

6. File copying, conversion and archiving

(This section last updated August 2004)

This section covers procedures for performing various types of file copying and conversion on our GNU/Linux systems. Major topics include:

- 6.1 Soundfile formats
- 6.2 Soundfile conversion operations
 - (Changing the format, number of channels, sampling rate or bit depth of a soundfile)
- 6.3 Remote logins to ECMC computers : ssh
- 6.4 Copying files to and from remote systems
- 6.5 Compressing files
 - 6.5.1 The Unix *tar* program and the ECMC *tarsf* utility
 - 6.5.2 Compressing files with *gzip* or *bzip2*

These topics may sound rather dull, and indeed they often are, but at times they will become vital to your work.

6.1. Soundfile formats

As noted in section 4.1, the *format* of a soundfile encompasses several elements that determine and how the samples represent sound and how they are written to a file.¹ Different computer and digital audio systems employ different formats, which generally are not compatible.

On **Macintosh** systems soundfiles generally are written in Apple's *AIFF* format. *AIFF-C* is a supposedly downwardly compatible extension to the *AIFF*, but some Macintosh audio applications cannot read *AIFF-C* soundfiles and it is best avoided. On **Windows** systems, Microsoft's *WAVE* format (sometimes also called *RIFF*) is the standard, and many utilities to play soundfiles, such as *WinAmp* and *Media Player*, cannot play soundfiles in *AIFF* or other formats. However, on both Macintosh and Windows platforms, major soundfile editor, mixing and sequencing applications (many of which run on both of these platforms) such as *Cubase* can import, read and play soundfiles in both of these major formats, although they may only be able to write to one of these formats.

There is some ambiguity in, and occasionally resulting problems in, certain aspects of the *WAVE* format specifications, particularly with respect to multichannel soundfiles and word sizes greater than 16 bits. *WAVE* format likely will be replaced — gradually, over several years — with Microsoft's new *WAVE-EX* format, which eliminates ambiguity regarding how 24 bit and 32 bit samples are written, defines loud-speaker locations and their mapping to audio channels for multichannel soundfiles, can handle 64 bit word sizes, and includes other updates to standard *WAVE* format. Currently, however, there are relatively few audio applications that can deal with *WAVE-EX* format, and for our purposes standard *WAVE* format generally remains quite serviceable so long as one is aware of its limitations.

On **GNU/Linux** systems standard *WAVE* format is the norm. Some Linux audio programs and applications can read (and in certain cases, as an option, also write) *AIFF* format as well, but others, including almost all Linux CD burning applications, require *WAVE* format. A *.wav* filename extension typically is appended to the names of *WAVE* format soundfiles on Linux systems. With most Linux audio applications this is not a requirement, as it is with many Windows applications. However, it is still a good practice, enabling us to recognize the file as a *WAVE* soundfile at a glance, and it is particularly important if you might someday use the soundfile on a Windows system. The extension *.aif* (or, occasionally, *.aiff*) is customarily appended to the names of *AIFF* format soundfiles. The names of Linux soundfiles should not include any blank spaces, nor any other characters (such as "*" or "#") that have a special meaning to Unix shells.

Soundfile headers

¹ for example, the order of the bytes, and whether or not the samples of stereo files are interleaved, or instead are written instead in two separate, contiguous blocks

Within most types of soundfile formats the sample data is preceded by a short **header**, which contains information about the sample data representation. This header information specifies the sampling rate, the number of channels, the number of samples in the soundfile, the duration of the soundfile (the number of samples divided by the sampling rate), and sometimes additional characteristics (such as read/write permissions on Unix systems). Most music applications (e.g. programs designed for playing, editing and mixing soundfiles) read this header so that they can process the samples correctly.

On ECMC Linux systems we can display a summary of the header data for user soundfiles with the *sfinfo* command (which can be abbreviated *si*), and for *sflib* soundfiles with the *sflibinfo* command (which can be abbreviated *sfl*). See the man page for *sfinfo* for information on both of these commands. On the ECMCLinux systems a utility called *sfcheck* will check soundfile headers and report on their formats, as well as identifying any problematic soundfiles that may require fixing. On **Windows** systems clicking on the *File>Properties* tab for a selected soundfile will display basic header information.

The most important information contained within a soundfile header includes:

- *number of channels*

With certain music applications on our Linux, Macintosh and Windows systems, it is possible to create soundfiles with 4, 6, 8 or even more channels. However, it obviously will do you little good to create an 8 channel piece on *masking* when only four full bandwidth loudspeakers are available in room 52.

- *sampling rate*

Common sampling rates in use in professional audio today include 44100 ("CD quality"), 48000 and 96000, which seems likely to become the eventual standard. 88200 and 192000 sampling rates are employed occasionally in professional audio, although almost never in the ECMC studios. For maximum audio quality use a sampling rate of 96000. For maximum portability, or if the master playback version of your composition will be audio compact discs, use 44100.

- *bit depth (word size)*

Common word sizes in use today include 16 bit "short" integers ("CD quality" again), 24 bit integers and 32 bit floats. With sampling rates of 44.1kHz and 48kHz, 16 bits are most common. When 96kHz is used, the word size generally will be 24 bit "ints," ("9624") or occasionally 32 bit floats. However, most sound cards cannot play floating point samples.

16 bit resolution provides a theoretical dynamic range of 96 dB. However, the actual dynamic range of 16 bit systems generally is somewhat lower —often between 85 and 90 dB —due to limitations in the analog circuitry. 24 bit systems typically increase this dynamic range to somewhere between about 104 and a theoretical maximum of about 140 dB, depending upon the quality of the analog circuitry and the sampling rate. This expanded dynamic range is most apparent in the reduced noise floor during soft passages and in the "long, smooth" decay of sounds.

