

---

## *An Overview of the Computing Facilities at DCC-IMECC-UNICAMP*

*Paulo Lício de Geus* (paulo@dcc.unicamp.br)

Departamento de Ciência da Computação  
IMECC - UNICAMP  
Campinas SP

---

### **Abstract**

The Department's network is presented, showing the lay-out and relevant machine configurations; connectivity problems and solutions are also treated. Stress is given to the solutions devised to allow a single, easily-manageable environment, which provides a friendly environment for users and a tolerable load to (part-time) managers.

### **Overview of the Network**

The Department's network is based on TCP/IP, and most Unix machines run either SunOS or Solaris. Initially a single computing environment, the Department is now expanding into separate research groups and setting up separate sub-domains. Currently there are two separate research Laboratories: LAC (Laboratory of Algorithms and Combinatorics) and AHAND (Project A\_Hand). The A-Hand lab is remotely connected through PPP at 14400 bps whereas LAC is directly connected to the main backbone at 10 Mbps. Figure 1 depicts the whole network.

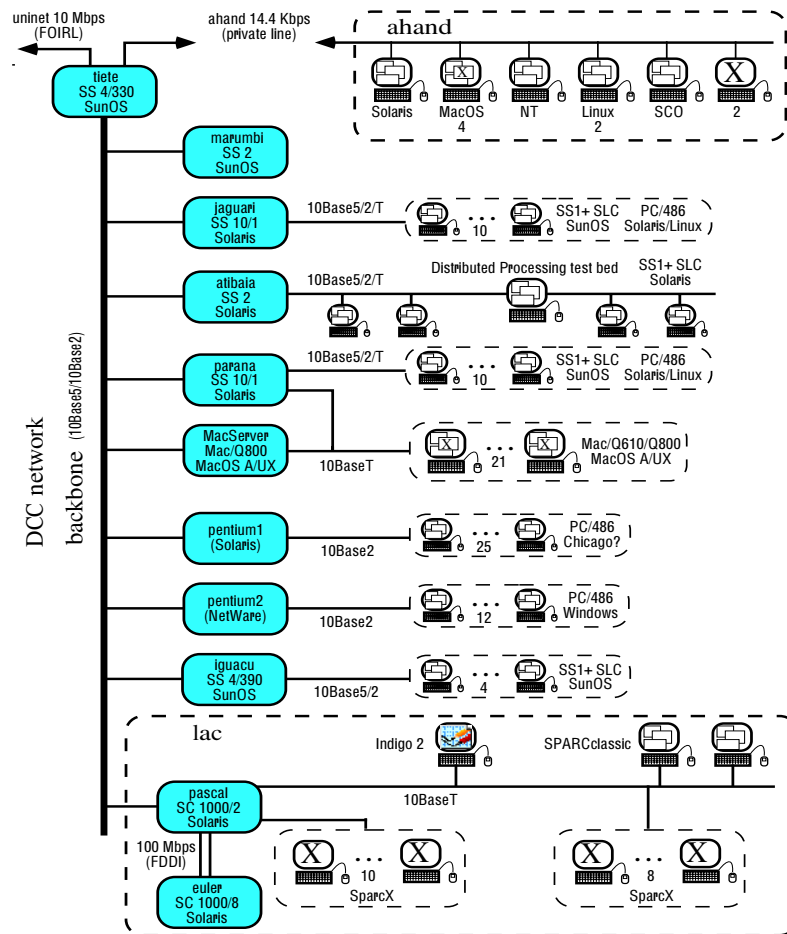
For historical reasons, the Department has not had a significant number of PC's till a month ago, and connectivity (PC-wise) is now becoming a problem. A Novell network is being integrated into the main network, as well as an assortment of Solaris-x86, Linux, Net-BSD and Windows/Chicago systems.

At the time of writing, the Macintosh network is mostly functional, providing reliable storage on the AppleShare server. CAP is at an experimental phase yet and the single homedir paradigm (hosted on Unix) is partly functional. Users overcome unreliable service using ftp. The most desirable feature of these consoles, apart from the native O.S., is fully functional: MacX allows X-Windows access to Unix servers, including services such as SUN's AnswerBook (despite some inconveniences).

