

User Guide for RAT v3.0.33

March 1999

Acknowledgements

This user guide is based on the UKERNA SHRIMP User Guide. Unlike the UKERNA guide, this document is updated to reflect developments in RAT.

Disclaimer

UCL cannot accept any responsibility for any loss or damage resulting from the use of the material contained in this document. The information is believed to be correct, but no liability can be accepted for any inaccuracies.

Copyright note

Copyright © 1995-1998 University College London
All rights reserved.

Redistribution and use in source and binary forms, with or without modification, is permitted, for non-commercial use only, provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
3. All advertising materials mentioning features or use of this software must display the following acknowledgement:
"This product includes software developed by the Computer Science Department at University College London."
4. Neither the name of the University nor of the Department may be used to endorse or promote products derived from this software without specific prior written permission. Use of this software for commercial purposes is explicitly forbidden unless prior written permission is obtained from the authors.

THIS SOFTWARE IS PROVIDED BY THE AUTHORS AND CONTRIBUTORS "AS IS" AND ANY EXPRESSED OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE AUTHORS OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

This software is derived, in part, from publically available source code with the following copyright:

Copyright © 1991-1993 Regents of the University of California

Copyright © 1992 Stichting Mathematisch Centrum, Amsterdam

Copyright © 1991-1992 RSA Data Security, Inc

Copyright © 1992 Jutta Degener and Carsten Bormann, TU Berlin

Copyright © 1994 Paul Stewart

Copyright © 1996 Regents of the University of California

This product includes software developed by the Computer Systems Engineering Group and by the Network Research Group at Lawrence Berkeley Laboratory.

Encryption features of this software use the RSA Data Security, Inc. MD5 Message-Digest Algorithm.

Contents

1	Introduction	1
1.1	Before you start	1
1.2	Getting started.....	2
1.2.1	Starting up RAT from SDR.....	2
1.2.2	Starting up RAT from the command line for a multicast conference	2
1.2.3	Starting up RAT from the command line for a unicast conference ..	3
1.2.4	Quick guide to using RAT.....	4
1.3	RAT user guide.....	6
1.3.1	RAT main window	6
1.3.2	RAT options window.....	9
1.4	Known bugs.....	16
	Further Reading.....	17
	Glossary.....	18

Boxes

Box 1: Unicast and Multicast 1

Box 2: TTL and Admin scope 3

Box 3: If you can't hear other people 5

Box 4: Silence Suppression 7

Box 5: Push to Talk 8

Box 6: Full Duplex and Half Duplex 8

Box 7: RAT Options Window 11

Box 8: RAT Synopsis 13

Box 9: RAT Command Line Options 13

Box 10: RAT Examples 15

Box 11: RAT CODECs 15

Box 12: RAT Packet Overheads 16

Figures

Figure 1: Main RAT Window 1

Figure 2: RAT User Information Window 4

Figure 3: RAT User Info Window 6

Figure 4: RAT Control Window 9

1 Introduction

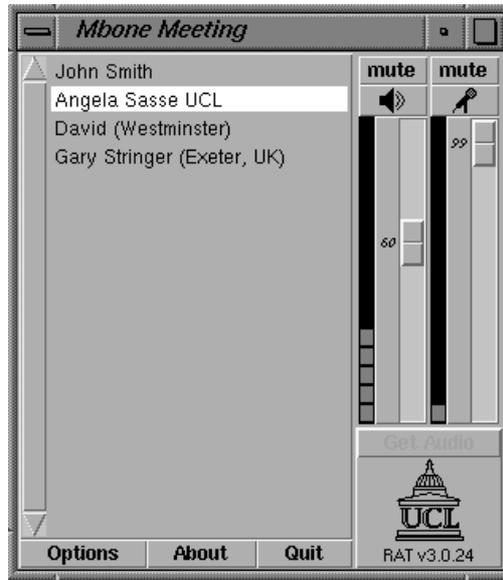


Figure 1: Main RAT Window

RAT (Robust Audio Tool) - unicast and multicast audio conferencing tool. RAT can be used for both point-to-point videoconferencing involving a direct link between two hosts or for multiparty conferencing with many participants via the Internet Mbone (see Box 1).

