive information such as seminar announcements, important information about the computing system, ...
Everybody from the chair can send a message to these lists. Simply send an email to

`richter-liste@@mailman.uni-regensburg.de`
`richter-profs@@mailman.uni-regensburg.de`

(Note: if you want to reach the professors, don't forget `richter-profs`!)
If you need to send an email to the members of the chair Grifoni, the address is

`qtd@@mailman.uni-regensburg.de`

2.4 Useful links

Besides this document, there is a vast amount of information on the websites of the university. A short list of useful links:

***Homepage Richter→Internal*:**
http://www.physik.uni-regensburg.de/forschung/richter/CMT/pages/internal.html
Computing pages of the department of physics (German only):
http://www.physik.uni-regensburg.de/edv/
Linux pages of the computing center (German only):
http://www.uni-regensburg.de/e/r/Benutzer/Speziell/Linux/index_de.phtml

2.5 Local devices

CDROMs, floppy disks and USB memory sticks are shown on the Desktop when using KDE or Gnome; they are mounted by mouse click. You can find these devices in `/media/`. The exact name differs from computer to computer, so it is best to call `ls /media/` in order to find the exact path.
IMPORTANT: KDE users must make sure that the field *Show device icons* in *Control Center → Desktop → Behaviour → Device Icons* is active.
On the console, mounting is done via `pmount`:

Floppy: `pmount /media/floppy` mounts the floppy at `/media/floppy`. Use `pumount /media/floppy` to unmount the floppy again.

CDROM: `pmount /dev/cdrom` mounts the CDROM device at `/media/cdrom` (Note `dev`↔`media`!). `pumount /media/cdrom` unmounts the device

USB-Sticks: Usually, a USB memory stick can be mounted by calling `pmount /dev/sda1` (or `/dev/sdb1`, ..., depending on the hardware configuration of the machine). However, in this case mounting through KDE or Gnome is the preferred method.

2.6 Software

The Linux system that is in use in the physics department has a large number of programs installed. If you need a certain application, most probably it will already be available to you. However, should you miss a certain, essential application, please contact your system administrators. Manuals for many Linux programs can be read with the command

`man` *program name*

Some packages (e.g. Mathematica) have to be enabled before usage. This is carried out by the command `en` (it gives a list of available packages). A package has to be enabled on every console anew.
Below we give a short overview of available programs. This list is far from being complete and shall only serve as an introduction.

2.6.1 Internet and eMail

The standard Internet browser is `iceweasel` (basically this is the firefox, due to copyright reasons it is called 'iceweasel' in Debian), the standard mailer and PIM (personal information manager) is `groupwise`. Both programs can be started from the terminal (`firefox` bzw. `groupwise`). In KDE and Gnome groupwise can be found in a special folder in the applications' menu called "ReX-Debian".
Groupwise is fully configured for access to the university-postbox. Detailed information about groupwise can be found at (german only):
http://www.uni-regensburg.de/e/r/Benutzer/Allgemein/Email/Groupwise/index_de.phtml
Note: Access to eMail from outside the university is possible via a web-based interface. It can be found at

https://webmail.uni-regensburg.de/

2.6.2 Desktop

The standard desktops are *KDE* and *Gnome*. You can choose your favorite desktop at login.

2.6.3 Programming languages

By default, there are compilers for C (`gcc`, C++ (`g++` and Fortran (`gfortran` and NagWare, to be enabled with `en`).

2.6.4 Scientific applications

Several math programs are available: Maple (`xmaple`), Matlab (`matlab`) and Mathematica (`mathematica`. For data processing and visualization you can use Gnuplot (`gnuplot` or `xmgrace`. Further applications can be found in the package repositories.
Note: Older versions of some packages (e.g. Maple) can be found in `/psi`.

3 Multimedia Windows PC

3.1 General remarks

The multimedia Windows PC in PHY 4.1.29 is available to all members of the chair. Please note the HowTo on the wall! A second and newer Windows PC can be found in the diploma students office (PHY 5.1.34)

The Windows PC can be used, amongst other things, to burn CDs or DVDs (via *Nero*) and to scan pictures. In addition to that, there are several graphics applications, such as *CorelDraw* or *Photoshop*.