## **System Management Approach**

DCC's network (DCCnet for short) is managed with the following goals in mind, in order of precedence:

- simplicity of management is top priority due to the unavailability of dedicated personnel. Some key items to achieve simple management are extensive use of NFS, NIS, automount, network distribution of configuration files, keeping /usr standard (only bug fixes are added to the partition after installation). Upgrading system software is much easier, since no large backups are needed. Only /etc and /var require caring for. Sometimes repartitioning will force backups, of course.
- consistent user view of filesystems throughout the network (NFS, automount).
- up to now, recoverability could not be enforced, due to shortage of required hardware. Newly received equipment will allow for a more reliable overall system.
- particular system usage conditions have spared us from a more secure policy. DCCnet and Uninet (campus network) have suffered a little hacking in the past, but recently hackers have managed to sniff packets on Uninet, thanks to improperly managed gateways; they now use scripts and packages with all the needed tools to exploit the latest security holes. New concern campus-wide, with dedicated support from the local Computer Centre. We are considering setting up a firewall.



## **Heterogeneous Systems**

Heterogeneity has not been a problem till now on DCCnet, since all Unix systems are SPARC-based. Problems started to appear about a year ago with the pilot Solaris system (a SS1+). The lack of both a bundled compiler in the system and resources in the Department have made it impossible to recompile lots of PD packages without extensive porting. Solaris 2.3 solves most of the problem with the enhanced compatibility with SunOS binaries. Nonetheless, an approach for dealing with different binaries has already been devised and is explained in the next Section.

DCCnet has little third-party commercial software. The main problem is going to be the high prices charged for the Solaris C compiler, i.e. one license for each concurrent user (gcc does not compile most PD packages); resources for software acquisition are traditionally harder to obtain. This helps disrupt a consistent environment for users.

## **Basic Network Services**

### **Network Hardware**

DCCnet has been using twisted pair cabling for over two years due to the advantages regarding the integrity of the network. These advantages are apparent when comparing with thin cabling; with thick cabling the advantages fall on price, since most hardware come with built-in TP transceivers, whereas thick transceivers are always extra. Hubs are very cheap at \$150 per 9-port unit. Our experience tells us that SNMP support is not needed on the hubs, and that makes the hardware several times cheaper.

The 10 Mbps campus link is optical and has not presented problems over a 4-year period, with only a few units failing in the campus despite thunderstorms. All routing is done by general-purpose machines. This contributes to increase eventual lock-ups, since dependency on other machines is high, but so far the cost factor has prevailed.

### **Naming Services**

Machine names are internally resolved by NIS. All campus and world machines are resolved by DNS. NIS is extensively used to propagate configuration and system data across all machines. NFS bounds are set with netgroups, account changes take seconds to validate across all machines, and automount maps are easily changed and propagated (making disc management a lot easier). Perhaps the most common

management task on DCCnet is finding new disc space for overgrown hierarchies, and changes are easily accomplished by the following "script" (notice that everything is done on-line except for the affected hierarchy):

- rsh AllHosts /usr/etc/umount net\_hierarchy
- exportfs -u local\_hierarchy; mv local\_hierarchy local\_hierarchy\_disabled
- change automount maps and propagate
- mount new\_host:/new\_place /mnt
- sync; dump 0bf 1000 - local\_hierarchy | (cd /mnt; restore rbf 1000 -)
- correct residual pathnames on new\_host; exportfs /new\_place

### File Service

Standard NFS is used throughout, coupled with standard automount and NIS to propagate maps. Since we already try to isolate system from third-party software, another place to deal with is /usr/local/bin, usually crowded by thousands of files that lead to file name conflicts and slowness. Our approach breaks /usr/local into several software category hierarchies with corresponding ./bin, ./lib ./etc directories, such as:

- /n/net -> all networking software such as mosaic/www, gopher, ftp, DNS utilities
- /n/dtp -> holds FrameMaker, TeX, dvi and PS previewers etc
- /n/draw -> holds AutoCAD, xfig and related packages
- /n/gnu -> holds all GNU software; users decide which version to use via PATH.
- /n/image -> holds khoros, image database etc
- /n/X -> holds all PD X material

/n can also be accessed as /network, for friendliness. Locally, all hierarchies are stored under filesystems whose names help configure a network backup scheme, such as //repXX, //homeXX and //projXX. /l stands for local filesystems, rep for general-purpose repositories (such as //rep01/dtp, //rep05/gnu-), home for homedirs of categories of users (accessible as /h/staff, /home/phd, /h/grad-), and proj for project hierarchies (accessible as /p/dicio, /proj/vid-).