Unicast and Multicast

A

A *unicast* connection is like a telephone call which has one connection between two telephones for the duration of the call. You should use unicast if there are only two people in the conference and if one of you does not have access to multicasting. In order to set up a unicast connection, you must know the name of the computer that the person you want to have a conference with is going to use.

A *multicast* connection can, in principle, have an unlimited number of participants. You should use multicast if it is likely that there will be more than two participants in the conference. Tools like RAT, VIC, NTE, WB and SDR have been specially developed for multicasting.

Box 1: Unicast and Multicast

1.1 Before you start

In order to receive audio, there must be either a speaker or a pair of headphones connected to your computer.

In order to transmit audio, you must have a microphone plugged into your computer.

It is recommended that headphones are used in preference to speakers (see Box 6). Using headphones allows you to listen and talk at the same time (this is known as full duplex audio). Without the use of headphones the output from the speakers can interfere with the microphone, causing unpleasant feedback. The only way of avoiding this is either to use an echo canceller, or to switch to half duplex audio, which means that you can only speak or listen at any one time, i.e. you cannot listen to others when you speak, or you cannot speak when you are listening. This is unnatural and can make conversation difficult.

! Some audio cards do not support full duplex.

1.2 Getting started

There are two ways of starting up RAT: from the command line and from SDR. Using SDR is easier as it automatically configures the tool for use.

1.2.1 Starting up RAT from SDR

When joining a conference with audio using SDR, RAT will start up automatically. All parameters will be set automatically.

1.2.2 Starting up RAT from the command line for a multicast conference

As with the other Mbone tools the format for setting up RAT from the command line is as follows:

Unix:

```
Prompt> rat [options] <address/port>
```

Windows95/NT4.0:

Click on the *Start* button in the bottom left hand corner of the screen. Select *Run...* In the Run Window, type the following in the *open:* field.

```
Open: rat [options] <address/port>
```

The multicast address should be in the range 224.2.0.0 - 224.2.255.255 (except when using admin scope – see Box 2). The port number should be at least 5002.

So, as an example, this line would start RAT with a TTL of 47 (see Box 2 for an explanation of TTL) on the multicast address 224.2.2.60 and port number 5004:

```
Prompt>/Open: rat -t 47 224.2.2.60/5004
```

In the example above the user has used the option '-t' to specify the TTL. TTL determines how far the audio packets will go, i.e. how far away you can talk to people. A TTL of 47 is generally large enough for communication within the UK.

 All participants must use the same multicast address and port number in order to talk to one another.

TTL

TTL stands for Time To Live and determines how far the audio that you transmit will travel. A TTL of 15 will reach other parts of your campus, a TTL of 47 will reach all parts of the JANET Mbone, a TTL of 63 reaches Europe, and 127 worldwide.

The default for all the Mbone tools is a TTL of 16.

Admin scope

There is currently a move away from using TTL for scoping conferences to using admin scopes. Using admin scope will limit conferences more accurately and is more sophisticated than specifying a TTL.

Like TTL scopes, admin scopes need to be set up by the network administrator. When the scopes have been configured, you will get a range of addresses for each of the scopes defined. Using an address in the range provided will limit your conference to that scope.

Example: If local scope is defined as addresses in the range of 239.0.0.0 - 239.100.255.255, any conference on addresses in this range will be restricted to the local area as defined by the network administrator.

The most widely used method for scoping is still TTL, so you should contact your local systems administrator to see if admin scopes have been set up. If not, you should ask them to do it.

Box 2: TTL and Admin scope

1.2.3 Starting up RAT from the command line for a unicast conference

As with the other Mbone tools the command format is as follows:

Unix:

```
Prompt> rat [options] <remote_hostname/port>
```

Windows95/NT4.0:

Click on the *Start* button in the bottom left hand corner of the screen. Select *Run...* In the Run Window, type the following in the *open:* field.

```
Open: rat [options] <remote_hostname/port>
```

So, as an example, this line would start RAT for a unicast conference with someone on the computer called myhost.example.ac.uk on port 5004:

```
Prompt>/Open: rat myhost.example.ac.uk/5004
```

The person at the other end, in this case the person who is working on myhost.example.ac.uk will have to do the same command but with the name of your computer instead of myhost.example.ac.uk. The port number must be the same on both machines.

