

MAVEN

These are some release notes for *Maven*, an Internet audioconferencing tool for the Macintosh computer. This is alpha-quality code, and this is version 2.0d21.

Changes in version 2.0d21:

- The most significant change is that Maven is now a fat binary application. Donations were received. :-)
- Some minor changes have been made to playout delay calculation. This will continue to evolve over the next few months.
- This is development-quality code and perhaps somewhat more likely to crash than the previous release (2.0a18). If problems occur, please report the nature of the problems and use 2.0a18 as a fallback position.

Changes in version 2.0a18:

- Maven should work on Powermacs now. Thanks to Jim Browne for letting me sit at one long enough to verify sort of what the problem was. It still has not been extensively tested on the PPC architecture or compiled natively for the PPC, and won't until I get one on my desk. Donations welcome. :-)
- Maven no longer forces you to dismiss the configuration dialog at startup; instead it comes up with the defaults and goes directly to the main window. A new Edit menu entry "Configuration..." brings up the configuration dialog. Changing either the network port or conference ID causes Maven to remove all current conference participants, because it has to close and re-open the network connection and the participants on the old port and ID would see you "go away" anyway. Changing only "Local Site ID" or "Ignore Unsolicited Streams" will make the change dynamically in the current conference. Note that a change in the Local Site ID results in the change appearing to other participants after a few seconds.
- Tired of fiddling with frustratingly CPU-intensive DSP algorithms like GSM and CELP and trying to cram them onto a 68040, I opened up a digital communication textbook and read up on something called Delta Modulation. I've implemented a codec for delta-mod, which uses only 16 kilobits per second of network bandwidth and doesn't require any fancy floating-point either. This *may* be in reach of SLIP connections, which is what a lot of people have been asking for. Quality is somewhat poor; there's a great deal of quantization noise in the received audio which sounds like an AM radio station tuned slightly off, but the speech is quite intelligible. Comments about this coding are welcome.
- Some people reported that sometimes the network audio was so quiet that it wasn't really intelligible even with the sound volume turned all the way up on the receiving Mac. To address this I've added "Louder" and "Softer" commands on the Audio menu which cause Maven to adjust the signal amplitude before it is sent to the Sound Manager. You should use these controls only if you have set the volume in the Sound control panel all the way up and it's still too quiet.
- I'm aware of the "freeing free audfrag" bug and am still tracking that one down.

Changes in version 2.0a16:

- Fixed bug where if a session was terminated and then reestablished, Maven would not recognize it due to a cache problem (Brian Lloyd).
- Fixed bug where audio received from *vat* mixers was getting ignored.
- Added a new Audio menu entry, "Antialias," which allows you to control whether or not Maven will oversample the receive audio. Turning it off can result in annoying buzzing due to Nyquist aliasing, but turning it off can also gain you a big improvement in CPU performance. Antialiasing is on by default.
- Maven will now run if no sound input device is found. At least it's good to listen in on MBONE conversations via a *vat* mixer this way.

Changes in version 2.0a15:

- Control of audio encoding, quantization, and the push-to-talk vs. squelch-controlled mode are now on an “Audio” menu that comes active after the main menu appears. These used to be in the configuration dialog.
- A new command, “Mute Speaker” also appears on the Audio menu. When checked, no network audio will be played out to the Sound Manager.
- Further improvements in the playout algorithm. It computes playout delay dynamically based on delay variance. It seems to work pretty well, but I’ll be interested in anyone’s experiences with unusual long speech delays or long gaps in the audio.
- An even coarser audio quantization (160 milliseconds) is now available. This is nice for sending things like lectures because the packet rate is quite low. Be careful, though; with this it’s easy to send packets larger than some pieces of the Internet can handle without fragmentation.
- (Some) session control. The main window now includes a list containing the names of the participants in the audioconference. The names of the participants are determined from ID messages coming from the other participants. New participants can be added to this list by choosing “New...” from the “Sessions” menu. In this case the list entry temporarily contains what you connected to, preceded by question marks, until an ID message is received from the connection partner. Participants can also be added by an unsolicited ID message being received when “Ignore unsolicited streams” was left unchecked in the config dialog. Maven automatically begins sending audio packets to each new participant no matter how they were added to the list. Participants get removed from the list (and audio stops getting sent to them) when they quit their audioconferencing tool, be it Maven or *vat*. There is still no way to close a single session in the list short of quitting Maven (unless the other participant quits). You can click on the list participants with the mouse, and they’ll highlight, but this doesn’t do anything interesting.

