User Guide for

RAT v4.2

© University College London, Computer Science Department

Acknowledgements

This user guide was originally written for the UKERNA SHRIMP project.

Version history

RAT 3.0	Original version	Louise Clark
RAT 3.2	Revision 1	Kristian Hasler
RAT 4.2	Revision 2	Roy Bennett

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.

THE SOFTWARE AND DOCUMENTATION 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
Copyright © 1992-1997 Karl Lehenbauer and Mark Diekhans
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

2 Before you start 2

3 Getting started 2

3.1 Starting RAT from SDR or SCS 2

3.2 Starting RAT from the command line 2

3.2.1 Command lines for a multicast conference 2

3.2.2 Command lines for a unicast conference 4

3.3 Quick guide to using RAT 4

4 Using RAT 6

4.1 RAT main window 6

4.2 RAT options window 9

4.3 RAT participant windows 18

4.4 Command line features 20

5 Known bugs as of RAT 4.2.6 21

Further Reading 24

Glossary 26

Boxes

Box 1: Unicast and Multicast 1

Box 2: TTL and Admin scope 3

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

Box 4: If other people can’t hear you 5

Box 5: If you have problems hearing other people 6

Box 6: Silence Suppression 8

Box 7: Push to Talk 8

Box 8: Personal Details 10

Box 9: Transmission Details 12

Box 10: Reception Details 13

Box 11: Audio Details 15

Box 12: Codecs 15

Box 13: Security Details 16

Box 14: Interface Details 17

Box 15: 3D Positioning 18

Box 16: RAT Command Line Options 20

Box 17 RAT Examples 21

Box 18 RAT Packet Overheads 21

Figures

Figure 1: Main RAT Window 1

Figure 2: RAT Participant information panel - Personal Details 7

Figure 3: RAT Preferences: Personal Details 10

Figure 4: RAT Preferences: Transmission Details 11

Figure 5: RAT Preferences: Reception Details 13

Figure 6: RAT Preferences: Audio Details 14

Figure 7: RAT Preference: Codecs 15

Figure 8: RAT Preference: Security Details 16

Figure 9: RAT Preferences: Interface Details 17

Figure 10: RAT Participant information panel – Playout and Decoder 18

Figure 11: RAT Participant information panel – 3D positioning and Audio 19

1 Introduction

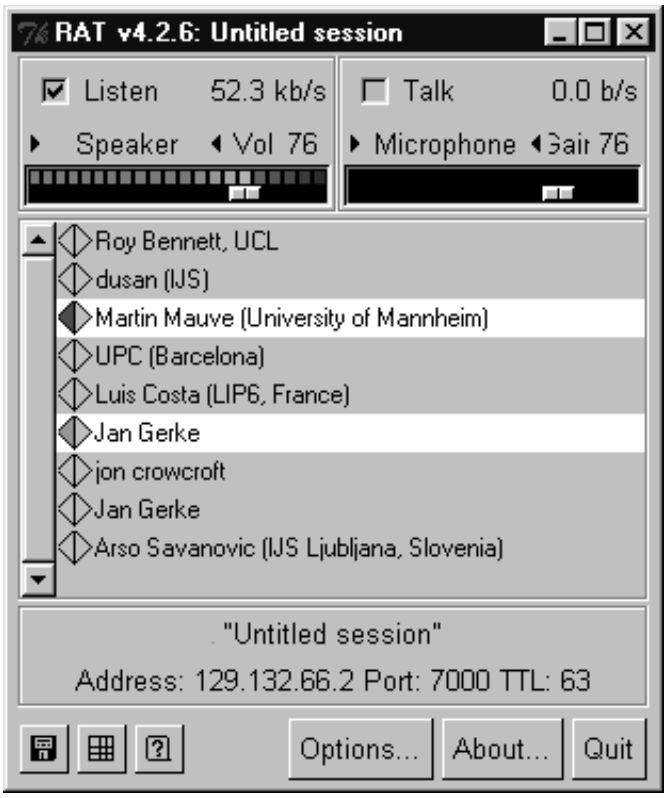

Figure 1: Main RAT Window

RAT (Robust-Audio Tool) is an audio conferencing tool. RAT can be used for either point-to-point (unicast) conferencing, involving a direct link between two computers, or for multiparty (multicast) conferencing when many participants in different locations are linked via a multicast-capable network (see Box 1).

Unicast and Multicast

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 to be used by the person with whom you want to have a conference.

A *multicast* connection can, in principle, have an unlimited number of participants. Multicast networking is available on Local Area Networks (LANs) and in parts of the Internet. You must use multicast if 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

2 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. You should also have a 'full duplex' soundcard installed.

% RAT only supports full-duplex sound cards. Check your sound card's instruction booklet to see if it has a full-duplex mode before using RAT.

It is recommended that headphones are used in preference to speakers. Using headphones allows you to listen and talk at the same time (this is known as full-duplex audio). If you do not use headphones, the output from the speakers may interfere with the input to the microphone, causing unpleasant feedback. The only other way of avoiding this feedback is to use an (expensive) echo cancellor.

3 Getting started

There are three ways of starting RAT: from the command line, from SDR and from the Secure Conference Store (SCS). Using one of the latter two is easier as they automatically configures the tool for use.

3.1 Starting RAT from SDR or SCS

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

3.2 Starting RAT from the command line

The commands vary with platform and conference type. A full list of the command line options is given in section 4.4 below.