TTL is not applicable in unicast sessions.

1.2.4 Quick guide to using RAT

The first time you start up RAT, a window might appear where you have to enter your name and then click the *Done* button (see Figure 2).

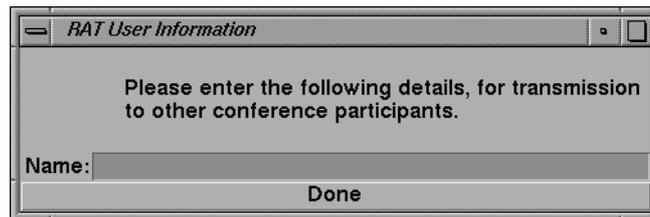


Figure 2: RAT User Information Window

These details will be stored in a file called *.RTPdefaults*. (The details are also used by the video tool VIC.) The file will be stored in your home directory. The name that you type in will appear in the RAT Main Window (see Figure 1), and will identify you to all other participants in the session. You can change your name at any time by changing the text in the name field of the RAT Options Window (see section 1.3.2).

In the Main Rat Window, you can use the left slider bar to adjust the volume of the speaker or headphones. When someone is talking, the left power meter will be activated, and the name of the speaker will be highlighted in the main window (see Figure 1).

A

If you can't hear other people:

Make sure that the left *mute* button is *not* pressed down. When the *mute* button is pressed down, the icon (depicting either a speaker, headphones, or the symbol for line out), will appear inside a grey box, and the audio will be muted.

Make sure that someone is actually talking. You can see that someone is talking when the left power meter in the RAT audio window is going up and down.

Make sure that the volume of the speaker or headphones is not 0. The volume can be adjusted on the slider bar next to the left power meter.

Make sure that the button with the little icon under the left *mute* button is depicting headsets if using headsets, and speakers if using speakers. You can change between headsets, speakers and line-out by clicking on the button with the icon. (Some audio cards will automatically send audio to headsets if they are plugged in and speakers if not. In this case, the button will not toggle, but will always display a speaker)

Check that your headsets are plugged into the computer correctly.

Box 3: If you can't hear other people

In order to *talk* to other people, you must unmute your microphone. This can be achieved by either clicking on the right *mute* button with your left mouse button, *or* holding down the *right* mouse button anywhere inside any RAT window while you are talking. If you are using headphones rather than speakers, you can leave the microphone unmuted all the time. If you are using speakers (without an echo cancellor connected), the microphone should only be unmuted when you wish to talk.

 If the other participants' audio is too faint, ask them to turn up their volume. You can *adjust the volume* of your microphone with the right slider bar. On most audio cards, you will not hear yourself in your headsets or through the speakers when you are talking. Whether you can hear yourself or not, you will to a certain extent have to rely on the power meter to adjust your microphone volume – if you can hear yourself, the volume you hear is not necessarily related to the volume other participants will hear. When you speak, the power meter should not go all the way into the red area at the top as this will sound distorted. Generally, you have to rely on the other

participants in the conference to tell you whether your volume is set correctly.

 *If the audio from another participant is very broken up* The most likely causes of this are either network congestion, or because the other user is speaking too quietly for the audio tool to distinguish them from the background noise. You can check if you are receiving packet loss from another participant by clicking on their name in the Main RAT Window. A small window will appear (see Figure 3). The last line in this window will tell you how much packet loss (instantaneous Loss Rate) you are receiving from that participant. If the packet loss is small, i.e. less than 5%, ask them to turn their volume up. If that does not help, ask them to turn off silence suppression (this option is available from the *Options* menu, accessible from the RAT audio window).

 *If everyone's audio keeps cutting out* When you move windows around on the screen, the processing power of the computer will be spent on window management rather than audio decoding, and you will lose some audio quality. If possible, try not to move windows around unnecessarily.

You should now be able to use RAT for taking part in conferences.

1.3 RAT user guide

RAT consists of a Main Window (see Figure 1), an Options Window (see Figure 4), and a number of smaller windows. The following sections discuss the Main Window and the Options Window in greater detail. The different command line options are described at the end of the chapter.