3.2 Accessing data from Linux

For data transfer from Linux to Windows, you can use a SSH client.
With *SSH Secure Shell Client* you can login interactively on a Linux PC, with *SSH Secure File Transfer Client* you can copy data from Linux to a Windows directory and vice versa. A connection is established via the menu *File→Connect...*. *Host name* takes a valid computer name (e.g. `pc56936`), *User name* is your user name (NDS short tag). It's advisable to choose your own PC for *Host name*. Finally, you have to enter your NDS password at the *Password* prompt.
Note: These instruction can also be found in the HowTo on the wall by the Windows PC!
Furthermore, the *G:*-drive can be mounted in Linux with the command `rexmountG.py` *`mountpoint`*.

4 Printing

4.1 Printing PostScript files

Under Linux, PostScript (PS) files can be directly printed from the command line by using the command

```
lpr -Pprinter name file.ps
    kprinter Datei.ps
     gtklp Datei.ps
```

In this context, *printer name* is one of the theoretical physics printers. `kprinter/gtklp` provide a graphical interface for processing printing jobs. They offer an easy way to do complex configurations, e.g. duplex-printing, printing on slides... Both programs also can deal witz pdf-files. Our printers are:

pt0, pt2, pt3, pt4: Black-and-white (BW) laser printer, duplex.
Location: PHY 5.1.20. These printers are intended for every-day's use.

clj4: Color laser printer, no duplex.
Location: PHY 5.1.20. This printer is intended for **color** prints. **Please do not send BW prints to this printer!** For BW printing, the color laser jet is disproportionate expensive and yields noticeably worse quality than the specialized BW laser printers.

Thus, if you have a document containing only few color pages, only print the color pages on the color printer, the rest on a BW laser printer.

clj2: Color laser printer, duplex.
Location: printers'-room of the chair Grifoni, close to PHY 3.1.13. These printer can produce double-sided color prints. Again: Printing on these

devices is very expensive and should be limited to what's absolutely necessary.

4.2 Printing from applications

Besides printing from the command line, you can print directly from common applications under Linux. Details are different from program to program. Therefore, we just explain some typical procedures using selected examples:

acroread (Acrobat reader for PDF files): The menu *File→Print* starts the printing process. In the pop-up window, you have to select the printer.

soffice (Star Office)/openoffice: Printing is done with a graphical user interface containing all necessary options.

firefox/mozilla: The menu *File→Print...* starts the printing process. A graphical user interface contains all necessary options.

other programs: Mainly the programs use gtklp or kprinter for the printing job administration, so their use is described above.

In addition, almost all applications can print to a PS file. As explained in the previous section, it can then be printed via `lpr` or `gtklp`

4.3 Printing on transparencies

Transparencies for the different printers can be obtained from Mrs. Reißer (PHY 4.1.34).
As a matter of principle, **only use transparencies suitable for a certain printer!** As an example, you must **not** use inkjet transparencies in a laser jet. Transparencies for copiers are also **not** suitable for laser printers. Furthermore, you must **not** use BW laser transparencies in the color laser jet! If you are in doubt, please contact your system administrators.
For printing on transparencies, make sure that you choose **one-sided** printing (**no duplex**). For the laser printers you have to use the manual paper tray. You can set these options using the graphical user interface of `gtklp`.

5 Storage

5.1 Available storage space

Every user has a personal home directory. The available storage space is approx. 200–400 MB (quota). The command `quota` returns the available disk space. This directory is backed up daily (cf. section 5.3).
In addition to the home directory, there is more storage space:

/data1/*group*/*user name*: Disk space on a hardware RAID (about 800 GB total). Daily back-up (cf. section 5.3).
This directory can be accessed from every Linux PC of the chair Richter.

Large scientific data should be saved here. There is, in principle, no quota, so please use this storage space in a responsible way!

Furthermore, **don't** save the output of ongoing calculations on `data1` since this interrupts the backup process: If files are written to `/data1` during the backup, this will lead to a incomplete backup of your directory. The output of ongoing calculations must be saved on `/richterscratch`.

