


# Tondo Instructions

## version: 1.0.2



## Introduction:

Welcome to **Tondo**!

**Tondo** lets you create and manipulate audio in a new way: by painting sound on a radial sonogram. With **Tondo**, you can learn to create sound out of its simplest elements, and how rhythm, timbre, melody, and harmony all are closely related.

The basic principle behind **Tondo** is that all sounds can be built by adding together very simple sounds (sinusoid waves) with varying frequencies and amplitudes. **Tondo** organizes this information in a way that is easy to manipulate, by painting with specialized tools. As you paint the sound, the corresponding audio is generated. At any point, you can save the information or the audio you have generated and export it in a number of ways.

**Tondo** can be used for making samples for other programs to use, rhythms and experimental explorations of sound. Tondo 1.0.2 can export audio onto the Clipboard and also uses AudioCopy by Sonoma Wire Works. [ <http://www.sonomawireworks.com/> ]

"AudioCopy and AudioPaste are trademarks and software of Sonoma Wire Works"

You can export the information which defines the sound onto the clipboard, which you can then paste into documents, mail to friends, post on blogs, or send to me to make public at [Tondo@jhh.net](mailto:Tondo@jhh.net) . You can see some demos, sample voices, and updated information at <http://www.jhh.net/iPhone/Tondo> .

## History:

I have a pretty extensive background of working with sonic spectral synthesis. Back in the 1980s, I wrote a fairly popular freeware

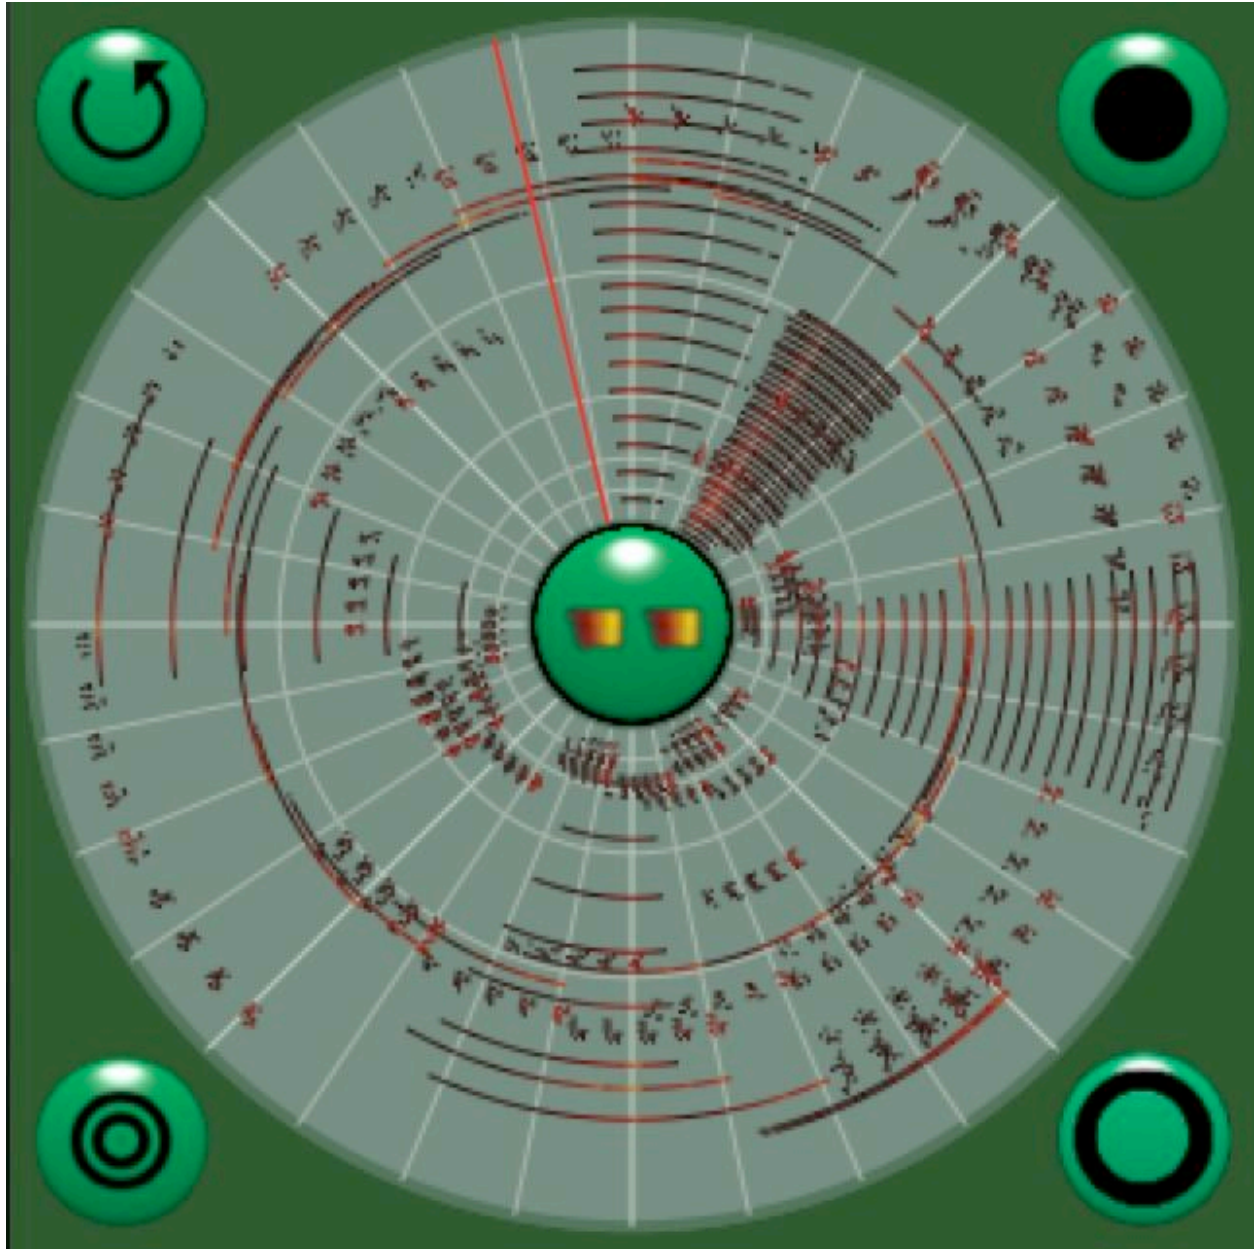
program called **RGS** for Commodore Amiga computers. Some of the ideas in **RGS** are present in **Tondo**: the color scheme of the spectrum, the color bar that can remap colors to amplitudes, and many of the tools. Since I've had about 20 years to think about it, I've added some ideas as well, such as the whole idea of displaying the spectrum radially, the better to understand it as a loop.