6.2. Soundfile conversion operations

Sometimes it is necessary to convert a soundfile from one format to another. The most frequent types of conversion operations required include

- (1) changing the system format of the soundfile (e.g. from AIFF to WAVE or vice versa)
- (2) changing the number of channels from stereo to mono or from mono to stereo, or else from some multichannel format to stereo
- (3) changing the sampling rate
- (4) changing the bit depth (e.g. from 24 bits to 16)
- (4) compressing the soundfile to reduce its size

Some soundfile editing and multipurpose utility applications or programs that can perform several of these operations simultaneously (e.g. converting a 44.1 k stereo WAVE soundfile to a 48 k mono AIFF version). The GNU/Linux soundfile editor *sweep* is an example.²

However, there is no single application that can perform all of the five types of operations above. And often it is quicker to use smaller *special purpose* utilities that perform only one or two of these conversion operations. These special purpose utilities generally load quickly, are easier to use and run more quickly, and can be used in sequence if necessary.

6.2.1. Utilities to change the format of a soundfile: *cpsf.wav* and *cpsf.aif*

Occasionally you may find it necessary to convert soundfiles between AIFF and WAVE format, or vice versa. Some soundfile editors can accomplish this, but the quickest way to perform such conversions on *madking* is to use the ECMC utilities

cpsf.wav —copies an AIFF (or SND) format input soundfile to a WAVE format output soundfile and

cpsf.aif —copies a WAVE (or SND) format input soundfile to an AIFF format output soundfile.

Example: To make WAVE format copies of two AIFF soundfiles, type

```
cpsf.wav inputfile1 inputfile2
```

Example: To make WAVE format copies of three AIFF soundfiles, type

```
cpsf.wav inputfile1 inputfile2 inputfile3
```

For additional options and more details and examples consult the *man* page for *cpsf.aif* or *cpsf.wav*.

6.2.2. Utilities for changing the number of channels in a soundfile

☞ Stereo to mono conversion

A few music applications, including *rt* and some other mixing applications as well as most cd rippers, do not provide the user with a choice of mono or stereo output, but rather always write their output samples into stereo soundfiles. If we have a two channel soundfile that is *not* true stereo —for example, if one of the channels is silent, or if the two channels contain identical sample data (sometimes called "split-mono") — we should convert the soundfile from stereo to mono. Beware, however, that if there are significant phase differences between the signal on the two stereo channels, you may not be happy with the mono folddown.

Stereo soundfiles are converted to mono by adding the left and right channel values for each sample, multiplying the sum by a *gain* factor, then writing the result as the output sample value. By default, most programs that perform stereo-to-mono conversion introduce considerable amplitude attenuation — typically a *gain* factor of .5 (-6 dB, reducing each left and right channel input sample by half) —to guard against clipping. This works fine if both channels of the source stereo soundfile are near maxamp at approximately the same point. Unfortunately, this often is not the case, and stereo-to-mono conversion may introduce considerably more amplitude attenuation than we would wish. Some programs allow us to adjust the *gain* factor, others do not.

Before performing stereo-to-mono conversion, you should know the approximate maximum amplitude of the input stereo soundfile and, after conversion, compare this value with the peak value of the mono output soundfile. To find out the peak amplitude of one or more soundfiles you can use the utility *sfpeak*:

```
sfpeak soundfile1 [soundfile2 soundfileN]
```

Some soundfile editors can "bounce down" the two channels of a stereo input soundfile to a mono output. The quickest and simplest way to perform this operation, however, is to use the ECMC *bounce* script, discussed earlier (section 4.7) in this *Users' Guide*. This utility not only "bounces" down (mixes) the two channels of a stereo input soundfile to a new mono output soundfile, but also normalizes the amplitude of the output soundfile to a peak value of about 98 % of maxamp, or a 16 bit integer value of 32000.³ You

² Another example is the command line utility *sox*, a self-proclaimed Swiss army knife both for performing a wide variety of soundfile conversion operations and for applying various types of effects, such as writing a soundfile backwards or adding echos. However, *sox* is an aging and currently unsupported program that does not support 96kHz or other higher sampling rates, and I no longer recommend it.

³ Occasionally, normalizing samples to full maxamp (a 16 bit integer value of 32767 or floating point value of 1.) can result in slight roundoff errors, and can make life more difficult for some spectral analysis/resynthesis

can normalize to some other level if you prefer. The usage syntax for *bounce* is simple:

```
bounce inputsoundfile outputsoundfile [output_amp]
```

The optional *output_amp* argument is required only if you wish to scale the mono output samples to some peak value other than the default.

Example: The command

```
bounce myvoice.wav myvoice.mono.wav
```

will create a mono soundfile called *myvoice.mono*, with a peak amplitude about 98 % of maxamp, from the source stereo input soundfile *myvoice.wav*. If, after playing *myvoice.mono.wav*, you are happy with it, and have no further use for the original stereo soundfile, you should delete it, possibly by typing

```
rmsf myvoice.wav
```

in a shell window, or perhaps, instead, overwriting it:

```
mv sf myvoice.mono.wav myvoice.wav
```

or else by dragging an icon for *myvoice.wav* to the trash bin.

☞ *Mono to two channel conversion*

Mono-to-"stereo" format conversion is performed infrequently, since it simply copies the input channel identically to both of the stereo outputs, with no left-right stereo distribution of sounds. (Note that converting a stereo soundfile to mono, then converting this mono file to stereo, will not restore the original; all left-right stereo imaging will be lost.)