/richterscratch/*group*/*user name*: Disk space on a hardware RAID (about 2 TB total). No Backup.
This directory is also accessible from every Linux PC of the chair. Temporary data and output of ongoing calculations should be saved here. In principle, data on `/richterscratch` can be deleted within a week after notification.

/misc/condmat/*Gruppe*/*Benutzername*: Storage on a hardware RAID (about 10 TB). No backup. Similar to `richterscratch`

/temp_local/*user name*: Disk space on the local hard disk. No Backup.

group is, for example, `richter`, `fabian`, ..., *user name* the NDS short tag, e.g. `dau12345`.

`G:\` This is the personal diskspace in Windows. Logging in to Windows (not Workstation only), one has automatically the network-drive `G:\`. The available disk-space on this drive is 1 GB. This also has a backup mechanism. Access to `G:\` from linux is established with the command:
`rexmountG.py mountpoint`
where `mountpoint` is an empty directory.

5.2 Private homepages

Every user has the possibility to put his/her private homepage in the WWW. There are two possibilities:

Linux: Login with ssh to `rex1` the homepage has to be located in the direcory `public_html` in the home directory. Don't forget to set read permissions for everyone (use `chmod`). The standard disk space is 50 MB. The URL for user vip12345 is:
http://homepages.uni-regensburg.de/∼vip12345/

Windows: In Windows you have your personal network drive `G:\`. The homepage should be placed in the directory: `G:\www\`. The available disk space in the whole G-drive is 1GB. The URL for user vip12345 is:
http://www-nw.uni-regensburg.de/∼vip12345/

To access `G:\` from Linux, see the description in section 5.1.

Further information can be found at (German only):
http://www.uni-regensburg.de/e/r/Webmaster/01536_de.phtml

5.3 Backup

Both the home directory and `/data1` are backed up daily. Data is stored within a time frame of 4 weeks. This means, that you can restore a file the way it was at any point within the last 4 weeks. However, this also means that you **cannot** restore a file that has been deleted more than 4 weeks ago.
The backup can be retrieved by using the *Tivoli* client:

home directory: Login onto `rex2` via `ssh rex2` and start the backup client via `dsmj -server=titan7`.
Titan7 → File level → t7h → physik → NDS-name
As *Restore destination* you can unfortunately not choose *Original location*. Instead, choose *Following location* and enter the complete path of the destination directory, such as `/home/abc12345/`.

/data1: `ssh pcph051`, `dsmj`.
pc50793 → File level → data50-backup

A short description of `dsmj` can be found at (German only)
http://www.uni-regensburg.de/e/r/Benutzer/Speziell/Linux/01682_de.phtml

6 Computing

6.1 Jobs on desktop PCs

Every user can start computing jobs on his/her desktop PC.
You can also start jobs on other desktop PCs of the chair Richter. However, programs must be started with `nice +19`. In addition only one job per PC is allowed.
A list of desktop PCs can be found at *Homepage Richter→Internal*.
or in the *Ganglia*-overview with some more details.

6.2 Jobs on cluster PCs

The chairs Richter, Schäfer and Braun operate a physics-owned Linux cluster.
A list of the cluster machines can be found at *Homepage Richter→Internal*.
Every member of the chair is authorized to use this cluster for scientific calculations. The number of jobs may not exceed the number of processors on the cluster PC. The list of cluster PCs tells you whether a certain machine is a single- or dual-processor PC.

Note: Further information about *Job submitting* can be found in the *Computing policy of the chair Richter*.

6.2.1 `ssh` without password

To start a compute job on another machine, a user has to log into that computer via `ssh`. In fact, the tools `rsubmit`, `rjobs` and `rkill` do just that, only in an automated fashion.

