

NEWSLETTER

December 2006

Welcome to the final DICE newsletter of 2006. We introduce new members of staff, explain more about the Informatics Development Meeting, feature articles on MATLAB and "DIY DICE and AFS" and offer you some hints and tips from your fellow users.

Morna Findlay <morna@inf>

Arrivals and Departures

We were sad to say farewell to Computing Officer Julieta Pineda, and to Computing Support Officers Sarah Reed and Charlie Penman, but we'd like to wish them well in their new pursuits - a new post in MIS, work as a full-time Mum and a life 'down-under', respectively. Staff and students at BP, FH and AT in particular will miss them and will want to thank them for all their hard work over the years.

We have been fortunate in being able to appoint replacements for both CSO positions. Jennifer Oxley joined us on 30th October and Ewan Grant on 13th November.

Ewan graduated from Napier University where he studied Network Computing and went on to complete a postgraduate diploma at Glasgow Caledonian University. Before joining the support team he worked for BT, as a first line support analyst, supporting Edinburgh City Council IT systems.



Jennifer has just graduated this year from Napier University with a degree in Network Computing. She spent several years working in France, followed by a year in sunny Queensland before deciding to return home to study and change career. She enjoys swimming, Tai Chi and is a fan of TV's Lost.



Once they are settled in and have undergone some initial training, Jennifer and Ewan will be staffing the AT and FH support offices on a week by week basis. I hope you will all make them feel very welcome. I'd like to thank everyone for their patience while we have been short-staffed, and to thank Lindsey, Richard, Carol, Ross and Morna who have been so flexible and co-operative in providing support at all four sites over recent months.

Desktop Upgrades

We have been very busy over the summer upgrading staff and lab desktops to FC5. We have upgraded in excess of 800 machines but we still have a few remaining to do and we'll be contacting the users of the few remaining fc3 machines to schedule a suitable date before Christmas for upgrade.

Site	FC5 Desktops
AT	192
BP	172
FH	101
KB	280
TOT	745

Alison Downie <alisond@inf>
User Support Unit

Matlab Usage

Over the last six months we have noticed an increased use of Matlab and its associated Toolboxes.

In particular, the use of multiple Matlab licences by single users due to running Matlab on more than one machine at the same time or on the compute clusters.

Each separate Matlab process requires its own licence and the licences we have are limited. Below is a table of the currently licensed products and the maximum concurrent users (or rather user processes) supported for each:

product	licences
MATLAB	50
MATLAB Compiler	1
Image Processing Toolbox	9
Neural Network Toolbox	5
Optimization Toolbox	1
Signal processing Toolbox	5
Statistics Toolbox	4
Wavelet Toolbox	1

Unfortunately it is costly to purchase more licences and it is hard to restrict licence usage to specific groups or individuals.

In preference to having to introduce a policy on how many licences one individual can use concurrently and to what degree Matlab can be run on the compute clusters we would like to draw attention to and promote the following alternatives.

You can use the Matlab Compiler to compile Matlab programs into standalone applications that do not require any licence to run.

The Matlab Compiler supports the full Matlab language and most of the toolboxes. Of our licensed toolboxes above all are fully supported by the compiler (although not usually the GUI components) except the Neural Network Toolbox which is only minimally supported.

Pre-compiling your Matlab programs for jobs you need to run across multiple machines would

bring a significant reduction in licence usage for small effort.

“R”

You could start developing code in 'R' [<http://www.r-project.org>] instead. This is not a drop in replacement for Matlab but offers similar functionality and is free. While Matlab is largely targeted at mathematical analysis and signal processing 'R' has a very broad coverage of different statistical analysis techniques (linear and nonlinear modelling, classical statistical tests, time-series analysis, classification, clustering) and graphical techniques. Otherwise they are similar tools, both providing a language and environment for data manipulation, calculation and graphical display. 'R' has comprehensive documentation and similar performance to Matlab.

'R' is actually an implementation of 'S' (as is S-PLUS) which is a very high level language and an environment for data analysis and graphics. Although there are important differences much code written for 'S' would run unaltered under 'R'. Certainly anyone using Matlab with the Statistics Toolbox could benefit from looking at 'R' (as could anyone currently using S-PLUS).

Like Matlab, 'R' is highly extensible and has a set of contributed packages, similar in a way to Matlab toolboxes. Some come with the 'R' distribution itself whereas hundreds more can be reached through CRAN (the Comprehensive R-Archive Network) [<http://cran.r-project.org/>].

Some researchers and teaching courses in the School of Informatics are already using 'R' and it is definitely worth consideration.

'R' is available on FC5 DICE machines using the command "R" (uppercase).

Other Alternatives

There are also some other alternatives to Matlab, such as Octave [<http://www.gnu.org/software/octave/>] (a code compatible equivalent) which is also available on FC5 DICE using the command "octave" and SciLab [<http://www.scilab.org/>] which provides some similar functionality.

Either may be suitable in place of Matlab for certain work and may be worth consideration.

Tim Colles <timc@inf>
Research and Teaching Unit

DIY DICE and AFS

One of the initial drawbacks of DIY DICE was that we could not make central NFS mounted filesystems available to it. This was for security reasons. Consequently it was not possible to access your system home directory. However, if you have an AFS home directory it is possible to mount this on a DIY DICE machine in exactly the same way as a normal DICE machine. Any other shared areas currently in NFS can also be made available to DIY DICE machines if moved to AFS. While not officially a production service AFS is now used by all computing staff and a number of research staff and groups. So far AFS has been very stable and we are encouraging other people to try it. If you want an AFS home directory or want to move research space across to AFS you should raise a normal support request.