This session control is still somewhat cheesy as it does not deal with *vat* mixers very well (doesn’t handle the IDLIST message), you still can’t tell who’s speaking, and there’s no way to manipulate the list. The code that implements it is also full of race conditions. I have a student working on this, and I have high hopes that I will soon have some session control at least as useful as what *vat* does.

- A new creator signature and a cooler icon. If you can do better, send me your improvement in the form of ‘icl8’ and ‘ics8’ resources.
- Balloon help for the menus and the windows (ooh, ahh). Actually some of the menu help provides a short tutorial on the tradeoffs of the various audio encodings and packet sizes, so you might want to check it out by turning balloon help on in the help menu.

If you have a development version of Maven (use the Finder’s “Get Info...” command to retrieve the version number—development versions include a letter ‘d’, as in 2.0d1), please do everyone a favor and throw it away. They were very incomplete and could easily crash. They also had some network behavior that was rather unfriendly.


A User's Manual (Such As It Is)

I don't have any real documentation written yet, sorry. I'll try to outline some of the features of Maven here. I'm willing to answer any questions you might have about this version. To keep up on future versions and discussions of Maven, please join the Maven discussion list, maven@cnidr.org. To join, send electronic mail to listserv@cnidr.org containing the single line `sub maven Your Name` substituting your full name for *Your Name*.

If you're on a non-AV Macintosh and you don't have version 3.0 of the Macintosh Sound Manager already, you should install it. It incorporates a lot of efficiency improvements over what you probably have in your ROMs. Sound Manager 3.0 is included in this distribution; it belongs in the Extensions folder in your System Folder. You'll have to restart to make it take effect. You may have Sound Manager 3.0 already if you've installed one of the hardware system updaters from Apple. I've heard that you don't need it on AV's and Powermacs, and that it might even make them crash.

Maven requires a Macintosh with sound input hardware to send audio to the network. It will run on Macintoshes without sound input hardware, but will only be able to receive other audio packets in this case. I haven't tested it on a lot of different Macintosh models, but I have noticed that many cannot input and output audio at the same time. Maven does the best it can with this shortcoming, but doesn't yet check for it at startup. If you find that no sound comes out of your speaker when not in Push-To-Talk mode, you'll have to run in Push-To-Talk mode at all times. Since the squelch-detect software requires that the sound input channel be open at all times, and if your Macintosh can't input and output at the same time, no sound output channels can be opened and no audio can be produced.

While Maven is running, it keeps its main window up at all times.



On the right hand side of the main window, a vertical bar graph shows the current sound input level (in push-to-talk mode this is only active when the mouse button is down). Below the bar graph is an icon for the sound input device. This will be displayed, as above, in reverse video when Maven is actually transmitting audio packets to the Internet.

If you're operating in squelched mode (see below), a small triangular slider will appear to the left of the bar graph. You can drag it with the mouse to the point where you want the silence detector's threshold to be. When the sound level at the microphone causes the bar graph to move above the slider, Maven will start transmitting. When it falls below, Maven will stop after about a half-second *squelch tail*.

When Maven starts up, it's not sending audio to anyone. To begin sending to a peer, whether another copy of Maven or a *vat* program on a Unix system, you can select "Create New Session..." from the File menu. You'll get a dialog box in which you can type the hostname or IP address of the other site. When you press return, Maven will send an audio stream to that site. Alternatively, if "Ignore unsolicited streams" was left unchecked, you can just wait for a peer to connect to you, and the session will be created automatically.

Warning! Maven will let you continue to add new sessions indefinitely. The more sessions you've added, the greater the load you are placing on the Internet as a separate copy of your transmitted audio is sent to each site you've added. Take special care when sending to more than one other site that you don't swamp the network with packets!