1.3.1 RAT main window

The Main Window (see Figure 1) of RAT is split into three sections. On the left is a list of conference participants. Your own name will always appear at the top. Whenever you speak, your name will be highlighted. Clicking on the name of a participant (except yourself) will display a user information panel, giving various statistics for that user, such as reception statistics, which audio tool they are using, which version of the tool they are using, and so on (see Figure 3).

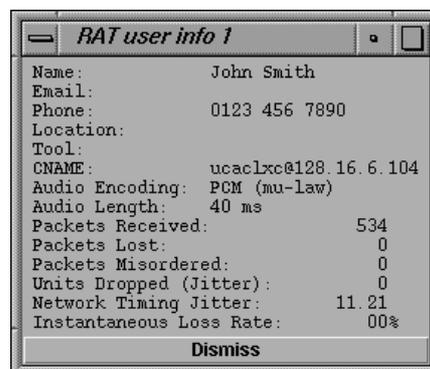


Figure 3: RAT User Info Window

On the right of the Main RAT Window are volume/microphone controls. There are two sets of controls, the right set is for outgoing audio (microphone) and the left set is for incoming audio (headphones or speakers). Each set of controls consists of a *mute* button, a button for selecting the input/output device, a volume control and a power meter.

 *If other people cannot hear you* Make sure that your microphone is unmuted when you talk. The microphone is unmuted when the right *mute* buttons in the RAT audio window is *not* depressed. The microphone icon immediately beneath the mute button will be greyed out when muted and not greyed out when unmuted.

Also make sure that the TTL is high enough (see Box 2). If your TTL is too low, you may still be able to hear a remote participant even though they cannot hear you. In order to change the TTL you will have to restart RAT.

A

Silence suppression

By default RAT starts up with silence suppression on. The silence suppression ensures that only audio above a certain level is transmitted. This means that you can leave your microphone open, i.e. un-muted, throughout the conference and silence suppression will ensure that background noise is not being transmitted.

However, sometimes the silence suppression will cut out speech as well as background noise. This can happen if your microphone is a long way away from your mouth, making your voice sound faint, or if the microphone is not powerful enough. If whatever is causing the silence suppression to cut you off cannot be rectified, you have to switch off silence suppression in the RAT Control Window (see Figure 4), and use push-to-talk (see Box 5).

Box 4: Silence Suppression

A

Push-to-talk

If you prefer to leave your microphone muted when you are not talking, for example if you have switched off silence suppression (see Box 4), or if you are listening to a lecture, you can use push-to-talk. It means that you keep your microphone muted, and only un-mute it when talking. The easiest way of achieving this is by right clicking anywhere within a RAT window, which will toggle the microphone mute button. This is why is it called push-to-talk – because you “push” the right mouse button when you talk.

Box 5: Push to Talk

A

Full duplex and half duplex.

Full duplex means that you can talk and listen at the same time and is the most convenient method for taking part in a conference. In order to achieve full duplex, you select *Full duplex* under *Mode* in the RAT Control Window. You must use headphones when using *full duplex* in order to avoid feedback from the speakers – alternatively, you can use an echo canceller.

NB: Full duplex is not supported by all audio cards.

Half duplex: if your audio cards does not support full duplex, you must use half duplex which means that at any point in time, you can only either *listen* to other participants or *talk* to them. When in half duplex mode you must select *Net mutes mike* or *Mike mutes net* under *Mode* in the RAT Control Window (see Figure 4).

When *Net mutes mike* is selected, your microphone will be muted automatically as soon as another participant is talking. This means that if someone is sneezing while you are talking, nobody will be able to hear you.

When *mike mutes net* is selected, you cannot hear other participant while you are talking. This means that you cannot hear people responding to what you are saying while you are talking.

If the audio card supports full duplex, this will be the default setting. If not, Net mutes mike will be the default.

Box 6: Full Duplex and Half Duplex

When pressed down, the *mute* buttons mute incoming and outgoing audio. When muted, the icons underneath the *mute* buttons will be greyed out. The mute button for the outgoing audio, i.e. microphone or line out, can be operated from anywhere within the Main RAT Window by holding down the right mouse button with the cursor anywhere in the Main RAT Window.