3.2.1 Command lines for a multicast conference

The format for setting up RAT from the command line is as follows:

Unix:

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

Windows95/98/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>
```

Multicast addresses 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 an even

number and at least 1024. RAT uses both the port specified and also the next port.

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. As a rough guide a TTL of 15 will reach other parts of a campus, a TTL of 47 will reach all parts of a country in Europe or a state in the USA), a TTL of 63 reaches a continent and 127 world-wide. These values can differ depending on your location and you should confirm them by contacting your local network administrator.

The default for RAT is a TTL of 16. The maximum TTL allowed is 255.

Admin scope

Admin scoping assigns different ranges of multicast addresses that determine how far your audio will be transmitted. 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.

TTL scoping is still widely used, but you should contact your local systems administrator to see if admin scopes have been set up and, if not, you should ask for it to be done.

Box 2: TTL and Admin scope

3.2.2 Command lines for a unicast conference

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*:

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

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*. A port number can be specified but, if one is not, then RAT uses the default port 5002.

% TTL is not applicable in unicast sessions.

3.3 Quick guide to using RAT

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

In order to talk to other people, you must unmute your microphone. This can be achieved either by checking the box labelled *Talk* at the top right of the main window with your left mouse button, or by holding down the *right* mouse button whilst the pointer is in the RAT window. If you are using headphones rather than a microphone and speakers, you can leave the microphone unmuted all the time. If you are using a microphone and speakers (without an echo cancellor connected), the microphone should only be unmuted when you wish to talk.

You should now be able to use RAT for taking part in conferences, but be aware that you may encounter problems with hearing and being heard. These problems can result in your not hearing or not being heard at all; often the more difficult problems to solve are those which cause degradation of the sound. The most likely problems, their possible causes and the solutions are given in boxes 4 to 6 below.

If you can't hear other people:

Make sure that the *Listen* button is checked. If the *Listen* button is not checked, the incoming audio will be muted and the power meter will be plain black.

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 active and a name in the participants list is highlighted.

Check that your headset is plugged into the computer correctly.

Check the output volume setting. Make sure that the volume slider for the speaker or headphones is not too far to the left.

Check the output type setting. Make sure that the description of the output is correct *Headset* if using a headset, and *Speaker* if using speakers. You can change between the output types available by clicking on the small triangles on either side of the description. (Some audio cards will automatically send audio to headsets if they are plugged in and speakers if not. In this case, the description will not toggle, but will always display *Speaker*)

Check that you are not running any other audio programs. If you are, you should exit from them and re-start RAT.

Check in the 'audio' options that your audio device is selected See Box 11: Audio Details.

Check sample rate of sender. If other participants are using a different 'sample rate', then check in the 'Reception' options that a sample rate converter has been selected (see Box 10: Reception Details).

Box 3: If you can't hear other people at all

If other people can't hear you:

Make sure that your microphone is not muted when you talk. The microphone is muted when the *Talk* button in the RAT audio window is checked and the powermeter is plain black.