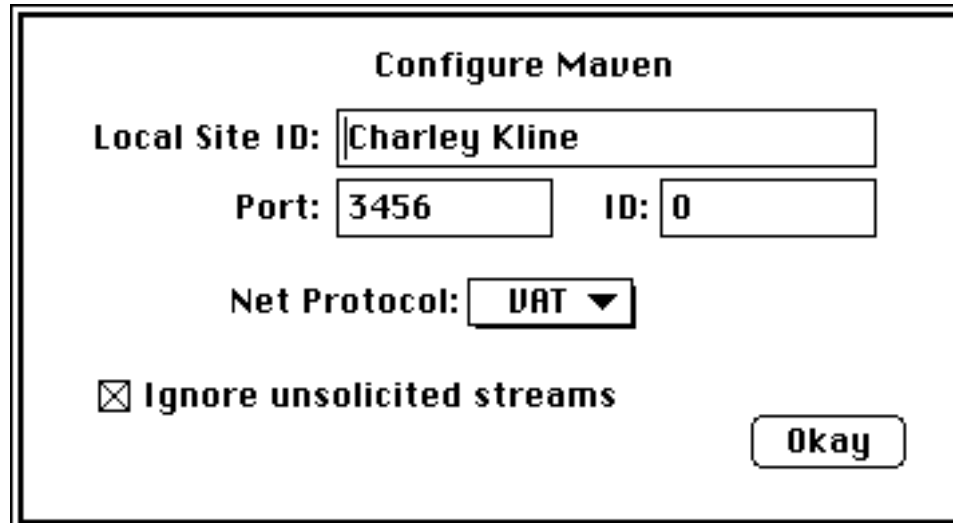
To the left of the input controls is a list containing the list of participants in the conference. Your name (as it appeared in the configuration dialog) always appears first on the list. Other names will appear as new participants are added.

A new participant appears in the list (and Maven begins sending audio to it) when you add a new session manually, or if you left "Ignore Unsolicited Audio" unchecked in the configuration dialog and an ID message arrives from another machine. In the latter case, Maven will automatically create the new session, add it to the list so you can see that a new person has joined the conference, and begin sending your audio to them.

There's currently no way to delete a session once it has been created; this will be a future enhancement. However, if a participant in a conference quits their conferencing tool (Maven or *vat*), they will be removed from the participant list of the other members of the conference.

There are several things you can do to modify Maven's behavior

when the main window is visible. To change some of the global settings, select **Configuration...** from the **Edit** menu.



The screenshot shows a dialog box titled "Configure Maven". It contains the following fields and controls:

- Local Site ID:** A text box containing "Charley Kline".
- Port:** A text box containing "3456".
- ID:** A text box containing "0".
- Net Protocol:** A dropdown menu with "UAT" selected.
- Ignore unsolicited streams**
- Okay** button

This dialog tells Maven what mode to operate in and how to exchange digital audio over the Internet.

Local Site ID is simply the identification that is sent with your audio stream on the network. Maven sends this out periodically along with your audio. Both Maven and *vat* use these ID messages sent to fill in a list of participants. Maven also uses ID messages to determine whom to send audio packets to. This field takes its default value from the Owner Name you typed into the Sharing Setup Control Panel.

Port is the UDP port number which Maven will send on. 3456 is the *de facto* Internet default, but it doesn't matter what you use as long as the parties you're communicating with all agree on the same port number.

ID is the conference ID within the port number. Since on multiuser machines several people may be trying to audioconference on the same port, the ID field is used to keep the different audio streams straight. Since a Mac can only be used by one person at a time, however, this field doesn't really do much so you may as well leave it zero unless you're talking to a *vat* process which is using a different ID.

Net Protocol is a popup menu which lets you select the network protocol Maven will use to send digital audio on the network. This isn't much of a selection, however, as support for RTP (the IETF Real Time Protocol) isn't ready for the public yet, so you're stuck with the VAT protocol, which is what you'd use anyway to interoperate with the Unix *vat* program.

Ignore Unsolicited Streams, if checked, will cause Maven to ignore any incoming ID messages from anything other than the sites to which you have manually created sessions. Leaving it unchecked causes Maven, when it receives an ID message from a site it doesn't know about, to create a new outbound stream to that site and a copy of your audio packets to be sent there.

When you're done setting up Maven's configuration, click the Okay button to dismiss the dialog.