However, there are a few occasions when it is necessary to create two channel "split-mono" soundfiles. If we want to record a mono soundfile to an audio cd or to a DAT tape, we first must create a two channel version of the soundfile. One way to accomplish this is with the *sweep* soundfile editor. Open the mono soundfile in *sweep*, select *Save as* from the *File* menu, and within the *Channels* box of *Save* dialog window change *Mono* to *Stereo*. Another, quicker way is to use the ECMC utility *cpsf.1to2*. See the *man* page for *cpsf.1to2* for usage details.

☞ *Extracting the individual channels of a stereo or multichannel soundfile*

Occasionally it is necessary to extract the individual channels of a stereo or multichannel soundfile into individual monophonic soundfiles. One example would be if we wish to burn a four channel soundfile created on *madking* to a DVD-A disc, using the *discWelder chrome* application on *gesualdo*. *discWelder chrome* requires mono or stereo input soundfiles, even for multichannel output, to assure that the input channels are mapped correctly to the output channels. Therefore, before we can import our four channel piece into *chrome* for burning, we must convert it into four mono soundfiles.⁴ The simplest way to accomplish this is to use the ECMC utility *splitchans*. Consult the *man* page for *splitchans* for usage information.

6.2.3. Changing the sampling rate and/or bit depth of a soundfile

Sample rate conversion is performed by resampling the input sound. Resampled soundfiles will differ —usually slightly, but occasionally significantly—in peak amplitude from the original. Generally, *downsampling* (resampling to a lower output sampling rate, e.g. from 96 kHz to 44.1 or 48 kHz) results in slight attenuation. More importantly, downsampling almost always will result in *some* signal degradation. (Some of the jagged edges of the input waveform will be missed by the resampling operation.) The resulting loss in signal resolution and audio quality—which may be heard as "graininess," harshness, or loss in crispness and clarity—is most often apparent in soundfiles that contain significant high frequency components, rapid attack transients, or complex textures, when the new sampling rate is half, or less, that of the original, and, with lower quality resampling algorithms, when there is a complex, non-integer ratio between the two sampling rates. *Upsampling*—increasing the audio sampling rate, for example from 44.1 kHz to 96kHz—will not result in better signal quality, but will merely "translate" the original signal to a higher numerical representation.

The same remarks apply to increasing or decreasing the *word size* of a soundfile or audio stream. Converting a 16 bit soundfile to 24 bits will not result in better audio quality. Decreasing a 24 or 32 bit soundfile to 16 bit representation will result in some degradation. Increasing the word size from 16 to 24

programs.

⁴ We probably will need to add MLP encoding to these mono soundfiles as well before we import them into *chrome* for burning.

bits will not improve signal quality, but simply will make the soundfile usable within a project where all of the other soundfiles are at 24 bit resolution.

Sample rate and word size conversion can be accomplished independently, but they often are performed together between 16 bit resolution at 44.1k or 48k and 96k 24 bit resolution. How appreciable — and, more importantly, how aurally perceptible — signal degradation will be when we downsample and/or reduce the word size will depend upon several factors, including the complexity of the source audio waveform. The most important determinant, however, is the quality of the resampling or bit conversion algorithm.

Whenever downsampling or bit depth reduction are performed, high quality dithering should be applied whenever possible to minimize quantization "noise" (roundoff error) and waveform staircasing within the least significant bits. Elementary dithering algorithms simply add one bit of white noise to the signal. More sophisticated dithering algorithms offer a choice of filtering and/or more complex, weighted noise generating procedures to concentrate the dither "noise" within frequency bands where it is less noticeable perceptually.

On Linux systems there currently is no handy utility for performing high quality sample rate and bit depth conversion. It is possible to use *ardour* to perform these operations, with high quality dithering. However, it likely will take you several minutes to create an *ardour* "project" folder, import one or more source soundfile, perform the conversion(s) and saving the output(s).

For situations where the highest quality sample rate and bit depth conversion is not required, *sweep* can be used. To change the sampling rate:

- open the source soundfile in *sweep*,
- select *Sample -> Resample*, set the new sampling rate in the *Custom* box and click on Resample box

To change the *bit depth*:

- select *File -> Save as* and choose a new name for the output soundfile
- in the *Save* dialog window, change the word size in the *Encoding* box and save the file

Note however, that no dithering is applied during these conversions.

The best utility application currently available in the studios for performing sample rate and word size conversion is a Windows application called *resample* that is available on *gesualdo*. The best way to learn how to use *resample* is to consult the graphical *help* tutorial available within the application. Be sure to read the section on applying dithering, and dithering options, near the end of the *help* document.

6.3. Remote logins to ECMC computer: ssh

All of the ECMC computer systems are connected by Ethernet hardware and networking software, with each individual computer comprising a *node* on this *Local Area Network (LAN)*. Connections between the ECMC Unix-based (Linux and Macintosh systems) provide for

- remote logins (e.g. logging on to *madking* from a shell window on *sound*),
- executing commands remotely on these machines (e.g. listing our Unix files or our soundfiles on *madking* while logged on to *sound* or *wozzeck*); and
- copying files back and forth between any of these machines

Connections between the ECMC Windows systems and any of our Linux or Macintosh systems provide a similar but more limited range of services, generally confined to remote logins and file copying —because the ECMC Windows systems can only act as *clients* (which initiate requests for services), rather than as network *servers* (also called *hosts*), which actually do the work).