If you have started rat from the command line, make sure that the TTL is high enough (see Box 2: TTL and Admin scope). If your TTL is too low, you may still be able to hear a remote participant even though they cannot hear you. To change the TTL you will have to restart RAT.

Box 4: If other people can't hear you

If you have problems hearing other people:

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 slider bar (beneath the word Microphone in the input segment at the top right of the main window). With most audio cards, you will not hear yourself in your headset or through the speakers when you are talking; if you do hear yourself in this way, you can turn off this effect. Whether you can hear yourself or not, you will have to rely, to a certain extent, 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 volume level will cause you to sound distorted. Generally, you have to rely on the other participants in the conference to tell you whether your volume is set correctly. The numeric indication (on a scale from 0 to 100) labelled *Gain* can be useful for indicating to others the level at which you are transmitting.

If the audio from another participant is very broken up. The most likely causes of this are either network congestion causing the loss of their data packets, or the other user speaking too quietly for the audio tool to distinguish them from the background noise. You can tell if you are losing packets from another participant by looking at the diamond to the left of the participants name. If the diamond is green (loss less that 5%) or orange (loss less than 10%), then the packet loss is small so ask them to turn their volume up. If that does not help, ask them to turn off silence suppression. Tthis option is available from the *Options* menu, accessible from the RAT main window (see Box 11: Audio Details).

If everyone's audio keeps cutting out. A possible cause is your moving windows around on the screen. When you do this, the processing power of the computer is spent on window management rather than audio decoding, and you lose audio quality. Try not to move windows around unnecessarily.

Box 5: If you have problems hearing other people

4 Using RAT

RAT consists of a *Main window* (see Figure 1), a *Preferences window* (see Boxes 6-10), and a number of *user-info* windows (see section 1.3.3 below). The following sections discuss the Main window and the Preferences window in greater detail. The different command line options are described at the end of the section.

4.1 RAT main window

The Main Window (see Figure 1) of RAT is split into three sections. In the middle is a list of conference participants. Your own name will always appear at the top of this section. Names are highlighted whenever anyone speaks. Clicking on the name of a participant, including your own, will display a Preferences panel, with the category *Personal Details* selected (see Figure 2).

In this panel, you can display other technical details by selecting other options on the Category menu: *Playout, Decoder, Audio* and *3D Positioning* (see section 4.3 below for details).

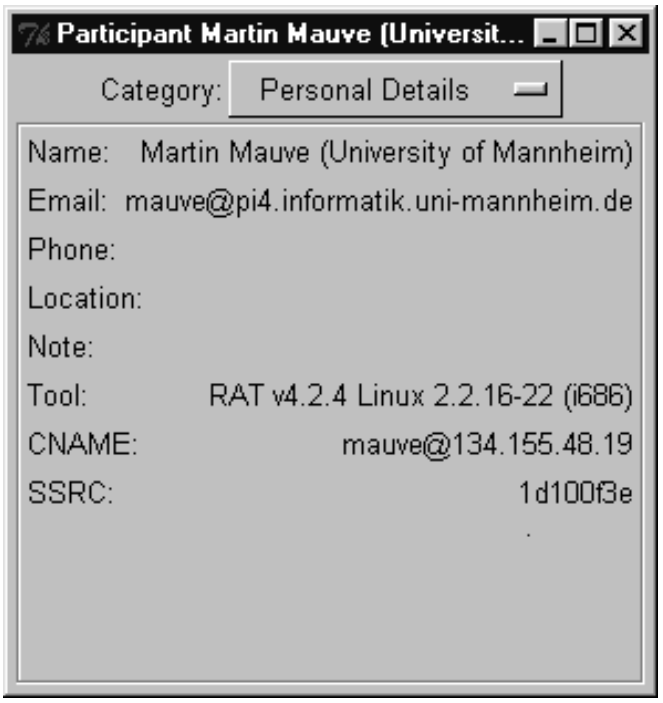

Figure 2: RAT Participant information panel - Personal Details

The top section of the main window is divided into sections, input to the left and output to the right. Checking the boxes labelled *Listen* and *Talk* will enable the devices shown below them. If you do not check the appropriate box, you will be unable to hear and/or to be heard. Individual participants can be selectively muted by clicking on their names with the middle mouse button or, on machines that have a two button mouse, by holding down the control key (on the keyboard) and left clicking on their names.

By default, the output device is *Speaker* and the input device is *Microphone*. Depending on your hardware configuration, *Line in* and *CD* may be available for input *and Line out* and *Headset* for output. To change between these options you click on the small triangular buttons on either side of the device description.

Beneath the input and the output device is a slider to set sound levels which also gives a numeric indication on a scale from 0 to 100 next to the device indicators. This numeric indication is useful when setting appropriate levels with other users at the start of a conference, providing a precise indication of your level settings. Immediately above the slider is a graphical indication of the output volume and the input gain.