DIY DICE provides essentially the same benefits as DICE except with more user control over the machine such as package management and root access. In order to encourage people to try it we are making available a number of old GX260 machines. These can be used to play with DIY DICE without having to convert your normal desktop. We will also be able to provide KVM switches for swapping easily between both machines. If you are a member of staff or a research postgraduate student and want to pursue this offer you should raise a normal support request.

Tim Colles <timc@inf>
Research and Teaching Unit

The Informatics Development Meeting

While a large proportion of the CO resources are dedicated to "responsive-mode" problem solving, a significant amount of time is spent on one-off "development projects". In the past, planning for these projects has been somewhat ad-hoc, and prioritisation and accounting for the resources has not been particularly transparent. The recent restructuring of the CO group has introduced a monthly "Development Meeting" to manage these projects in a more organised way, and to clarify their prioritisation and the allocation of resources.

We wanted to establish a lightweight mechanism which would allow us to control and monitor these projects without introducing the overhead of very formal project management procedures. As an interim solution, we have been using a Wiki (1) to record the details of the various projects, and their progress, as well as the usual

meeting minutes and general information. We have gradually been evolving a set of procedures (2) for the development meeting and capturing the requirements for a simple application to record the project information in a more useful way. (3).

Projects which are a direct response to teaching or research requirements now have a ring-fenced resource and are prioritised in consultation with the "customers". Projects involving the computing infrastructure, such as the networking for the new building, or the upgrade to FC5, are prioritised by the Computing Executive Group.

This new process has already highlighted some interesting issues; for example, it seems there was probably a tendency in the past to accept considerably more projects than those which could be reasonably undertaken simultaneously. This could lead to people having too many projects active, and difficulty in planning, or completing projects on time. We are therefore trying to restrict the number of simultaneously active projects; there are currently 22 projects active (4) and a further 46 projects pending (5). These are documented (currently on the Wiki) and open to comments from all staff on their proposed implementation, or priority.

- (1) <https://wiki.inf.ed.ac.uk/DICE/DevelopmentMeeting>
- (2) <https://wiki.inf.ed.ac.uk/DICE/DevelopmentMeetingProcedures>
- (3) <https://wiki.inf.ed.ac.uk/DICE/DevelopmentMeetingSoftwareProject>
- (4) <https://wiki.inf.ed.ac.uk/DICE/DevelopmentMeetingInTerimProjects>
- (5) <https://wiki.inf.ed.ac.uk/DICE/DevelopmentMeetingPendingProjects>

Paul Anderson <dcspaul@inf>

Users Hints and Tips

Help for LaTeX Users

Researchers & students, increase your productivity in paper/thesis writing now that DICE FC5 has both "lyx" and "JabRef"!

Lyx allows quick easy WYSIWYG production of LaTeX documents and is particularly good for complex equations.

JabRef is a GUI reference manager for BiBTeX (LikeEndnote in windows) which works with lyx. JabRef built in querying and importing from cite-seer, IEEE online and pubmed saves time over using the web.

Tim Hospedales <t.hospedales@ed>

I Can't Find a Machine to Log Into

At busy times the computer labs do get crowded. The machines in the Informatics computer labs are available on a "first come first served" basis. However, this means that a machine is only yours while you are working at it. Please do **not** go off to a lecture, lunch, etc., leaving the screen locked.

If there are no spare machines, and you find a machine that has been locked for 15 minutes or more, it is legitimate practice to end the login, so that you can use the machine yourself. In order to find out how long it has been locked, press the 'ctrl', 'alt' and 'f2' together. This will provide you with a text login prompt. Log in as normal and type:

```
who
```

to see who is using the machine. Then type the command

```
ps aux | grep lock  
(or for kde screen savers).
```

```
ps aux | grep *.kss
```

This will then show you how long the lock has been running. Now log out by typing 'exit' or pressing 'ctrl' 'd', then press 'ctrl','alt' and 'f7' together to get back to the locked screen.

If the user has been logged in for 15 minutes or more, you can press the 'ctrl', 'alt' and 'back-space' keys together to end the locked login session.

*Support FAQ entry recommended to users by:
Andrew Emerson <a.o.emerson@sms>*

Copying Text From PDF files

Sometimes when copying text from PDF files you end up with non-standard characters in the text, which can give problems with programs that don't properly support unicode. For example, if I copy & paste abstracts from PDF papers into my bibliographic database, non standard characters that appear cause it to not save properly. It can be hard to spot these characters to change them, so I use the "iconv" tool.

```
"iconv -t ISO-8859-15 temp.txt"
```

This tries to convert the (presumably unicode) text in temp.txt and send it to stdout, but it will exit when it hits an odd character in the input stream, so you can identify it and change it.

Iain Wallace <i.a.wallace@sms>

Sharing Resources¹

I often see a lot of computers with 'experiment running, do not log out' on them, frequently left unattended for days at a time.

While some of these are clearly just lazy final years/Msc students [user's own words! - Ed] there is no need for this at all. All the machines come with '**screen**' and so processes can be left running remotely without the need to maintain a connection. For more information the man pages are very useful, but here is the short version:

- (1) Type "screen" (this detaches your session from the remote connection and runs it locally instead)
- (2) Set your program running
- (3) Close the terminal you logged in with, log out and go home. Then some time later..
- (4) ssh back to the computer you were using and type "screen -rd" (this restores the remote link to the computer, where your process has been running in the background.

NOTE: if you are using a shared computer, it is considered nice to use the nice command to share the resources with anyone else using the computer. When no-one else is using it, your job will be at 100%, and when other people need to use it your job will slow down, giving them some CPU cycles to play with. Again the man page for nice tells you more, but simply putting nice in front of the command is often enough.

Craig Tongue <s0348924@sms>

¹ See also our Good Practice Guide at http://www.inf.ed.ac.uk/systems/good_practice.html [Ed]