With appropriate software, most of the network operations that you can perform between the ECMC Linux and Macintosh systems can also be performed from a home system, or from most any networked computer, whether it be a Windows PC, a Mac, or a Linux or Unix system.

However, our Windows systems cannot be accessed remotely, since they lack *server* software. All remote connections from the ECMC Windows systems must be made while you are logged onto the Windows machine.

Security issues

In the past some of you may have employed networking client applications such as *telnet* to log on remotely from one computer system to another, and applications such as *ftp* to transfer files between these computer systems. However, both of these applications are insecure. They are no longer installed in the server software of our Linux and Macintosh systems, and we have closed their ports. (It may still be possible —although certainly NOT recommended —to use *telnet* and *ftp* client software to get OUT of any ECMC systems to a host that still supports these creaking protocols, but you cannot use them to get IN to an ECMC system.)

It is a new and sometimes nasty world out there on the internet, with computer networks beginning to resemble walled medieval cities, and system firwalls have become a paramount concern at the ECMC as just about every place else. In the spring of 2001 a cracker from Korea, exploiting an *ftp* security hole on our SGI systems, installed "backdoor root kit programs that gave him super-user privileges on these systems. He then installed hot copies of the not-yet released Windows 2000, which were dutifully downloaded by others from Korea, slowing our flagship SGI machine to a crawl. System restoration was agonizing and required more than three weeks. All ECMC users must take security issues very seriously. Make certain your password is NOT crackable; and whenever you have a choice, always use a "secure" networking application, even if it is somewhat slower or more cumbersome than an "insecure" alternative such as *ftp*.

All connections into and between ECMC systems are handled by an authentication and protocol system called *ssh* ("secure shell"), which comes in several favors. *ssh* actually is proprietary commercial software developed by a Finnish company. So on the ECMC Linux systems we are running a widely used open source derivative known as "Open SSH." Unlike *telnet* and *ftp*, *ssh* encrypts data sent over the net. And, when fully implemented, *ssh* employs a system of "public" and "private" keys to offer two layers of protection for all data you send over the internet against packet "sniffing," "man in the middle" and either types of attacks by crackers. The OpenSSH web site at

<http://www.openssh.org> (also available at <http://www.openssh.com>)

includes links to compatible downloadable client software for Windows and Mac as well as Unix/Linux systems. Mac OSX systems come bundled with *ssh* software. To connect to an ECMC Linux or Macintosh computer from your home system, you will need to have *ssh* client software. For Windows, systems, we currently recommend the *Putty* SSH client application and the *WinSCP* file transfer application, both of which are installed on the ECMC Windows systems.

On Linux and Mac systems the command line *ssh* syntax is

```
ssh [options] host [command]
```

To remotely log on to madking from one of our other ECMC Linux or Macintosh systems, type:

```
ssh madking.esm.rochester.edu
```

The first time you try to connect from one machine to another with *ssh* the program will engage you in a brief dialogue, and you will see something like:

```
The authenticity of host "machinename" can't be established.
```

```
Key fingerprint is 1024
```

```
Blah blah blah ... gibberish numbers ...
```

```
Are you sure you want to continue connectiong (yes/no?)
```

Since this is the first time you are trying to connect to *madking* from the machine from which you submitted this command, *ssh* doesn't have any information on which to authenticate your request. Type "yes" to the query above, and the *ssh* software will add a line with tons of numbers in the file *known_hosts* in a directory called either *.ssh* (or perhaps *.ssh2*) on your local machine. The stream of numbers within this file contains information about your account on *madking*. You won't be bothered again by this detour in future attempts to connect to *madking* from this machine. Next you will be prompted for your password on *madking*. (This will happen every time you use *ssh* unless you take steps to automate encrypted password forwarding.) Eventually you should find yourself logged on to your home directory on *madking*. Celebrate by listing the files in your home Unix directory on *madking*.

If you wish to remotely log on to a serve on which your user ID (login name) differs from your UID the local machine, use the *-l* flag followed by your UID on the host. In other words, if I have an account with a UID of *mighty mouse* on U of R machine *troi*, I would type

```
ssh -l mighty mouse troi.cc.rochester.edu
```

To execute a command, rather than log on, to a remote host, include the command at the end of the call to *ssh*:

```
ssh sound.esm.rochester.edu ls -l
```

(or *ssh -l youruserID sound.esm.rochester.edu ls -l* if your login name differs on this machine)

This will display a long listing of the files in your home Unix directory on machine *sound*. To execute two or more commands, separate the commands with a semicolon (;), just as with a local shell:

```
ssh madking.esm.rochester.edu lsf; cat Section2/Bmix.wav Section2/Cmix.wav
```

This would list the soundfiles in your home soundfile directory on *madking*, then would display the files *Bmix.wav* and *Cmix.wav* in your Unix subdirectory *Sections2*.

Perhaps you already are growing weary of typing the refrain "machine.esm.rochester.edu." If so, you can create a bash function or csh alias abbreviation for machines that you access frequently with *ssh*. Ask a staff member for help in doing this.

Some complications

☞ We keep the ECMC machines fairly close to the current version of *ssh*. As of this writing we are running version 3.9 of *ssh* on *madking*, but this may have changed by the time you read this. However, the versions of *ssh* that we are running at any given time on all of our ECMC systems may not be identical. Some *ssh2* versions store all of a user's *ssh* information within a directory other than *~/ssh*. Additionally, the file names within this directory may not have exactly the same names as those given in the examples here.

Furthermore, as of this writing some servers are still employing older *ssh 2.x* protocols, which may not be fully compatible with some features of *ssh3*. If you try to connect with a non-ECMC machine using *ssh* and some things don't work, this probably is the reason. Type *ssh -V* while logged on to the host machine to find out what version of *ssh* it is running. You can consult *ssh* online documentation to find alternative commands that work with older *ssh* systems.