In the bottom section of the window is the session name, the conference address and port number, the TTL and six control buttons. From left to right, these buttons have the following functionality:

1. File control when playing from or recording to files
2. Enabling display of the reception quality matrix
3. Enabling the balloon help facility
4. *Options* Displaying the various preferences which can be set (see section 4.2 below)
5. *About* Shows credits, copyright information and feedback methods
6. *Quit* Closes RAT and ends the session.

Silence suppression

By default, RAT starts up with silence suppression turned on. The silence suppression attempts to ensure 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 the silence suppression feature will ensure that background noise is not 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 disable silence suppression in the RAT Control Window (see Figure 6: RAT Preferences: Audio Details), and use push-to-talk (see Box 7: Push to Talk).

Box 6: Silence Suppression

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, 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 *Talk* button check. This is why is it called push-to-talk – because you “push” the right mouse button when you talk.

Box 7: Push to Talk

When checked, the *Listen* and *Talk* buttons mute incoming and/or outgoing audio. You can mute the outgoing audio, i.e. microphone or line out, by holding down the right mouse button with the pointer anywhere within the RAT Window. This muting will continue only while the mouse button is depressed.

4.2 RAT options window

Pressing the *options* button in the Main RAT Window brings up a Preferences window, allowing the operation of RAT to be modified. There are several categories of options that can be altered and you can browse through them by using the *Category* Menu. Options can be set for the current session by clicking on the *Apply Preferences button* or kept for future sessions by clicking on the *Save & Apply Preferences* button. Clicking on the *cancel* button will cancel any changes made.

These are some of the options you may need to change often:

- 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 6: Silence Suppression). If you turn silence suppression off, you should mute your microphone whenever you are not speaking, otherwise you will be transmitting background noise continuously (see Box 7: Push to talk).

- 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, allowing more time for packets to be received. *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. 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.

For a thorough description of the options available in the RAT Preferences Window, see Boxes 6 -10.

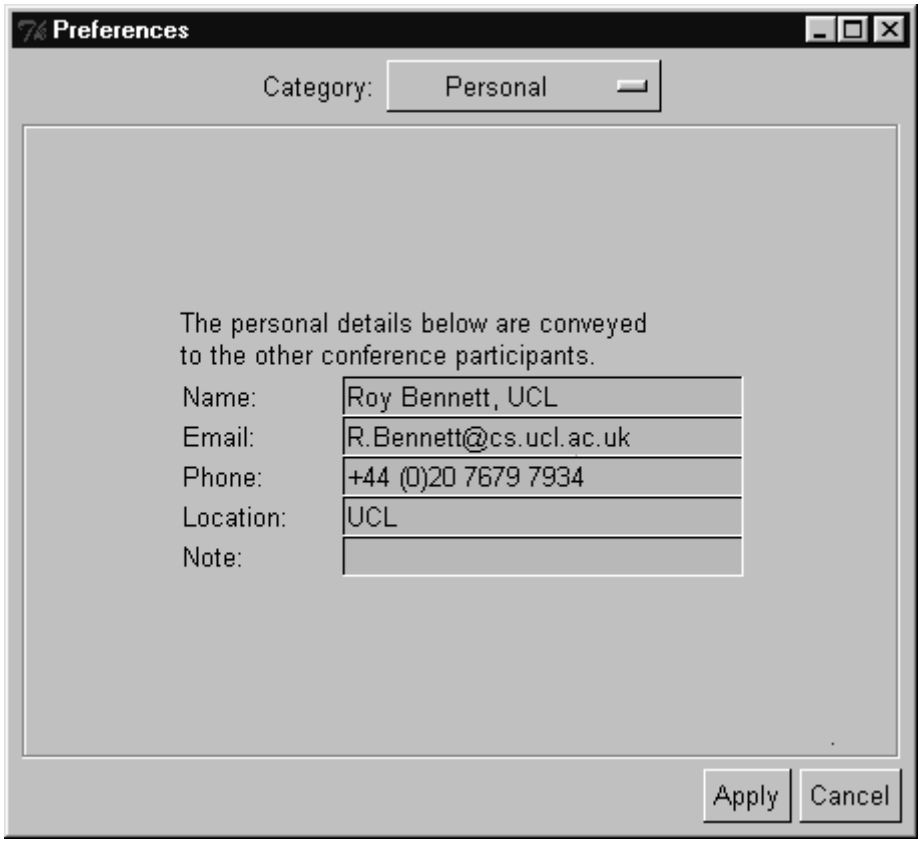

Figure 3: RAT Preferences: Personal Details

The details that you enter here will be conveyed to all other conference participants and can be recalled by them by clicking on your name. The information is sent periodically and may take some time to update at all receivers.

Name
Enter your name here.

Email
Enter your Email address here.

Phone
Enter your telephone number here.

Location
Enter the location of your site here.

Box 8: Personal Details

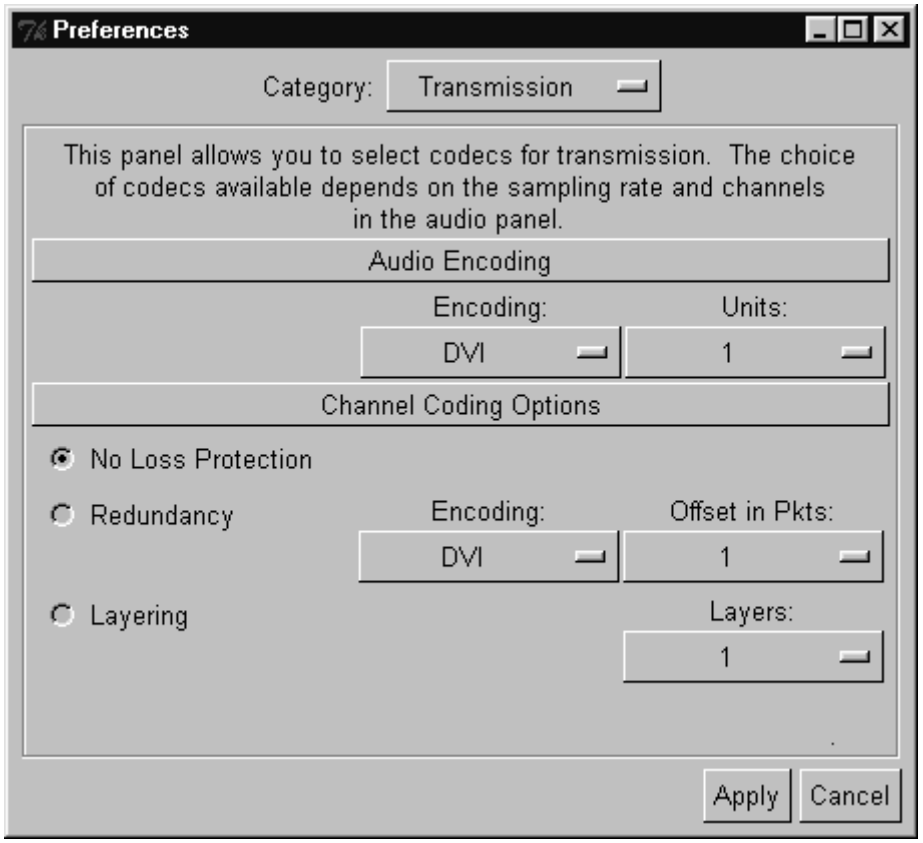

Figure 4: RAT Preferences: Transmission Details

Encoding

Sets the format for data transmitted to the network. The different types of encoding use different amounts of bandwidth. Using a 8Khz sampling rate these are:

Linear16: 128 kbit/s: Uncompressed 16 bit samples, with linear spacing between sample values.

A-Law: : 64 kbit/s: a-law companded 8 bit samples. Very similar to Mu-Law but only supported by RAT (i.e. not VAT).

Mu-Law: 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. Quality is very similar to L16 (the difference is imperceptible with most soundcards/headphones).

DVI-ADPCM: 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 marginally worse than Mu-Law, but is generally acceptable. Probably the best compromise between processing requirements and bandwidth.

GSM: 13.2 kbit/s: This encoding is 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. It works by the same principle as GSM but less information is calculated and sent. Should only be used if bandwidth is *very* limited.