The buttons for selecting input and output devices are toggle buttons. The button for outgoing audio toggles between microphone and line in. The button for incoming audio toggles between speaker, headset and line out.

Individual participants can be selectively muted by clicking on their names with the middle mouse button (or by holding both mouse buttons down at the same time on a two-button mouse if using Linux). This function is not available on Windows95/NT4.0 versions.

Between the volume/microphone controls and the logo, there is a greyed out button called *Get Audio*. You can use this if you have more than one instance of RAT running on your computer. Clicking *Get Audio* in one of the RAT windows will switch the control of the input and output devices to that instance of RAT.

1.3.2 RAT options window

Pressing the *options* button in the Main RAT Window brings up a control panel, allowing the operation of RAT to be modified. Options which are not available are greyed out.



Figure 4: RAT Control Window

The options that you might want to set are the following (for a full list of options see Box 7 at the end of this chapter):

- If the audio quality is bad due to packet loss and there are no VAT users in your conference (VAT is an audio tool, similar to RAT), set *Secondary Encoding* to *DVI*.

- If other participants complain that your audio keeps cutting out and it is not due to packet loss, try turning up your microphone volume. If that does not help, turn off *Suppress Silence* (see Box 4). If you turn silence suppression off, you should mute your microphone whenever you are not speaking. Otherwise you will continuously be transmitting background noise (see Box 5).
- If you are only listening, for example to a lecture, you should turn *Lecture Mode* on. When *Lecture Mode* is switched on, RAT delays the playout of the audio for a short while, enabling computers with limited processing power to provide better quality audio. *Lecture Mode* is automatically switched off when you talk.
- You may want to change your name as it appears in the Main RAT Window. You can do that by clicking in the *Name* field, editing the information, and pressing Enter on the keyboard. Likewise, you can change your e-mail address, phone number and location. These details will be available to other participants if they click on your name in the Main RAT Window.
- Sometimes you might want to confirm which address, port and TTL you are running RAT on. This information is also displayed in the RAT Options Window.

For a thorough description of the options available in the RAT Options Window, see Box 7.

RAT Options Window

Duration

Selects the amount of audio data, in milliseconds, which is sent in each network packet. Larger values result in greater end-to-end delay over the network, but reduce the per-packet header overhead. You should not have to change the default setting.

Encodings

Sets the format for data transmitted to the network. Two types of encoding are possible: primary and secondary (secondary encoding is generally referred to as redundancy). If the conference includes any users of VAT (a widely used audio tool, similar to RAT), redundancy should be turned off, i.e. *Secondary Encoding* should be *NONE*; however if all users are using RAT, and packet loss is being experienced, the use of redundancy will vastly improve sound quality. It is hoped that future versions of VAT will be able to decode redundant audio data. DVI is recommended as secondary encoding.

The different types of encoding use different amounts of bandwidth:

L16: 128 kbit/s: Uncompressed 16 bit samples, with linear spacing between sample values.

PCM: 64 kbit/s: μ -law companded 8 bit samples. This means that the spacing between sample values is not constant, but has more samples in the loudness levels where the ear is most sensitive. The result is similar to using 11 bit sampling, but at a lower bit-rate. Quality is very similar to L16 (the difference is imperceptible with most soundcards/headphones).

DVI: 32 kbit/s: 16 bit samples, compressed to 4 bits per sample. The basic idea is to send the first sample, and then difference from it to the next sample, and so on, although in practice it's not quite as simple as this. This sounds slightly worse than PCM, but is generally acceptable. Probably the best compromise between processing requirements and bandwidth.

GSM: 13.2 kbit/s: This is the CODEC used in mobile phones, and is specially designed for voice signals (it works by modelling the way the vocal-tract works, and hence is not suitable for music). Uses a lot of processing power, but works over a modem connection.

LPC: 5.8 kbit/s: Generates a synthetic representation of speech, which sounds slightly artificial. Should probably only be used as secondary encoding.

For general purpose use DVI+DVI (or DVI+GSM for speech only).