☞ Perhaps more importantly —and here is something you won't want to hear —everything we have done so far only implements the first layer of *ssh* authentication, which beats *telnet* and *ftp* by miles, but is not really secure against a determined cracker. (And most crackers are determined!) Unfortunately, implementing the second level, through the creation and distribution of "keys," can be tedious, especially if you must do it on several machines. So, take a deep cleansing breath, roll up your sleeves, and get help from one of the teaching assistants if necessary.

6.3.1. Creating keypairs and passphrases to fully implement ssh security

I am logged onto my home Linux system and want to create a "public" and "private" pairs of *ssh* keys, and then distribute these keys to *madking* and then to some other remote machines I access frequently. Here we go!⁵

- 1) Generate the keys.

Type: *ssh-keygen -t rsa*

The key-generating *ssh-keygen* program now will engage you in lively conversation, and will respond with something like this:

```
Enter file in which to save the key (/home/allan/.ssh/id_rsa): <hit carriage return>
```

```
Enter passphrase (empty for no passphrase):
```

```
(Type in a passphrase)
```

```
Enter same passphrase again: (re-enter the passphrase)
```

```
Your public key has been saved in /home/allan/.ssh/id_rsa.pub.
```

```
The key fingerprint is:
```

```
29:cy:a7:d3:16:13:28:34:4d:7f:6c:a0:00:19:ed:f2 allan@localhost.localdomain
```

Note that within this dialog you are asked to type in not just a *password*, but rather an entire **passphrase**. Unlike login passwords (which it will replace for remote logins), this *passphrase* can (and should) include several words, spaces, and special characters. Think before you type —you ultimately may want to use this same passphrase on several computers, and you may type it a LOT. You don't want to forget it, but you want it to be completely secure.

⁵ Depending upon the current version of *ssh* running on the local machine your dialogue with *ssh-keygen* may not go exactly (word-for-word) as shown here, and the file names may be different.

3) Distributing your public key.

Still with me? Now comes the onerous part.

Next, we need to distribute this "public" key to all computer systems to which we will want to connect with *ssh* from our current machine. First I am going to make a temporary copy of this file, and call the copy *homepubkey*:

```
cp ~/.ssh/id_rsa_pub homepubkey
```

You can look at this ASCII file if you wish. Now, copy this temporary file to your home accounts on all hosts to which you will want to connect from this machine. In our example here, I will copy this file to my home directory on *madking*.

```
scp homepubkey madking.esm.rochester.edu:/home/UID/.
```

4) Install the public key

Next I log on remotely to *madking*:

```
ssh madking.esm.rochester.edu (or ssh -l UID madking.esm.rochester.edu
```

- I make certain that the copy of my *homepubkey* public key was copied successfully to *madking*:

```
ls -l homepubkey
```

- Determine whether your *ssh* directory on the host is called *.ssh* or something else:

```
ls -a | grep ssh and see if you already have a file called authorized_keys within that directory:
```

```
ls -l .ssh/authorized/keys or ls -l .ssh2/authorized/keys
```

- Now I paste the key for *homepubkey* to the end of you *authorized_keys* file:

```
cat homepubkey >> .ssh/authorized_keys
```

If the file *authorized_keys* did not exist before, it will be created now.

When executing this command, be careful to use the double redirect symbol `>>` (which appends to a file if it already exists) rather than the single redirect `>`, which will cause the file to be overwritten, destroying any information it previously contained about other machines.

- I check to make certain that the paste has been successful:

```
cat .ssh/authorized_keys
```

If so, I can delete the temporary file: `rm homepubkey`

I can now log out of *madking*.

From now on, whenever I remotely log on to *madking* or use *scp* to *madking* to copy files to or from one of these machines, I will be prompted for the *passphrase* I created with *ssh-keygen*, rather than for my password on the host.

However, if you use *ssh* frequently—especially for copying files between machines—you likely will find it increasingly irksome to type your passphrase every time you wish to access a remote system (and for every file transfer). Fortunately, this is not necessary. It is possible to have the *ssh-agent* program "inform" your shell of your passphrase so that it automatically is transmitted whenever you issue an *ssh*-based command. If you have a *bash* window open (*bash* currently is the default ECMC shell), type:

```
ssh-agent /bin/bash
```

(*tsh* users instead must type `ssh-agent /bin/tsh`)

```
ssh-add
```

When prompted for your *passphrase* type it in.

Now, within this shell window, you can use *ssh* and *scp* with any machine that has your keypairs without typing in your passphrase.

It also is possible to set up your login process so that *ssh-agent* starts your *KDE* or *Gnome* window manager, so that every shell you open will "know" your passphrase and transmit it automatically as required. However, this can be tricky to set up the first time, so consult a staff member for help if you want to try it.

6.4. Copying files to and from remote systems

ssh-based software provides various ways in which we can copy files securely from a hard disk on one computer system to a hard disk on another system (or "node"). In remote system file copying operations, the machine on which you are logged on is called the *local* system, and the computer to which, or from which, you copy files is known as the *remote* system.