Loss Protection (Redundancy and Interleaving)

Two types of loss protection are possible: redundancy and interleaving. If the conference includes any users of other audio tools that don't support redundancy, redundancy should be turned off i.e. *No Loss Protection* should be set. However if all users are using RAT, and packet loss is being experienced, the use of a loss protection will vastly improve sound quality.

Redundancy works by 'piggy backing' a low quality duplicate sample onto the next packet, so that if a packet is lost then the next packet has a backup sample as a replacement. You can specify how to encode the backup sample and also set the offset from the original packet the backup sample should be sent in. This type of loss protection is also supported by FreePhone from INRIA.
Interleaving works by interleaving the order in which the packets can be sent. This tries to prevent against large gaps in network loss. The size of the packets to interleave and the separation between each interleave can be altered using *Units* and *Separation* respectively. This should only be used with other participants using RAT3.2 or above.

Box 9: Transmission Details

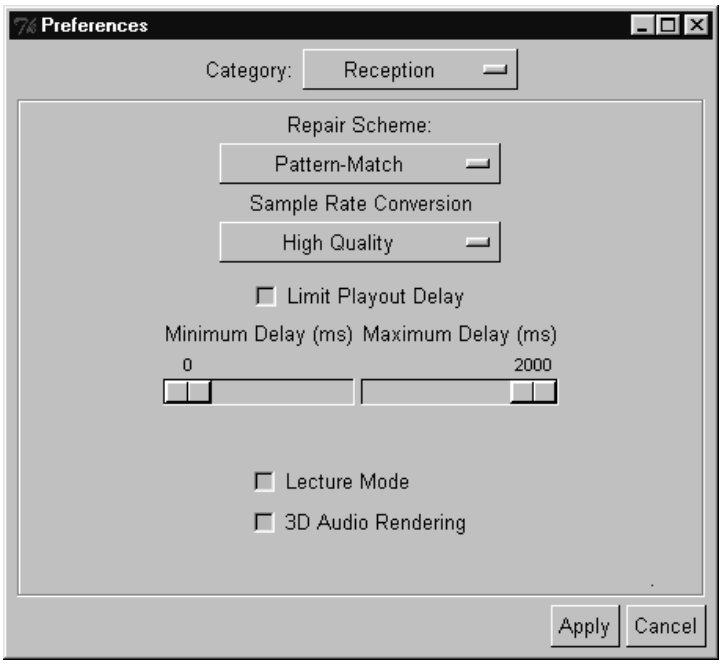

Figure 5: RAT Preferences: Reception Details

Repair Scheme:

There are five options, none of which affect the data sent to the network.

Pattern-match repeats a suitable segment of audio from before the loss occurred and with no discontinuity in phase. This sounds better than packet repetition but involves more calculation.

Repeat patches over the gap left by a missing packet with the previous packet.

Repeat+ similar scheme, but more sophisticated

Noise introduces a random pattern

None means that no attempt is made to repair losses.

Sample Rate Conversion

Allows the conversion of audio sources which are not using the sampling rate supported by the sound card. This requires significantly more processing power. The options are: high, medium and low quality and the Microsoft converter.

Playout Delay

You can limit the playout delay within the range 0 to 2000 ms.

Lecture Mode

If *checked*, 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.

3D Audio Rendering

If *checked*, allows you to alter the sound for each participant, giving the effect that they are virtually placed around you when using headphones. See also section 4.3.

Box 10: Reception Details

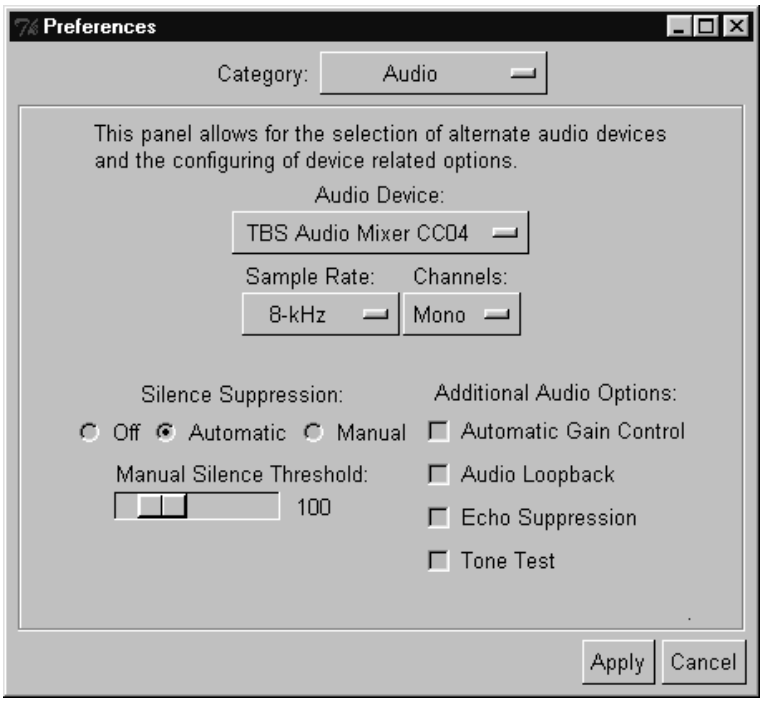

Figure 6: RAT Preferences: Audio Details

Audio Device

Use this to change between multiple audio devices. If 'No Audio Device' is selected then no audio can be recorded or played out. This will happen automatically if other software is using the audio device when RAT is started.

Sample Rate

Changes the sampling rate. Higher sampling rates improve the sound quality but also require greater network resources. When using higher sampling rates, the encoding schemes available are limited to ones that support that sampling rate.

Channels

Allows you to change between mono and stereo.

Silence Suppression

This feature ensures that periods of silence within a conversation (including background noise) will not be transmitted, hence reducing network traffic. If you leave the default setting, *Automatic*, RAT will take care of it. If you have problems with the automatic setting, you can switch to *Manual* and increase the silence threshold using the slider. It is recommend that you use some method of silence suppression (see also Box 6: Silence Suppression)

Automatic Gain Control

If *checked*, the microphone gain will be controlled automatically, i.e. the volume will remain relatively level despite your shouting or speaking quietly.

Audio Loopback

If *checked* and if the soundcard supports this option, the audio is looped back so that you hear your own audio.

Echo Suppression

If *checked*, RAT will suppress (i.e. not transmit) your audio whilst another participant is active. This prevents problems with feedback or echo when using headphones.

Tone Test

This option is for research use only.

Box 11: Audio Details

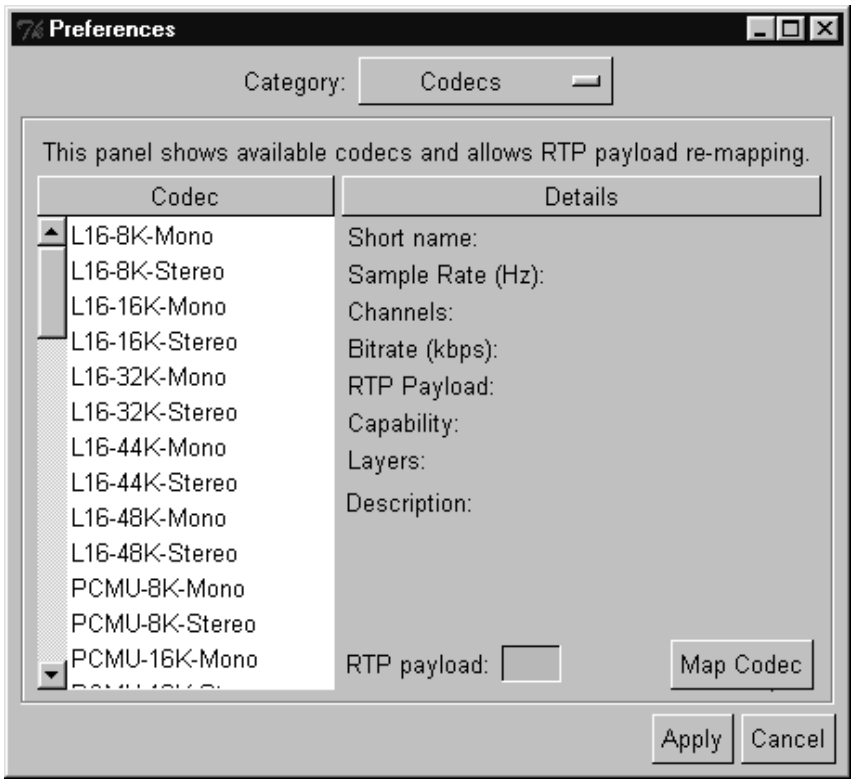

Figure 7: RAT Preference: Codecs

Codec/Details

When you click on a codec in the window, its details are displayed. To use the selected codec, click on *Map Codec*

RTP payload

You can remap the RTP payload using this feature. It is not often necessary to do this.