Changes you make in the Configuration dialog take effect immediately. If you have changed only the "Ignore Unsolicited Streams" or "Local Site ID" settings, the current conference in progress is not disturbed, although a change in the Local Site ID results in the change appearing to other participants after a few seconds. Changing the Port or ID causes Maven to remove all current conference participants, because it has to close and re-open the network connection and the participants on the old port and ID would see you "go away" anyway since you would no longer be in the "same place" in the address space.

The Audio menu allows you to control how Maven deals with both analog audio signals from your sound input device and the digital audio signals being received from the network.

Audio	
Mute Speaker	⌘M
Antialias	
Louder	⌘L
Softer	⌘S

Push-To-Talk	⌘P
Encoding	▶
Quantization	▶

Mute Speaker, if checked, will silence playout of incoming audio packets. This is useful, for instance, if your phone rings while you're listening to a conversation.

Antialias, if checked, will cause Maven to oversample the received audio up to 22kHz when it is converted to analog form at the speaker. Without antialiasing, the 8kHz samples are converted verbatim, which can result in an annoying buzzing sound which results from harmonics of frequencies between 0 and 8kHz being synthesized by the D/A conversion process. Obviously, leaving **Antialias** checked will result in the higher sound quality, but this also causes the Sound Manager to do more work (linear interpolation between samples). Thus you pay a CPU penalty when antialiasing, and you may want to turn it off if your Macintosh is behaving sluggishly.

Louder will cause Maven to increase the amplitude of the received audio before sending it to the speaker, causing it to seem louder. You should *not* use this as a substitute for the volume control in the Sound control panel. However, if you have the volume up all the way and you are still having difficulty understanding the other participants (this sometimes happens if their microphone signal is weak), you can use **Louder** to boost the signal. Boosting the signal too far will cause it to clip, resulting in the same distorted sound you hear when turning

up the volume on a radio beyond what the speaker can handle. Each time you select **Louder**, the signal is boosted more—the effects are cumulative.

Softer will cause Maven to decrease the amplitude of the received audio. It is the exact opposite of **Louder**. Again, do not use this as a substitute for the volume control in the Sound control panel. However if the signal is sounding mushy or distorted, it may be clipping and softening it may make it more understandable.

Push-To-Talk controls how Maven decides to send audio packets to the network.

If unchecked, Maven operates in squelched mode. In this mode, no packets are sent if the average sound level being detected at the microphone is less than some threshold. The idea is that when you speak, the sound level exceeds this threshold and Maven begins transmitting. When you stop speaking, the sound level drops below the threshold and Maven stops transmitting. This is natural and convenient, but you'll find that if your environment is loud, it doesn't work very well. Also, keep your microphone far from your Macintosh speaker, or the received audio will trigger the silence detector.

As alluded to above, squelched mode depends on the Sound Manager being able to keep the sound input channel open at all times (so Maven can monitor the sound level at the microphone). If you have a model of Macintosh which doesn't allow sound to be input and output at the same time, you can't operate in squelched mode as the Sound Manager won't be able to create any sound output channels for you to hear other people.

If **Push-To-Talk** is checked, Maven will not send any audio except when the mouse button is pressed in Maven's main window. If you've ever used a two-way radio, this will be somewhat familiar to you as you press the mouse button to talk and release it to receive, just like the PTT switch on the side of a radio microphone.

Push-to-talk mode isn't as cool as squelched mode, but I have to say it works a lot more smoothly in practice. But use whatever you can get to work best for you.

Encoding

μ-law (64kb/s)
Intel DVI (32kb/s)
Linear (64kb/s)
LPC4 (5kb/s)
CELP (5kb/s)
GSM (13kb/s)
Δ-mod (16kb/s)

The **Encoding** submenu of the audio menu lets you select the digital encoding Maven will use when it digitizes sound.

μ -Law is a format similar to Linear but which allows 12 bits per audio sample to be coded as 8. This provides greater dynamic range which the Mac can't appreciate in any event as the sound hardware in current Macs only uses 8 bits per sample anyway. It is, however, what you'll have to use to interoperate with most other audioconferencing systems as it is a standard encoding in the telephone business. μ -Law streams occupy 64 kilobits per second of network bandwidth.

Intel DVI is the simplest of the compressed audio encodings. It provides nearly the same audio quality as Linear while using only 50% of the network bandwidth. You should use DVI coding if you're concerned about network load. DVI coding is understood both by other Mavens and by *vat*. Each DVI stream uses 32 kilobits per second of bandwidth.