The basic program for copying files between any two computers that include *ssh* encryption and authentication software is called *scp* ("secure copy"). *scp* can be used either to upload files from the local system to the remote system (the easier and more common procedure), or to download files from the remote system to the local system.⁶ The basic *scp* syntax is:

```
scp file1 [file2 fileN] UID@host:. (copies files from the local to the remote system)
```

or

```
scp UID@host:file1 [file2 fileN] . (copies files from the remote to the local system)
```

Note the concluding dot (.) in these lines, which means "to here" (your current directory on the local system) or else "to there" (your home directory on the remote system). *host* (followed by a colon) is the full name of the remote system. The *UID@* can be omitted unless your login name differs on the two systems. Unix wildcard characters such as * can be used to specify groups of files with similar names.

Thus, if I am logged on to some Unix-based system and want to copy a file from my *home* Unix directory on *madking* to the machine on which I am logged in, I would type:

```
scp madking.esm.rochester.edu:/home/allan/filename
```

Things get a little uglier if I want to copy several files from *madking* to my remote system:

```
scp madking.esm.rochester.edu:/home/allan/file1 madking.esm.rochester.edu:/home/allan/file2 madking.esm.rochester.edu:/home/allan/file3 .
```

To copy a soundfile from *madking* to my local system, I would move into the directory on the local system where I want to place the soundfile, then type:

```
scp madking.esm.rochester.edu:/snd/allan/filename.wav
```

To copy three soundfiles from my local system to my soundfile directory on *madking*, I would type:

```
scp soundfile1.wav soundfile2.wav soundfile3.wav allan@madking.esm.rochester.edu:/snd/allan/.
```

Another command line program that can be used to copy files to and from a remote system is *sftp*. If you know *ftp*, you can learn commands for *sftp* fairly quickly. If you have never used *ftp*, however, it probably is not worth your time to memorize *sftp* commands. However, an easier way to use *sftp* is by means of the *Konqueror* browser. To open a connection to another computer system from the *Konqueror* browser type:

```
sftp://UID@machinename
```

in the *Location* box. For example, to get to my home directory on *madking* from my home Linux system, I would type:

```
sftp://allan@madking.esm.rochester.edu
```

I will be asked for my *madking* password, and then should see my *madking* files displayed in the window. I can navigate to any of my directories on *madking* and view their contents. To copy a file from *madking* to my home system, I need to open another *Konqueror* browser window and navigate within it to the folder where I want to place the files to be copied from *madking*. Then I need only *drag* any files I want to copy from one browser window to the other.

Note that all of the examples above only work between Unix-based systems. If you are logged onto an ECMC Windows system — *gesualdo* in room 52 or *igor* in the MIDI studio — you can use the *putty* client to log on remotely to *madking*. You then can execute almost any command that you could execute directly on *madking* — list and view your files, read an *ecmchelp* file, and so on. To copy files from the Windows system to *madking*, or from *madking* to *gesualdo* or *igor*, use the graphical Windows version of *scp* called *Win SCP*.

⁶ Actually, *scp* also can be used to copy files between two remote systems (for example, between *madking* and *wozzeck* while you are logged onto some other machine. Consult the *scp man* page if you are interested in this possibility.

6.5. Compressing files

A *compression* program reduces the size of a file. Compression can be applied to files of all types, including soundfiles, ASCII text files, executable binary files and, especially (as discussed in the following section) to archive files. Compression is frequently applied to files before they are archived to a data CD, DVD or removable disk, enabling us to squeeze more files onto the archive storage medium. However, compression can also be usefully applied to large files that will remain on a hard disk, but which will not be needed for a while. This reduces the amount of disk space consumed by the file.

Files that have been compressed are not immediately or fully usable with many standard Unix or music software commands while in this compressed state. You cannot view a compressed ASCII file with *cat* or a text editor, and you cannot play a compressed soundfile. Compressed files are restored to their original, fully usable state by applying an *uncompress* operation.

General purpose Unix/Linux compression programs such as *gzip* and *bzip2* can be applied most any type of file. However, we do *not* recommend using these general purpose utilities as the principal means of reducing the size of individual *soundfiles*, because they are not very effective for this task. Substantially greater soundfile compression can be obtained by using programs specifically designed for this purpose. However, consolidating groups of soundfiles into a single "tarball" archive file and then copying this tarball to a data CD or some other archive medium can significantly reduce the data space required to store these soundfiles.

6.5.1. The Unix TAR program and ECMC tarsf utility

tar is a multi-purpose Unix program used to archive (save and, eventually, restore) multiple files, compacting several individual source files into a compact continuous data stream. As noted below, *tar* is similar, in some respects, to the Windows applications *pkzip* and *WinZip*, and — most importantly for cross-platform purposes at the ECMC — is partially compatible with *WinZip*. On ECMC systems the medium to which *tar* writes its output is usually an archive file, which consolidates many individual source files within a single physical file on a hard disk. We frequently use *tar* (and a local ECMC variant called *tarsf*) to collect groups of Unix files, soundfiles, combinations of ASCII files and soundfiles, or even all of the files within a directory or several directories, into a compact *tar* archive file, which can be further compressed by using a Unix compression program such as *gzip* or *bzip2*.

Once we have created a *tar* archive file, we can do several things with it:

- Move or copy the *tar* file to another directory, or to another ECMC system or our home machine, and then extract its contents. This is generally much quicker and more efficient than copying a group of files one at a time.
- Archive one or more of these *tar* files to a data CD or DVD. A *tar* file will take up less space on the archive disc than the original files, and even less space if we apply compression to the *tar* file. This enables us to store more data onto the data CD or other archive medium.