Box 12: Codecs

Figure 8: RAT Preference: Security Details

Key

This is the key used to encrypt the audio data using DES. The key can be any combination of letters and number but you must have the same key as those you wish to communicate with.

Enabled

When *checked*, encryption is enabled and your communication with others will be limited to those having the same key.

Box 13: Security Details

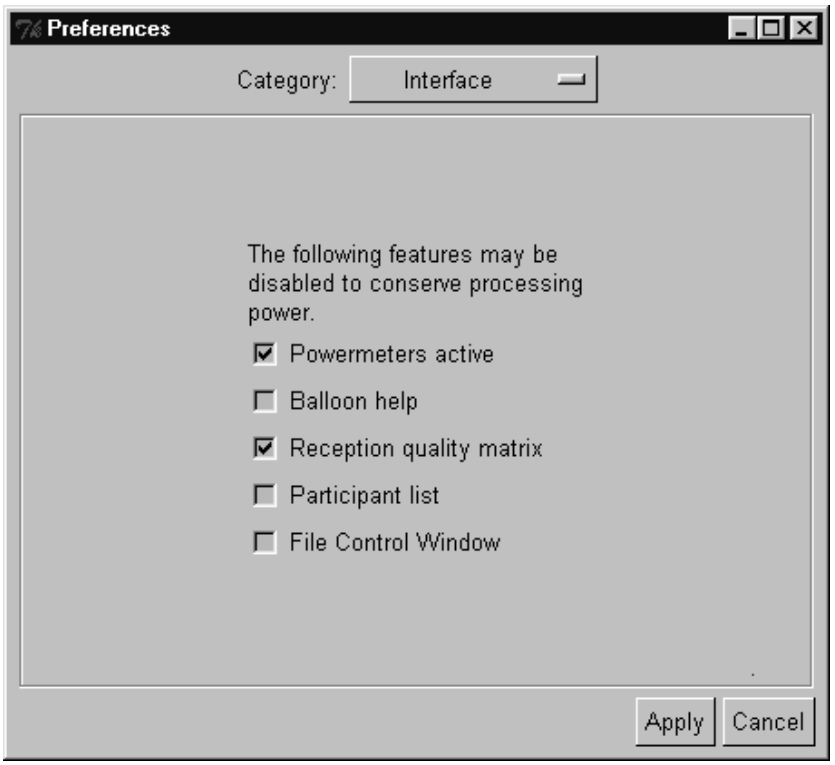

Figure 9: RAT Preferences: Interface Details

Powermeters active

If *checked*, audio power-meters will be displayed in the main RAT window. Disabling power-meters is not usually necessary unless you are on a very slow machine.

Balloon help

If *checked*, help text is displayed in small balloons when you pass the mouse over buttons. This can be controlled also with the icon on the main RAT window.

Reception quality matrix

Toggles a window that displays a matrix based on each participant and the quality of the audio received from them. This can be done also with the icon on the main RAT window.

Participant list

Toggles the participant list in the main RAT window.

File Control Window

Toggles a window that allows you to play and record sound samples. To play or record a sample you first need to click on the relevant disk icon to select the file. You can then use the buttons, much like a standard CD player, to control the sample. This can be done also with the icon on the main RAT window.

Box 14: Interface Details

4.3 RAT participant windows

The Playout and Decoder information options show technical details which may be of use in the event of problems, but will not normally be of interest to users.

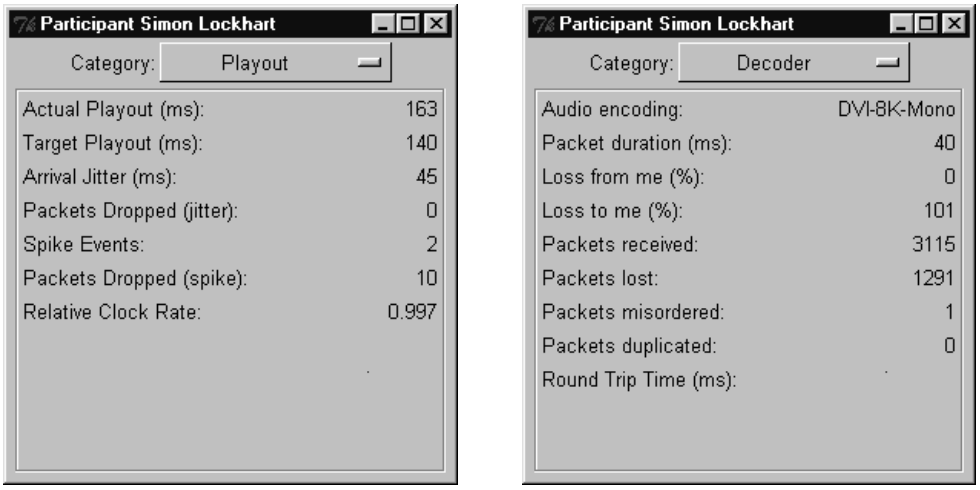

Figure 10: RAT Participant information panel – Playout and Decoder

The 3D positioning and Audio options allow the user to change RAT settings. The images in Figure 11 below show the default settings.

If you activate 3D positioning, by checking the appropriate box, you can then select the further options:

Filter type

HRTF (Head Reflected Transfer Function) is the default filter and should be the one you use. It models the signal to simulate the normal hearing experience where sound is reflected off ears and shoulders.
Identity provides the basic signal and should be used if HRTF gives any problems.
Echo is a filter which is being developed in a research environment and should not be selected.

Length

This is a feature which, when implemented fully, will enable users to adjust quality; at present it is fixed at 32.

Azimuth

This setting allows you to set individual participants on a virtual half circle in front of you. The scale runs from -90 representing your extreme left to +90 on your extreme right. Zero is the position straight in front. If you are viewing individual video images of the participants, you can set the sound to appear to come from the same relative positions as the images on the screen.

Box 15: 3D Positioning

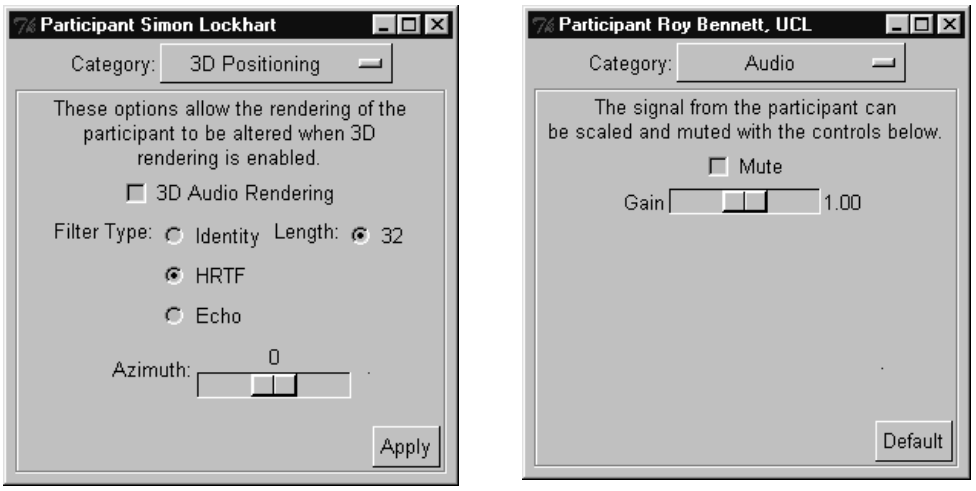

Figure 11: RAT Participant information panel – 3D positioning and Audio

The audio category allows you to adjust the volume (gain) for yourself and for other participants in the conference. You can also silence individuals by checking the mute button, an option available from the main window also.

4.4 Command line features

RAT Synopsis

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

The following command line options are supported:

-**crypt** *key* or **–k** *key*
Enable encryption, with the specified key. Encryption is done using DES.

-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. See Box 12 for description of the CODECs and their recommended usage.

-allowloopback
Causes RTP packets to be looped back during multicast sessions. Mainly used for debugging.

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

-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, 64 continent, 47 country and 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). The option is useful, for example, for transcoding between low- and high-bandwidth sessions, over a slow link.

-version
Displays the version number of RAT.

Box 16: RAT Command Line Options

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 224.2.2.2 and port 8000 type, using primary encoding of DVI and a secondary encoding of LPC:

```
rat -f dvi/lpc 224.2.2.2/8000
```

Box 17 RAT Examples

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).

Initially RAT starts with '2 units' (represents 40ms) packets of DVI coded data at 8Khz. This represents a trade off between bandwidth, quality, and the ability of receiver based repair techniques to work successfully.

Box 18 RAT Packet Overheads

5 Known bugs as of RAT 4.2.6

There are a number of known bugs and a list of desirable improvements for this version of RAT. Those that are relevant to users are listed here. The full list is included in the documentation with the software at the end of the file named *MODS* under the heading *TODO.*

Assorted audio driver problems

- FreeBSD driver bug on SoundBlaster 16 has small write buffers and stops working full duplex mode after a time. Hard to fix because of cushion.
- Crashes on FreeBSD 4.0 (reported by Lars Eggert and Joe Love). Due to bugs in FreeBSD

- Broken auddev_pca with adding of audio interface conversion code. FreeBSD 3.1 pca audio does not seem to work anymore.

- SunVideoPlus interface code does not work properly. The driver Sun distributes only works in 8k mono. MMAC produce the card and wrote a driver for Sun: an exceptionally buggy and unstable driver. MMAC now has a better version, but expect you to pay for the developer kit to access features. $1200 for a card that does not perform as advertised. The marketing blurb states it has G728, G722, and "CD Quality" audio: it may do, but you can't access them.

- SGI O2's don't have autosensing audio interface selection like Indy's.

- Thierry Turletti still has problems with the NM256 card and device pairing

- Maestro2 cards only support stereo recording; you need to hack auddev_oss to deal with this (Joseph Eggleston)

Disabling participant list doesn't work if the window has been resized: not sure if this is a feature of Tk or the window manager, but there's little we can do about it.

Layering is broken (it hasn't been updated for multiple processes, etc)

Weird interaction with rat-3.0.x - seems that we can generate packets which some instances of rat-3.0 won't play. You get some people in the session whose tools report receiving your data, but they don't play it out. Suspect we're hitting the limitations of the old playout algorithm, rather than a specific bug, and the correct solution is to get everyone to update to rat-4.x

Sometimes see messages of the form "source.c:767 Spike (160 > 2708390)" and hear nothing.

If a receiving stream contains a lot of jitter and out-of-order packets it crashes sometimes. (jitter around 1000 ms and random out-of-order). Reported by Peter Parnes.

Send doesn't always send. Sometimes, when you turn on the mike, nothing is sent. The mike volume bar doesn't move. Turning on and off 1 to 3 times can help. Reported by Peter Parnes.

When RAT starts, it turns off "playback-mute" on the mike input which gives feedback for the mike. Muting it via the audio control turns off the feedback. Note that record-mute and playback-mute is not the same thing. Reported by Peter Parnes.

Request outstanding to allow the changing of RTP sessions whilst RAT is running. Requested by many people (Robert Olson, Joerg Ott, etc).

Crashes on WinNT if another version of Tcl/Tk is installed on the machine? Reported by Didier Lamballais <Didier.Lamballais@irisa.fr>.

Occasional crashes due to "too many blocks allocated" with many allocations in channel_types.c:29, channel_types.c:34, playout.c:188 and source.c:893. Seems to be due to packets being received off the network, but never played out. Reported by Angel Mateo, and also observed at UCL. Not yet reproducable.

Problems with the configure script on Irix 6.2. This works for some users, and not for others. Must be something in the environment, or a software version issue but we don't know what...

Excessive end-to-end delay on Windows? Reported by Gerry Wolfson.

Bob Olson reported:
While changing settings from 48khz GSM to 16khz L16, I got this:

```
rat-4.2.4-media: source.c:1588: source_process:
Assertion `md->nrep < 5 && md->nrep > 0' failed.
```

FYI, I just triggered this same failure inadvertently at another site by transmitting at 48khz when they were set to receive at 16khz.

Occasional weird behaviour of the playout buffer, where sources are buffered for several _minutes_ before being played out? Reported by Steven Dawson on the CAIRN list.

Further Reading

2000

- O. Hodson, C. S. Perkins & V. Hardman, Skew detection and compensation for Internet audio applications, Proceedings of the IEEE International Conference on Multimedia and Expo, New York, July 2000.

1999

- A. Bouch & M. A. Sasse, Network Quality of Service: What do users need?, Proceedings of the 4th International Distributed Conference (IDC'99). 21-23 September, Madrid, Spain 1999
- L. Lambrinos, P. T. Kirstein & V. Hardman, Improving the Quality of Recorded Mbone Sessions using a Distributed Model, to be presented at IDMS, 1999
- V. Hardman & O. Hodson, "Speech and Music Transmission over the Best-Effort Internet" to be published in Handbook of Communications Technologies: The Next Decade, published by CRC Press LLC.

1998

- C. S. Perkins, O. Hodson & V. Hardman, A Survey of Packet-Loss Recovery Techniques for Streaming Audio, IEEE Network Magazine, September/October 1998.
- I. Kouvelas, V. Hardman & J. Crowcroft, Network Adaptive Continuous-Media Applications Through Self Organised Transcoding, Proceedings of Network and Operating Systems Support for Digital Audio and Video (NOSSDAV 98), 8-10 July 1998, Cambridge, UK.
- C. S. Perkins & O. Hodson, Options for Repair of Streaming Media, IETF Audio/Video Transport Working Group, June 1998, RFC2354.
- V. Hardman, M. A. Sasse & I. Kouvelas, Successful Multiparty Audio Communication over the Internet, Communications of the ACM, Vol. 41(5), 1998.

1997

- I. Kouvelas, O. Hodson, V. Hardman & J. Crowcroft, Redundancy Control in Real-Time Internet Audio Conferencing, Proceedings of the 1997 International Workshop on Audio-Visual Services Over Packet Networks 15-16 September 1997.

- C. S. Perkins, I. Kouvelas, O. Hodson, V. Hardman, M. Handley, J.-C. Bolot, A. Vega-Garcia & S. Fosse-Parisis, RTP Payload for Redundant Audio Data, IETF Audio/Video Transport Working Group, September 1997, RFC2198.

- C. S. Perkins, V. Hardman, I. Kouvelas & M. A. Sasse, Multicast Audio: The Next Generation, In Proceedings of INET'97, Putra World Trade Centre, Kuala Lumpur, Malaysia, June 1997.

- C. S. Perkins & J. Crowcroft, Real-Time Audio and Video Transmission of IEEE Globecom'96 over the Internet, Presented at the IEEE Engineer's Conference at Networld+Interop, Las Vegas, May 1997.

- C. S. Perkins & J. Crowcroft, Real-Time Audio and Video Transmission of IEEE Globecom'96 over the Internet, IEEE Communications Magazine, Volume 35, Number 4, April 1997.

- I. Kouvelas & V. Hardman, Overcoming Workstation Scheduling Problems in a Real-Time Audio Tool, In Proceedings of the USENIX Annual Technical Conference, Anaheim, January 1997.

1996

- V. Hardman & M. Iken, Enhanced Reality Audio in Interactive Network Environments in Proceedings of the FIVE Technical Conference, Pisa, Italy, December 1996.

- I. Kouvelas, V. Hardman & A. Watson, Lip Synchronisation for Use Over the Internet: Analysis and Implementation, Proceedings of IEEE Globecom'96, November 1996, London UK.

1995

- V. Hardman, M. A. Sasse, M. Handley & A. Watson, Reliable Audio for use over the Internet, In Proceedings INET'95, Hawaii, 1995.

See also:

- http://www-mice.cs.ucl.ac.uk/multimedia/software/

- http://www.ip-multicast.com

- RTP spec: RFC 1889/1890

Glossary

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

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.

GSM
: Audio encoding standard.

IP address
: IP (Internet Protocol) addresses are unique numerical identifiers for each networked host computer.

Kilobits per second (kbit/s)
: The kbit/s rate enables the participants to track how much bandwidth is being used during a session.

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.

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 a tool such as SDR.

Network congestion

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.

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.

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. See RFC 2198.

RTP

Real-Time Protocol is an application level protocol that provides real-time support for multimedia communication. See RFC 1889/1890

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.

TTL

TTL stands for time to live, and determines how far multicast packages can travel over the Internet. Roughly the values can be interpreted as:
16 - Site
47 - UK
63 - Europe
127 - World

Unicasting

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