Painting and observing spectrograms is a great way to see what's happening in sound, even though it's nothing like physical processes of sound production, unless you have a rather big and perfectly tuned panpipe orchestra.

**Tondo** is unconcerned with being able to build musical compositions as they are usually thought of in Western music, although with great care you can do it. It's much more a tool for exploring sound itself.

## **Terminology:**


In this document, I refer to the Apple iOs devices – iPhone, iPod Touch, iPhone 4, or iPad – as the “device,” so that I don't have to repeat the tedious list every time.



Sonogram Canvas


The main part of the screen is a radially presented sonic **sonogram canvas**. It acts as a canvas for displaying and painting audio information. The frequency increases from the center of the circle, while time increases when moving clockwise. Color is used to represent the amplitude of the underlying sinusoid wave at that point in time and frequency.

The first slider box, the **amplitude color bar**, under the **sonogram** is used to select the amplitude that you are painting with.



Amplitude Color Bar


The mapping between colors and amplitude can be changed by using two fingers to set the limits of the colors on the color bar, which will redraw the screen. This is so you can see details in the sonogram.



Amplitude Color Bar, remapping the colors

A double touch will restore it to full amplitude mapping.

The second slider box, the **size setting box**, is used to control the size of the **painting tool** that you are currently using.



Size Setting Box

It looks a little like a guitar fretboard, where the “frets” mark integral divisions of the sonogram, like  $1/2$ ,  $1/3$ ,  $1/4$  etc. It is not equally spaced because you generally want more accuracy with shorter sounds than longer ones. Touching the **size setting box** sets the length in time, and in some cases the width in frequency for the tool you are using. This is displayed as the width of a colored band between two dimmer bands.



Size Setting Box showing “height” bands.

For instance, the **Filter** and **Echo** tools affect frequency as well as time. Multiple taps on the box will snap the length successively to the frets: a double tap to 1/2, a triple to 1/3, four taps to 1/4, etc.