Normally, you have to enter your password for every `ssh` connection. however, you can avoid that by generating a *SSH key pair*:

```
ssh-keygen -t dsa
```

creates a key pair `id_dsa` (private key) and `id_dsa.pub` (public key) in `.ssh` in the home directory, simply accept the paths that `ssh-keygen` proposes. In any case, do enter a password for the keys.
Next, add the public key `id_dsa.pub` to the file `.ssh/authorized_keys`, for exmple by calling

```
touch .ssh/authorized_keys
cat .ssh/id_dsa.pub >> .ssh/authorized_keys
```

Finally activate the key pair by calling `ssh-add` and entering the password of the keys. `ssh-add` must be called after every reboot.
A detailed explanation (German only) can be found at

http://www.physik.uni-regensburg.de/edv/unix/ux_saccess/ssh.html/

6.3 Athene-HPC-Cluster

The computing center (RZ) of the University of Regensburg provides a Linux-cluster for the whole university.

http://www.uni-regensburg.de/e/r/Benutzer/Speziell/Linux/Cluster/01787_de.phtml

In order to do calculations on this cluster, you have to register an account at the URL written above.
For using the Athene-Cluster, you have to login on its login-node `athene1` by typing `ssh dau12345@athene`. There you won't find the directories known from the physics department `/data1`, `/richterscratch` and `/misc/condmat`. Additionally you have a special home directory within the Athene-cluster. Your normal home directory can be reached as `/lhome/dau12345`, but only on the login-node `athene1`. The jobs are being processed on ather machines and thus should not access `/lhome`. The Athene home directory has a storage limitation of 2GB and a maximum of 10000 files, it is backed up every night. Further storage exists under `/data` (15TB), `/scratch1` (15TB) and `/scratch2` (15TB), which is not backed up.
The Athene-cluster is equipped with a queueing system (PBS). An overview over the current queue is given by `qstat`. Administration tasks on jobs are done by `qalter`.
New jobs can be submitted with the command Neue Jobs k"onnen mit dem Befehl

`qsub` *script file.*

script file is a plain text file containing information about estimated memory usage, duration of the job, etc.

6.3.1 Compiler, MPIs, Libraries

Module Many libraries, especially those needed for MPI, are organized in so-called modules which can be loaded on request. For this purpose there exists a special command `module` (only on `athene1`). Some examples for its usage:

module	avail	(Which modules do exist?)
module	load parastation-gcc	(load a module - here parastation-gcc)
module	list	(Which modules are loaded?)
module	unload parastation-gcc	(removal of a module - here parastation-gcc)

Example - Intel Compiler With the command `module load intel-11.0` the Intel compiler is loaded and the commands `icc`,`icpc` and `ifort` can be used. With `module load impi-intel` the corresponding multi-processor-compiler`mpiicc`, `mpiicpc` and `mpiifort` are loaded.

Math Kernel Library For using Intel-MKL librariessome environment variables have to be set. To do so, there exist some scripts in the path `/opt/intel/mkl/10.1.1.019/tools/environment/`, which contain the appropriate commands for your shell and architecture. The are exectued by the command
`source /opt/intel/mkl/10.1.1.019/tools/environment/skript`.
Your compile-command has to be complemented with: `-lmkl_gf_lp64 -lmkl_sequential -lmkl_core -lpthread -lm`
`(-lmkl_gf_lp64 -lmkl_gnu_thread -lmkl_core -lguide for multithreading)`
Analogous for the Intel compiler::
`-lmkl_intel_lp64 -lmkl_sequential -lmkl_core -lpthread`
`(-lmkl_intel_lp64 -lmkl_intel_thread -lmkl_core -lguide -lpthread for multithreading; to control the number of threads, one can use: export OMP_NUM_THREADS=2)`

ATLAS ATLAS libraries can be found in
`/opt/lapack/atlas-3.8.2/lib`.

6.3.2 Example

Login

```
> ssh athene1
```

Example file "Hello World!"

```
/* hello_mpi.c */
#include <stdio.h>
#include <stdlib.h>
#include <mpi.h>
int main (int argc, char **argv)
{
```

```
    int rank, size;
    MPI_Init (&argc, &argv);                /* starts MPI */
    MPI_Comm_rank (MPI_COMM_WORLD, &rank); /* get current process id */
    MPI_Comm_size (MPI_COMM_WORLD, &size); /* get number of processes */

    printf ( "Hello world from process %d of %d\n" , rank, size);
    MPI_Finalize ();
    return EXIT_SUCCESS;
}
```

Compiling with GNU compiler and Parastation MPI

```
> module load parastation parastation-gcc
> mpicc -o hello_mpi hello_mpi.c
```