To transparently deal with different binaries, DCCnet automounts the correct NFS hierarchy. Each machine is defined a shell variable that automount can use, such as OS. The host of the given hierarchy holds all OS versions of that hierarchy:

```
hierarchy_host$ ls -F //rep00/gnu
solaris2/  sunos4/  linux/   solarisX86/  irix/
```

```
hierarchy_host$ ls -F /l/rep00/gnu/solaris2
bin/   etc/   lisp/  info/  and so on...
```

Disc space taken by identical files can be saved by a script after the installation of a package (for all OS's). The script descends and compares files in both hierarchies and substitutes a link for the original file. Hard instead of soft links will also save inodes.

### **Backups**

We at DCCnet have developed a backup package to take care of all our filesystems (currently several tens of GB). It is mostly written in shell scripts and makes use of standard dump/restore for actual file storage. Its development was triggered by the need to back up several GB on 150 MB tape drives, and the very expensive prices of DC-6150 tapes in the local market.

CBackup was developed to use tapes more effectively and automate the process as much as possible. CBackup pipes dump and restore to a C program that takes care of tape lengths and asks for substitution whenever needed. The data stream can also be compressed before getting into tape. DCCnet survived a lot more time with the same set of tapes, while backing up a lot more filesystems. Nowadays DCCnet has standardized on DAT drives, each tape storing typically 4 GB of data (data compression by tape drive hardware). CBackup manages all dumps onto a single set of tapes but is still hard to configure. The next version will make use of a high level backup language to make the process much more friendly.

Printing has been done with standard BSD, but we need accounting on pages printed by each user. The SPARCprinter software has been hacked to log pages and toner coverage, but surely it will break when the new version of NeWSprint is installed. For normal PostScript printers the only accountable activity is the number of pages printed. A small utility can read the internal register and help a package such as PLP do the job.

### **Communication Services**

DCCnet uses standard sendmail from SunOS for external delivery. Due to several problems in the past we now decided to Eric Allman's distribution. Since sendmail is one of the largest security holes in Unix today (every other week a new hole is found:-), at least one can have a quick fix and keep the network exposed to a minimum.

Some normally found services are also run, such as news, WWW. Anonymous ftp is being deployed using wuarchive's daemon. As for user connections via modem,

DCCnet has only a handful of phone lines and no solution can be seen in a short term basis. The current PPP link to Projeto A\_Hand uses a PD implementation which is not stable. The built-in Solaris 2.3 version is going to be used shortly.

## **"Fancy" Filesystem Features**

Little has been possible so far concerning load balancing, recoverability and caching. Load balancing assumes availability of disc space for local executable copies. Recoverability depends on several other factors such as routing, NIS service, keeping filesystems local etc. All this requires extra hardware which only now is becoming available.

## **Machine Configuration and Installation**

New machines are installed not from CDROM, but by copying the system software from another machine. This has the advantage of time: a dump/restore pipe, even over the network, is much faster than a full install from CDROM, and also saves a lot of time applying OS patches. The disadvantage is that any filesystem problem might degenerate systems derived from the affected one. A good idea is to keep a disc containing the operating system as it should be on all related machines, but untouched otherwise. For security's sake, it should be kept unmounted.

## **References**

- [1] SunOS 4.1 System and Network Administration
- [2] Solaris 2.2 Administrator's Guide
- [3] Geus, P. L., Administração do Sistema Unix 2, Cursos Unix - Projeto A\_Hand, Campinas SP, March 1994
- [4] sun-managers mailing list.