## Corner Buttons:

There are four buttons surrounding the sonogram canvas.



Undo once



Clear

The **Undo Button**. Touching this button once undoes the work of the previous gesture. Each subsequent single touch will undo another level. A double touch will clear the whole screen. To help distinguish between a double touch and a series of single touches, the button changes its appearance to when it’s primed to be a “clear” button. I used to be a bit hasty when undoing a few levels and wipe things out without this reminder. I didn't want to put up a confirm dialog which is annoying.

**Undo** is immensely useful when you are trying to build a precise sound!



Record



Ready To Record



Now Recording

The **Record Button**. If your device supports recording, you can record a sound and have it converted into a sonogram for you to save or alter. When touched, it turns yellow, and when released, the button turns red to indicate that it is recording. Touch it again to stop recording. **Tondo** will then analyze the recording, and while processing, it turns purple. It will play the recorded sound while analyzing, and a purple iris will expand out as the analysis proceeds. This processing takes a bit of time, depending on the speed of the device. When it's done, it resynthesizes the new sound, extending a red iris as it progresses.

The recording length is set by the **Seconds slider** on the Configuration screen. It normalizes the level of the sound after recording, so that it can pick up quieter sounds. Recording wipes out existing sounds and “undoes”.

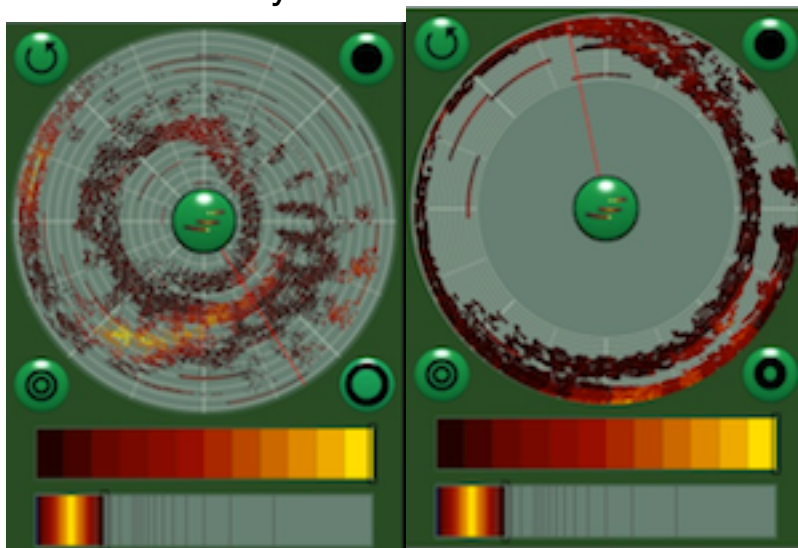


Zoom in



Zoom out

The **Zoom Button**. If you want more accuracy in the low frequency part of the spectrum, this button will zoom in so you can draw them more accurately.





Hide Grid



Show Grid

The **Grid** button. This shows or hides the white grid under the sonogram. Sometimes it's a little easier to see what you are drawing without it. The grid adapts itself to the **Frequency Lock** and **Time Lock** settings.



# PAINTING TOOLS

Touching the center button pops up a pie chooser, which lets you select the spectrum painting tool for painting and editing sonograms. It's hard to come up with icons for these things, but I'm trying to avoid words as much as possible!



Tool picking pie chooser

To pick a new drawing or editing tool, just touch the appropriate pie slice. The central button icon will change to confirm your selection. When you've made your choice, the pie chooser will go away.

Beyond picking the next tool to use, the position of the touch may also determine other characteristics of the tool. The **ramp**, **line**, **noise** and **harmonic** tools can have their envelopes skewed by sliding clockwise and counter clockwise within the pie slice. You can see this affect the colors in the **size setting box** at the bottom of the screen. A **ramp** with a fast attack and slow decay sounds like a bell, with a slow attack and fast decay like a backward bell; it's very adjustable.

The **echo**, **filter**, **eraser**, **increment**, and **decrement** tools have a special area marked at the edge of their sections which will apply that tool (or a reasonable variant) to the entire sonogram. I refer to this as the **whole sound mode**.

Each painting tool has its own way of interpreting the **amplitude color bar**, **size setting box**, **time lock**, and **frequency lock** settings when they are painting the sonogram. See the description of each tool for details!



