

MAVEN

Due to popular demand, this is the first officially released version of Maven, an Internet audioconferencing tool for the Macintosh computer. This is alpha-quality code, and this is version 2.0a11.

This version incorporates several enhancements over the “toy” that made it out to the Global Schoolhouse a few weeks ago. 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. It was very incomplete and could easily crash. It also had some network behaviour that was rather unfriendly.

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 don't have it already, you should install version 3.0 of the Macintosh Sound Manager, which 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.

Maven requires a Macintosh with sound input hardware and will refuse to run if it can't find any. 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.

When Maven is first launched, it presents you with a configuration dialog:

Configure Maven

Local Site ID:

Port: ID:

Net Protocol:

Audio Encoding:

Audio Quantization:

Ignore unsolicited streams
 Push-to-talk

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, but doesn't yet do anything with identification received from other Mavens on the network. The Unix audio tool, *vat*, however, does. 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.

Audio Encoding is a popup menu which lets you select the digital encoding Maven will use when it digitizes sound.

The default, *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 uses 64 kilobits per second of network bandwidth.

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 also 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.

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 popup.

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.

Audio 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 100 milliseconds bad at all, but

some people find it hard to communicate with that large a delay.

Ignore Unsolicited Streams, if checked, will cause Maven to ignore any incoming audio packets from anything other than the sites you have created sessions to. You probably want to leave this checked. Unchecking it causes Maven, when it receives audio packets 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. Maven doesn't tell you when it creates new streams this way, so it's easy to be generating large amounts of network traffic (a copy of each audio packet for each site Maven is sending to) without being aware of it.

Push-To-Talk is a checkbox which controls how Maven decides to send audio packets to the network.

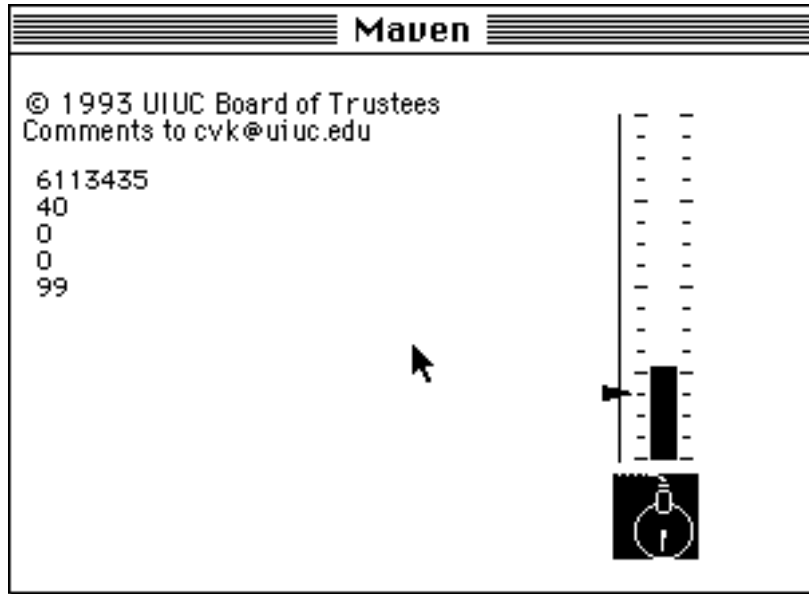
If unchecked, Maven operates in Silence-Detection mode. In this mode, no packets are sent if the average sound level being detected at the microphone is less than some threshold. The idea is that when you speak, the sound level exceeds this threshold and Maven begins transmitting. When you stop speaking, 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 I alluded to above, Silence-Detection 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 Silence-Detection 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 Silence-Detection 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.

When you're done setting up Maven's configuration, click the Okay button and Maven will go on to its main window.



The main window needs a lot of work; there are lots of things I want to put in it that I haven't gotten round to yet, which is why it looks so empty. The only interesting things in the main window are on the right hand side, where 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 Silence-Detection mode, a small triangular slider will appear next to 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'll need to 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.

Warning! Maven will let you continue to add new sessions this way 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. Maven does not tell you how many other sites you are sending to. There's no way to remove a session once you've created it, short of quitting Maven. Take special care when sending to more than one other site that you don't swamp the network with packets!

The big white area to the left is where the list of participants in the audioconference will be displayed, as soon as I can figure out the List Manager. Right now I just have some debugging information there, which I've left in just so you know Maven is still alive. If you care, the debugging numbers, from top to bottom, are the current value of the master clock, the current playout delay in samples, the number of times a quiet element was inserted into the output channel, the total number of bytes queued in the playout buffer, and the number of free memory buffers.

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 `cvk@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 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 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. "Thanks" to Greg German at Ohiolink for letting the cat out of the bag and releasing the development version of Maven to the Global Schoolhouse. This was actually good as it forced me to finish up this alpha release, something I've been shanking for months.

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.