If bandwidth is limited, and for speech only: GSM+GSM.

If using a modem: GSM+LPC.

All use 8kHz sampling rate at present.

Box 7: RAT Options Window

Mode

Net mutes mike: Half duplex. Your speech will be cut off whenever you are receiving speech from other participants.

Mike mutes net: Half duplex. Whenever you speak, you cannot hear audio from other participants.

Full duplex: In this mode you can talk and listen at the same time. This option is greyed out (text appears in grey rather than black) if it is not supported by the computer's audio card.

Loss repair

If set to *packet repetition* the receiver attempts to patch over missing packets with the replay of the previous packet. This is a receiver-based solution to the problem of network packet loss. We recommend that redundancy is used in these cases, but if that is not possible, use of this option may help. This does not affect the data sent to the network.

Suppress Silence

If *on* is selected, periods of silence within a conversation will not be transmitted, hence reducing network traffic. We recommend that this option is on.

Powermeters

If *on*, audio power meters will be displayed in the main RAT window.

Lecture Mode

If *on*, the playout delay at the receiver is increased. This results in better performance in the presence of variable network delay, at the expense of reduced interactivity. It is most useful when listening to broadcast lectures, hence the name. This is automatically turned off if you start to transmit audio.

Video Synchronisation

If using a modified version of VIC, this enables lip synchronisation between audio and video streams. This option is not currently available.

Automatic Gain Control

If *on*, the microphone gain will be controlled automatically, i.e. the volume will remain relatively level despite shouting or speaking quietly. Currently unavailable, but will be available in future releases.

Acoustic Feedback

If *on*, audio loopback is enabled. This option is not currently available.

Encryption key

Entering a key and enabling this will encrypt the audio data using DES. This encryption is generally compatible with encrypting versions of VAT.

Box 7: RAT Options Window - Cont...

Play file/Rec file

Allows the playback/recording of audio data in the conference. Format of the files is raw 8kHz 16bit linear PCM data, with no headers.

RTP Configuration

Displays the address, port number and TTL value that RAT has been started up with, and allows the setting of the name, e-mail address, phone number and location to identify you to other members in the conference.

Box 7: RAT Options Window - Cont...

RAT Synopsis

`rat [options] addr/port`

`rat -T [options] addr/port/TTL/CODEC addr/port/TTL/CODEC`

Box 8: RAT Synopsis

RAT Command Line Options

The following command line options are supported:

-crypt *key*

Enable encryption, with the specified key. Encryption is done using DES, and is believed compatible with encrypting versions of VAT.

-f *c1/c2/./cn*

Specifies the encodings used when transmitting audio. The order is primary encoding, secondary encoding, etc. A maximum of seven levels of encoding are possible, although only two are available from the user interface. The allowed values are 116, pcm, dvi, gsm, lpc. See Box 11 below for description of the CODECs and their recommended usage.

-loopback_rtp

Causes RTP packets to be loopbacked during multicast sessions.

-lbl_channel *channel*

Specifies the channel identifier for use with the LBL (Lawrence Berkeley Laboratories) Conference Bus. RAT always listens to the base channel for audio device trading with other tools, like VAT. When a channel is specified it allows RAT to communicate with other tools, like VIC, providing that they have the same channel number. This is particularly useful for voice switching the video sources.

Box 9: RAT Command Line Options

-lbl_priority *priority*

Specifies the processing priority of the code that handles the LBL Conference Bus. The default value is 100 and the permitted range of values is 0 to 200.

-name *name*

Sets the RAT window title to the value of *name*.

-no_ui

Do not display user interface.

-p *priority*

Sets the thread priority on the Windows 32-bit version. The following values are observed:

- 1 above normal priority level
- 2 time critical thread priority
- 3 highest thread priority

You should not need to change this option.

-repair *method*.

Specifies the use of receiver based repair techniques based on method. This can currently be either none or repeat.