After capturing several files into a *tar* archive file or tape, we can delete the original source files if they will not be needed for some time, in order to conserve hard disk space and clean up the clutter on our work space. We can easily restore one, a few or all of these individual files from our *tar* archive at any time.

tar command lines

tar is a command line program run in a shell window. Command line arguments to *tar* consist of

- (1) a **flag** argument that tells *tar* what to do;
- (2) either an output *storage device* or else a *file name*, which tells *tar* where to write its output or find its input; and
- (3) the names of one or more files to be saved, listed or restored.

tarsf

Remember that on the ECMC system *making* your shells always have two working directories:

- 1) your current working **Unix** directory, where your ASCII files are stored; and
- 2) your current working **soundfile** directory (*\$\$FDIR*), where your soundfiles and other large music-related files (e.g. phase vocoder and *sms* analysis files) are stored

tar operates from your current Unix directory. The ECMC utility **tarsf** is identical to *tar* in every respect except that it runs *tar* from your current working *soundfile* directory. Running *tarsf* instead of *tar* is equivalent to the following sequence of commands:

- *pushd \$SFDIR* (move into your current working *soundfile* directory)
- run the specified *tar* command
- *popd* (return to your previous directory)

Thus, in all of the *tar* command examples that follow,

- ☞ use *tar* when you are operating on ASCII (or other types of small) files in your current Unix directory, and use
- ☞ *tarsf* instead when you are working with *soundfiles* and related large files in your *\$SFDIR*

Basic tar commands

Original Unix versions of *tar* employ a single character for each command flag or option, and a complete flag sequence usually consists of two, three or more of these characters butted together immediately after the call to *tar*. These traditional single character command flags can be used with the GNU/Linux version of *tar* as well, and some of us prefer them. However, the GNU/Linux version of *tar* also provides longer, mnemonic alternatives, preceded by a single or double dash and indicated in [italicized square brackets] below, which you can employ instead if you prefer. The following summary of basic *tar* commands will be sufficient for most users:

c [*--create*] Create a new data tape (overwriting any current contents of this tape) or a new archive file, and write the named files onto it.

f [*--file*] Use the argument that follows as the name of the output device driver (e.g. the driver that controls the operation of a tape drive) or else as an archive file name.

r [*-A --append*] Read (append) the named files onto the end of an existing tape or archive file

u [*--update*] Only append files that are newer than the copy in the archive

t [*--list*] List all of the files contained in a *tar* tape or archive file.

x [*--extract --get*] Extract the named files from the tape or archive file.

If no file names are specified, the entire contents of the tape or archive file will be extracted

v [*--verbose*] By default *tar* does its work silently, but with this verbose option it will display the name and size of each file it acts upon. You will generally want to include this flag so you will know what *tar* is doing.

z [*--gzip --ungzip*] Apply *gzip* compression or uncompression

I [*--bzip*] Apply *bzip2* compression or uncompression

For more complete information on using the GNU version of *tar*, type

info tar or *man tar*

The *ecmhhelp* file *tar* contains an abbreviated summary, similar to that presented here.

Working with tar archive files

To capture a group of files into a *tar* archive file, you must include the **f** flag on your command line, followed by the name you give to this archive file. In naming this "tarball" archive file, it is highly recommended that you include a *.tar* filename extension.

ASCII file example:⁷

*tar cvf scorefiles.tar file1 file2 file3 mar**

or else this equivalent alternative:

*tar --create --verbose --file scorefiles.tar file1 file2 file3 mar**

⁷ Note: On GNU/Linux systems, it also is possible, and in fact generally recommended, to apply compression or uncompression to all the commands that follow (a procedure discussed in the next section) so long as you will not be appending additional files to the archive in the future.

Result: ASCII fi les *fi le1*, *fi le2* and *fi le3* in your current working Unix directory, and all fi les within this directory that begin with the character string *mar*, are written to an archive fi le called *scorefi les.tar*, also located in your current Unix working directory.

To list the contents of this archive fi le, type:

```
tar tvf scorefi les.tar
or else: tar --list --verbose --fi le scorefi les.tar
```

Once we see that all of the fi les have been written successfully to the archive fi le, with no error messages, we can safely delete the original fi les:

```
rm fi le1 fi le2 fi le3 mar* (be careful with that asterisk!)
```

At some later time we can extract only *fi le2* from the archive by typing:

```
tar xvf scorefi les.tar fi le2
or else tar --extract --verbose --fi le scorefi les.tar fi le2 (Linux only)
```

To extract all of the individual fi les and/or folders within the *tar* fi le, type:

```
tar scorefi les.tar xvf
or else: tar --extract --verbose --fi le scorefi les.tar
```

A soundfi le example:

```
tarsf cvf mixes.tar *mix*
or its equivalent: tarsf --create --verbose --fi le mixes.tar *mix*
```

Result: All soundfi les within your current working soundfi le directory whose names include the character string *mix* will be written to an archive fi le named *mixes.tar* in your *\$\$FDIR*.

To list the contents of this archive fi le, type:

```
tarsf tvf mixfi les.tar
or tar --list --verbose --fi le mixfi les.tar (Linux only)
```

Here again, if this command lists all the soundfi les we believe should be in this archive, we can delete the originals:

```
rmsf *mix*
```

To extract only soundfi le *Bmix.wav* from this archive, type:

```
tarsf xvf mixfi les.tar Bmix.wav
or else: tarsf --extract --verbose --fi le mixfi les.tar Bmix.wav
```

To extract all of the soundfi les within the archive, type

```
tarsf xvf mixfi les.tar or tarsf --extract --verbose --fi le mixfi les.tar
```

Once we have created a tar archive fi le, and perhaps compressed it (see below), we might decide to move it to another directory, or to archive it to a DAT tape or to a Zip disk. Or we could copy this fi le from one ECMC computer to another, or to our home system.