Linear is the most natural format for the Macintosh, but unfortunately isn't understood by the Unix *vat* program. Don't use Linear if you're talking to something other than another copy of Maven. Each stream of Linear audio also uses 64 kilobits per second of network bandwidth.

Δ -mod is a somewhat crude delta-modulation encoding. It is quite bandwidth-efficient, using only half the network bandwidth of DVI and one quarter that of μ -law. It does not however go to any great pains to preserve audio quality, so using this encoding will result in somewhat noisy audio, not unlike that of an AM radio tuned slightly off-station. The quality is however plenty good for speech. You should use this when you have very low bandwidth available to you and making efficient use of this bandwidth is more important than the quality of the audio. Since this encoding algorithm was dreamed up by the author from some information in Cooper and McGillem, *Modern Communications and Spread Spectrum*, it is not likely to be

understood by other conferencing tools. It was also not present in versions of Maven older than 2.0a18. Therefore, use this encoding only when you're sure all the other participants in the conference are using Maven 2.0a18 or newer.

Other encodings, such as GSM and LPC4, can provide even greater compression of the digitized audio data, but they're harder to implement and I haven't gotten round to them yet. They appear grayed-out on the menu. Also, these encodings are very CPU-intensive, and will never work at all on 680x0-based Macintoshes.

Maven can understand every encoding it can send, and there's no requirement that others send to you in the same format you're sending to them.

Quantization
20 ms
40 ms
80 ms
160 ms

Quantization controls how much sound information Maven will put into each outgoing network packet. The longer the quantization interval, the fewer packets you'll send out and the nicer you'll be treating routers and other network switching equipment. Longer quantization intervals, however, result in longer delays between when you speak and your voice is heard on the remote end. Use the largest value you find tolerable; to be honest I don't find 80 milliseconds bad at all, but some people find it hard to communicate with large delays.

Addendum: While writing this document, I played the entire Nine Inch Nails Pretty Hate Machine album into Maven, which was sending across the UIUC campus network to a Sparcstation running *vat*, which was playing the album in my office. So it's apparently fairly stable.

Please send comments and bug reports to me. I'd rather they come in email as I really hate talking on the phone. My email address is `kline@uiuc.edu`.

Thanks to George Brett and Jane Smith at CNIDR for offering to take some of the heat of people clamoring for Maven, which is heat I've been ignoring for a year or so. Thanks to Pete "Network Time" Resnick for help with some more obscure parts of the Macintosh Operating System; it's because of him telling me about something called the Deferred Task Manager that Maven continues to receive

and play audio even when you're pulling down a menu or off in another application doing something else. Thanks to the IETF Audio-Video Transport Working Group for the necessary mental stimulation to begin working on this in the first place. And thanks to Yvonne-Marie Andres at the Global Schoolhouse for the nifty T-shirt.

Maven is Copyright © 1994 by Charley Kline and the University of Illinois Board of Trustees.

Redistribution and use are permitted provided that the following conditions are met:

1. Redistributions 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.
2. All advertising materials mentioning features or use of this software must display the following acknowledgement: This product includes software developed by the University of Illinois, Urbana and its contributors.
3. Neither the name of the University nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE TRUSTEES AND CONTRIBUTORS "AS IS" AND ANY EXPRESS 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 TRUSTEES 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.

The Intel DVI codec is Copyright © 1992 by Stichting Mathematisch Centrum, Amsterdam, The Netherlands.

STICHTING MATHEMATISCH CENTRUM DISCLAIMS ALL WARRANTIES WITH REGARD TO THIS SOFTWARE, INCLUDING ALL IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS. IN NO EVENT SHALL STICHTING MATHEMATISCH CENTRUM BE LIABLE FOR ANY SPECIAL, INDIRECT OR CONSEQUENTIAL DAMAGES OR ANY DAMAGES WHATSOEVER RESULTING FROM LOSS OF USE, DATA OR PROFITS, WHETHER IN AN ACTION OF CONTRACT, NEGLIGENCE OR OTHER TORTIOUS ACTION, ARISING OUT OF OR IN CONNECTION WITH THE USE OR PERFORMANCE OF THIS SOFTWARE.