-t *TTL*

Specifies the TTL (Time To Live) value set in the packet headers. This limits the scope of the packets. The following values are generally considered appropriate:

- 16 campus/organisation
- 47 country
- 64 continent
- 127 world

-T *addr/port/TTL/CODEC addr/port/TTL/CODEC*

Causes RAT to operate as an RTP transcoder/mixer. In this mode the user-interface is not displayed, and no audio is played out. Instead, RTP packets received from either group are transcoded into the format specified for the other group, multiple sources are mixed together, and the result transmitted to the other group. In each case, *addr* may be either unicast or multicast, and the TTL and CODEC specifiers may be omitted (the default is TTL 16, DVI coding). This option is useful for transcoding between low- and high-bandwidth sessions, for use over a slow link, for example.

-version

Displays the version number of RAT.

Box 9: RAT Command Line Options Cont...

Examples

To start a unicast session between the current host and the host myhost.example.ac.uk using port number 12000 type:

```
Rat myhost.example.ac.uk/12000
```

To join a multicast session on group address 224.5.6.7 and port 8000 type, using primary encoding of DVI and a secondary encoding of LPC:

```
Rat -f dvi/lpc 224.5.6.7/8000
```

Box 10: RAT Examples

CODECS

Five types of audio encoding are currently possible with RAT, although more are in development. The types of encoding are:

Name	Bit rate	Description
L16	128 kbit/s	Linear PCM at 16 bits per sample
PCM	64 kbit/s	μ -law companded PCM at 8 bits per sample (G711)
DVI	32 kbit/s	Intel's DVI ADPCM at 4 bits per sample
GSM	13.2 kbit/s	EDSI Group Systeme Mobile CODEC
LPC	5.8 kbit/s	Ron Zuckerman's 10 pole LPC CODEC

Box 11: RAT CODECS

Packet Overheads

There is an overhead associated with each audio packet sent. When RAT is operating as an audio tool the RTP header is 96 bits per packet (it is longer when RAT is acting as mixer and there are multiple contributing sources in each packet). The UDP/IP overhead is 224 bits per packet. Thus transmitting with 20 ms packets has an overhead of 11.2 kbit/s, whereas 80 ms only incur an overhead of 2.8 kbit/s).

By default RAT starts with 40 ms packets of DVI coded data. This represents a trade off between bandwidth, quality, and the ability of receiver based repair techniques to work successfully.

Box 12: RAT Packet Overheads

1.4 *Known bugs*

RAT may not work reliably with half duplex audio cards.

Further Reading

RAT

Hardman, V., Sasse, M.A., Handley, M., and Watson, A. 1995. Reliable Audio for Use over the Internet. *Proceedings of INET'95*, Honolulu, Hawaii.

See also: <http://www-mice.cs.ucl.ac.uk/multimedia/projects/rat/>

Mbone

For further information about the JANET Mbone service, see the web site at:

<http://www.ja.net/documents/mbone.html>

Glossary

Bandwidth

A term used to show the amount of information flowing through a communications channel. Expressed in units of $n \times 1000$ bits per second, kbit/s.

CellB

A proprietary low bit rate video encoding developed by Sun Microsystems.

CIF

Common Image Format, fitting an image into dimensions of 352x288 pixels.

CODEC

COder/DECoder. A hardware or software processor converting between analogue audio or video and the digital format used for transmission, in both directions. The term is also used to describe the major hardware component of a videoconferencing system.

DES

Data Encryption Standard.

DVI

Audio encoding standard.

Framegrabber

A device which captures video one frame at a time from an analogue video source.

Frames per second (fps)

The frame rate for video image transmission, measured in frames per second (fps). The higher the frame rate, the better the motion rendition of the video image. 30 fps is the standard necessary for full motion video. To date, most communication channels are not capable of transmitting large numbers of video streams concurrently. With compression and the use of other signal processing algorithms, a videoconferencing VIC between 2 - 10 kilobits per second (kbit/s) is perfectly acceptable.

Gamma correction

Gamma correction is the correction of intensity values of cameras and monitors. In the context of displays, the luminescence of phosphor pixel is a non-linear function of the intensity of the electrons bombarding it. Conversely, both film and CCD cells have non-linear photon intensity versus recorded intensity functions. Gamma correction is a non-linear correcting function to overcome these problems.

GSM

Audio encoding standard.

H.261

ITU-T video encoding standard in narrowband audiovisual systems.