Beware, however, of moving or copying very large tar fi les between disks, or between computers, during periods of heavy system traffi c (such as when you or some other user is playing or mixing soundfi les). The massive I/O involved in such transfers, if frequently interrupted by other system traffi c, can cause sluggish system performance.

For advanced users: The *tar* command also can also be used to copy several fi les directly from one directory to another, without creating a *tar* archive fi le. To copy several fi les from your current working Uix directory to some other directory (here called *NEWDIR*), type

```
tar cf - fi le1 fi le2 fi le3 | ( cd ^NEWDIR ; tar xf - )
```

(The GNU/Linux version of *tar* may complain with an error message while performing this command, but it will work.)

6.5.2. Compressing fi les with gzip or bzip2

gzip currently is the most widely used general purpose compression program on Linux/Unix systems. One reason for this popularity is cross-platform support. Certain types of gzipped fi les —notably gzipped *tar* archive fi les —also can be uncompressed and extracted on Windows systems with the *Winzip* utility.

bzip2 is a more recent general purpose Linux/Unix compression program that often provides 20 % or more greater compression than *gzip*. However, it is less frequently used, and bziped files currently cannot be uncompressed and extracted by *Winzip* on Windows systems. Unix-based Macintosh systems (OS X and higher) include both *gzip* and *bzip2*.

The commands used to compress and uncompress files with *gzip* and *bzip2* are very similar. To *compress* one or more files:

```
gzip filename [filename2 filenameN]
```

or

```
bzip2 filename [filename2 filenameN]
```

gzip will append the extension *.gz* to the compressed file names. *bzip2* will append the file name extension *.bz2* to files it compresses.

To view an ASCII file that has been compressed you can type

```
zcat filename for gzipped files or
```

```
bzcat filename for files that have been compressed with bzip2.
```

To *uncompress* (decompress) files, you can use any of the variant commands below. *gzip* allows you to omit the *.gz* file name extension of the compressed input files. With *bzip2*, however, you must type in the complete input file names (you cannot omit the *.bz2* extension).

```
gzip -d filename [filename2 filenameN]
```

or

```
gzip --decompress filename [filename2 filenameN]
```

or

```
gunzip filename [filename2 filenameN]
```

```
bzip2 -d filename.bz2 [filename2.bz2 filenameN.bz2]
```

or

```
bzip2 --decompress filename.bz2 [filename2.bz2 filenameN.bz2]
```

or

```
bunzip2 filename.bz2 [filename2.bz2 filenameN.bz2]
```

Example: To compress an archive file previously created with *tar*:

```
gzip $$FDIR/SECTION2.tar
```

Result: The *tar* archive file *SECTION2.tar*, located in our current working soundfile directory, is compressed, and renamed *SECTION2.tar.gz*. (Note that for maximum compression, the source soundfiles within the *tar* archive should have been compressed with *shorten/runsfcompress* before the archive was created.)

Later, to uncompress this archive and extract the source soundfiles, we could issue this sequence of commands:

```
gunzip $$FDIR/SECTION2.tar (or gunzip SECTION2.tar.gz)
```

and then

```
tarsf xvf SECTION2.tar
```

gzipped tarballs like this example are very commonly encountered on web and *ftp* download sites. Consolidating all of the files required by a downloadable application, or group of applications, into a single compressed archive file not only conserves disk space on the server, and also greatly reduces downloading times across the net. Downloadable files with the extension *.tgz* (an abbreviation for *.tar.gz*) are always in this format.

Wild card characters (chiefly ***) can be used to specify groups of files:

```
gzip *.orc *.sco
```

will cause *gzip* to compress all files in your current Unix directory that end with the character string ".orc" or ".sco"

```
bzip2 -d $$FDIR/*.bz2
```

will cause *bzip2* to decompress all files in your current working soundfile directory that have been previously compressed with *bzip2*.

To display a list of additional command options type:

```
gzip -h or bzip2 -h
```

For full details, consult the *man* page for *gzip* or *bzip2*.

Applying tar and compression/uncompression simultaneously

In examples above, groups of files were consolidated into a *tar* archive file, and this archive file was then further compressed with *gzip* or *bzip2*. However, because *tar-and-compress/uncompress* procedures are so common, the GNU/Linux version of *tar* enables us to consolidate these two processes in a single operation, by including the flag *z* (or *--gzip*) for *gzip* compression/uncompression, or else the flag *Z* (or *--bzip2*) for *bzip2* compression/uncompression, while executing other *tar* commands.

Example:

```
tarsf czvf B.tgz B* b* or else tarsf --create --zip --verbose --file B.tgz B* b*
```

Result: All soundfiles (hopefully already compressed with *shorten*) within our current working soundfile directory whose names begin either with a capital or lower case "B" are written to a *tar* archive file named *B.tgz*, which is compressed with *gzip*.

Later, to view the names of the files within this archive, we can type:

```
tarsf tzf B.tgz or else tarsf --list --gzip --file B.tgz
```

(Note that we have omitted the common "verbose" flag here, since we only want to list the names of the files within the archive, rather than to obtain a "verbose" listing of each of these files.)

Later, to extract two of the soundfiles in this archive, we could type:

```
tarsf xzvf B.tgz b2vln.shn b2tuba.shn or tarsf --extract --gzip --verbose --file B.tgz b2vln.shn b2tuba.shn
```

This will extract the soundfiles *b2vln.shn* and *b2tuba.shn* from the archive. From the *.shn* extension on these two file names we can tell that they have been compressed with *shorten*, so to restore them to a usable state, we then would type

```
runsfuncompress b1vln b1tuba
```