Submit-Skript (job.sh)

```
#!/bin/bash
#PBS -l nodes=1:ppn=8
#PBS -l walltime=00:03:00
#PBS -l pmem=1024mb

cd ${PBS_O_WORKDIR}
mpiexec -np 8 ./hello_mpi
```

The first lines after the magic line with the interpreter are obligatory. The first line fixes the number of nodes, the job is running on and how many cores are used on each node. The walltime in the line after defines the maximum amount of time the job can use. After this time, the job is automatically killed. The third "PBS" line gives the amount of memory, each child process uses.
Jobs can't get arbitrarily big. Besides physical limitations (numer of cores and memory) there are some limitations from the queueing-system. Each user can submit a maximum number of 20 jobs at once and job can maximally use 32 nodes (256 cores).

Execution and output

```
> qsub job.sh
12345.athene2
> cat job.sh.o12345
Hello world from process 0 of 8
Hello world from process 1 of 8
...
Hello world from process 6 of 8
Hello world from process 7 of 8
```

The output of the job is written into `skript.o12345` and `skript.e12345` respectively for the error-messages.

6.4 External ressources

Besides the computers at the university of Regensburg, there are also external institutions providing computing time. Among them are:

Leibniz Rechenzentrum (LRZ) München

http://www.lrz-muenchen.de/

John von Neumann Institut für Computing (NIC), FZ Jülich

http://www.fz-juelich.de/nic/

To obtain some computing time on one of the super computers or clusters of these institutions you have to submit a scientific project proposal. Important information and dates can be found on the websites. If you have further questions, please contact Michael Wimmer.
Nach der Anmeldung wird eine Email zugeschickt, sobald der Rechner beim DHCP-Server registriert ist, das dauert etwa 20 Minuten. Die Registrierung gilt immer nur befristet maximal bis zum 15. Februar eines Jahres, kann aber online verl"angert werden. Kabelgebundenes Netzwerk funktioniert dann ohne weitere Schritte, um das Funknetz zu nutzen gibt es zwei verschiedene Mglichkeiten:

- **unifunk1:** Unverschl"usseltes Netz, das nur mit Hilfe eines VPN-Clients benutzt werden kann.
- **802.11i:** WPA/WPA2-verschl"usseltes Netzwerk.

Detaillierte Anweisungen zur Konfiguration des Computers zur Nutzung des Funknetzes finden sich unter *http://www-wlan.uni-regensburg.de/konfiguration.html* und *http://www-wlan.uni-regensburg.de/winxp/nutzung.html*.

7 Private notebooks

It is possible to have internet access with your own notebook at university. But it is necessary to register your notebook.
In order to have LAN- or WLAN-access the notebook has to be registered with the hardware address of your ethernet card. You need to fillout the online form at:
http://register.uni-regensburg.de

After registration you recieve an email as soon as the notebook is registered at the DHCP server. This takes rougly 20 minutes. The registration automatically ends on the next february, 15th. It can be prolonged online. Ethernet then works without further steps, for the wireless network there exist to possibilities:

- **unifunk1:** Unencrypted network which only can be used with the help of a VPN client
- **802.11i:** WPA/WPA2-encrypted network

Detailed instructions for your notebook configuration can be found at: *http://www-wlan.uni-regensburg.d*
and *http://www-wlan.uni-regensburg.de/winxp/nutzung.html*.

8 Glossary

NDS stands for *Novell Directory Service*. It carries out the administration of users at the university of Regensburg.

NAS stands for *Network Attached Storage*. This means storage available on all the PCs in a network.

PostScript is a file format for documents. This kind of file is produced, for example, by LaTeX(there: `dvips`). PostScript files can be printed directly under Linux.

Quota means a per-user limitation of storage space. By that the amount of data is kept reasonable.

Cluster is a pool of computers for the purpose of running calculations on them.

Queueing means that calculations cannot be started directly on a cluster. The queueing system is rather provided with information about memory needs and CPU time of a job. It then decides depending on the load on the cluster when to start a job.

Ganglia is a system for data collection in a computing cluster. In particular it collects computer load and the distribution of the cluster usage between different work groups.

MPI stands for *Message Passing Interface*. It is a library for parallelizing programs. MPI programs must be started with `mpirun`. Further information can be found at
http://www-unix.mcs.anl.gov/mpi/