IP address

IP (Internet Protocol) addresses are unique numerical identifiers for each networked host computer. The IP address is used in conferencing when the conference is point-to-point rather than multiparty.

JPEG

Video encoding standard for still images. It has a high compression rate; it computes by dithering and losing “unnecessary” information, hence considered lossy. Decoding speed can be traded off against image quality.

Kilobits per second (kbit/s)

The kbit/s rate enables the participants to track how much bandwidth is being used during a session. For videoconferencing, the kbit/s should not be raised above 128 kbit/s, unless using a point-to-point conference or a Local Area Network (TTL=16). See TTL.

LBL

Lawrence Berkeley Laboratories, University of California, Berkeley, USA.

Lecture mode

An optimisation for one-way transmission which increases the delay on the layout of the media in order to minimise the loss of data in the network. The longer the playout is delayed, the more time there is to retransmit lost packets.

Loss Rate

The percentage of packet information lost during transmission.

LPC

Audio encoding standard. Linear Predictive Coder, 9 kbit/s.

Mbone

The multicast capable backbone of the Internet. It currently consists of a network of tunnels linking the islands of multicast capable sub-networks around the world.

MJPEG

Motion JPEG. The use of JPEG to encode motion video. A number of manufacturers have developed M-JPEG systems. There is no systems level standard for M-JPEG systems, and different manufacturers' systems will not interoperate.

MPEG

Motion Picture Experts Group.

MTU

Maximum Transmission Unit. A measurement, in Bytes, of the largest packet size which a network is capable of transmitting.

Multicasting

Multicasting is sending audio, video etc. on the Internet in way which ensures that anybody who is interested in receiving the information, *can* receive it, but only people who *are* interested will receive it. Think of it as being in between unicast (like most telephone calls - between two telephones only) and broadcast (TV - the signals are sent to you whether you want to watch or not).

Multicast address

Addresses used for multicast conferencing should be in the range 224.2.0.0 - 224.2.255.255

Note: Multicast addresses must be agreed upon in advance of the multicast session, either through e-mail or phone or assigned using SDR.

Network congestion

RAT User Guide

A condition in an IP network where the amount of traffic injected into the network is too great for the routers to handle and some packets are discarded.

NTE

Network Text Editor.

NTSC

National Television Standard Code. The television encoding used in the USA.

NV

Network Video. A software package, in the public domain, which supports videoconferencing over Internet Packet Networks.

PAL

Phase Alternate Line. The television encoding used in Europe.

PCM

Audio encoding standard. Pulse Code Modulation, 64 kbit/s 8-bit μ -law encoded 8 kHz.

Port

A UDP (User Datagram Protocol) number unique to the broadcast session, and should be over 5000.

Note: The port number must be agreed upon in advance, either through e-mail or phone or assigned using SDR.

Protocol

A set of standards that govern the transfer of information between computers over a network or via telecommunications systems. To reduce errors, the computers at both ends of a communications link must follow the same protocol.

QCIF

Quarter Common Image Format with dimensions of 176x144 pixels.

RAT

Robust Audio Tool.

Receive-only

A condition where a tool is used to receive information, but cannot transmit.

Redundant audio encoding

Also known as redundancy. A technique to protect against packet loss where a second, low bandwidth version of the original encoding is piggy-backed onto the preceding packet so that, when single packets are lost, the redundant version is played back instead of silence.

RTP

Real-Time Protocol is an application level protocol that provides real-time support for multimedia communication.

RTIP

Real-Time Internetwork Protocol.

SCIF

Super Common Image Format with dimensions of 704x576 pixels.

SDR

Session Directory Tool.

Silence suppression

Silence suppression is used in packet voice systems to reduce the output bandwidth of an audio system by exploiting the on-off nature of speech; 50% savings are possible.

SIP

Simple Internet Protocol.

TTL

TTL stands for time to live, and determines how far multicast packages can travel over the Internet. The standard values are:

16 - Site

47 - UK

63 - Europe

127 - World

Unicasting

A unicast connection is a point to point connection between two specific machines.

VAT

Visual Audio Tool.

VIC

Video Conference tool. The tool included in the SHRIMP package for the transmission and reception of video streams.

WB

Whiteboard.