Size Setting Box with grayed out settings

For tools which alter the existing sonogram (**filter**, **echo**, **increment**, **decrement**), the **size setting box** turns grayscale, indicating that it's not drawing that amplitude, but changing an amplitude.

Some painting tools act immediately and turn into sound immediately. Others wait until the gesture is done and then resynthesize the whole canvas into sound. This is so slow operations don't slow up the drawing. In general, if you want to draw something complicated, you should make the audio sample as short as you are comfortable with, draw it, and then lengthen it using the **Seconds slider** in the **Configuration screen**. When resynthesizing, the screen shows an expanding red iris as it builds the sound.

## DRAWING TOOLS DETAILS

### RAMP:



The **ramp** adds a triangularly enveloped line of amplitudes to the sonogram. The peak of that triangle amplitude is selected by the **amplitude color bar**. *Where you click in the sonogram canvas designates where that peak goes, not where the sound “starts.”* I find this much more accurate than “starting” with a nearly inaudible sound.

The position of the pointer within the pie slice changes where the peak of the triangle appears in the envelope: counter clockwise puts it toward the front, clockwise toward the back. This corresponds to a sharp attack with a long decay vs. a long attack with a sharp decay and all variants in between. Because amplitudes are always summed into the canvas, if you set the color of the **ramp** fairly low, you can build up some subtly changing amplitudes.

**Ramp** comes up as the default painting tool.

### LINE:



The **line** tool draw a line - actually a series of short, closely spaced ramps - on the canvas. You can affect the size of these ramps with the size setting box. Like the **ramp** tool, it adds its amplitudes to the canvas, so setting the color to something low will give you a more satisfying experience. **Line** is useful for sketching out sounds like bird calls.

**Line** observes the settings of the **Frequency Lock** and **Time Lock** settings, and therefore is an easy way to make up harmonic material

## HARMONICS:



The **harmonics** tool draws several ramps spaced equally apart. If you set up the canvas to use the **Linear Frequency Type**, these will roughly correspond to harmonics of a sound. This can then be altered with other tools for more interesting effects. If you are using the **Logarithmic Frequency Type**, you get a kind of abstract, spacey sound. The place you touch originally is set to draw at the eighth harmonic, in order to give more accuracy. This is referred to as the **harmonic handle**, and can be changed in the **configuration screen**. You may notice the the peak amplitude fades out the further the harmonic is from the harmonic handle. That's to make the resulting waveform more interesting.

## BLOB:



The **blob** adds in a uniform, wide area of amplitude to the canvas. This generally results in some heavy “beating” noises, but it’s good raw material for further work. A really good way to work with it is to make the amplitude low and build up the sound bit by bit.

## ERASER:



The **eraser** erases an area from the screen. This can help clean up an analyzed sound or create a kind of negative rhythm.

The eraser's **whole sound mode**, accessed at the outer edge of the pie slice, removes all amplitudes less than the current setting in the **color bar**. This is really useful for cleaning up noisy spectrograms!

### NOISE:



The **noise** tool randomly deposits single-pixel dots at the current amplitude or less in the area specified by the **Size box**.

### INCREMENT:



This tool would be a lot like the BLOB tool, except that only adds amplitude to painted areas of the spectrogram. That is, it does not paint on empty space.

Its **whole sound mode** takes all the non-empty space in the canvas and sets it to the current amplitude. This is useful for the following trick: analyze a recording, use the **increment whole sound** mode to set it all to a very quiet setting, and then use the **increment tool** to paint in the sound as it plays.

### DECREMENT:



This tool, like the increment tool, works only on non-zero areas of the canvas, quieting down whatever it paints over.

Its **whole sound mode** rescales the sound to the current amplitude color.

### ECHO:



The **echo** tool paints a scaled echo of the area below the touch at a distance set by the **size box**. You can get a lot of great effects this way! Since you can selectively apply it, it opens up a lot of possibilities.

Its **whole sound mode** applies the echo to the whole sound.

### **FILTER:**



The filter tool smears the amplitude in its size area a little to the left and right. Sharp attacks are thereby smoothed somewhat. I mostly use it to turn noise into a more blurred sound.


Its **whole sound mode** applies the filter to the whole sound.

# TOOL BAR




The **tool bar buttons** let you do the following:

## PLAY / PAUSE

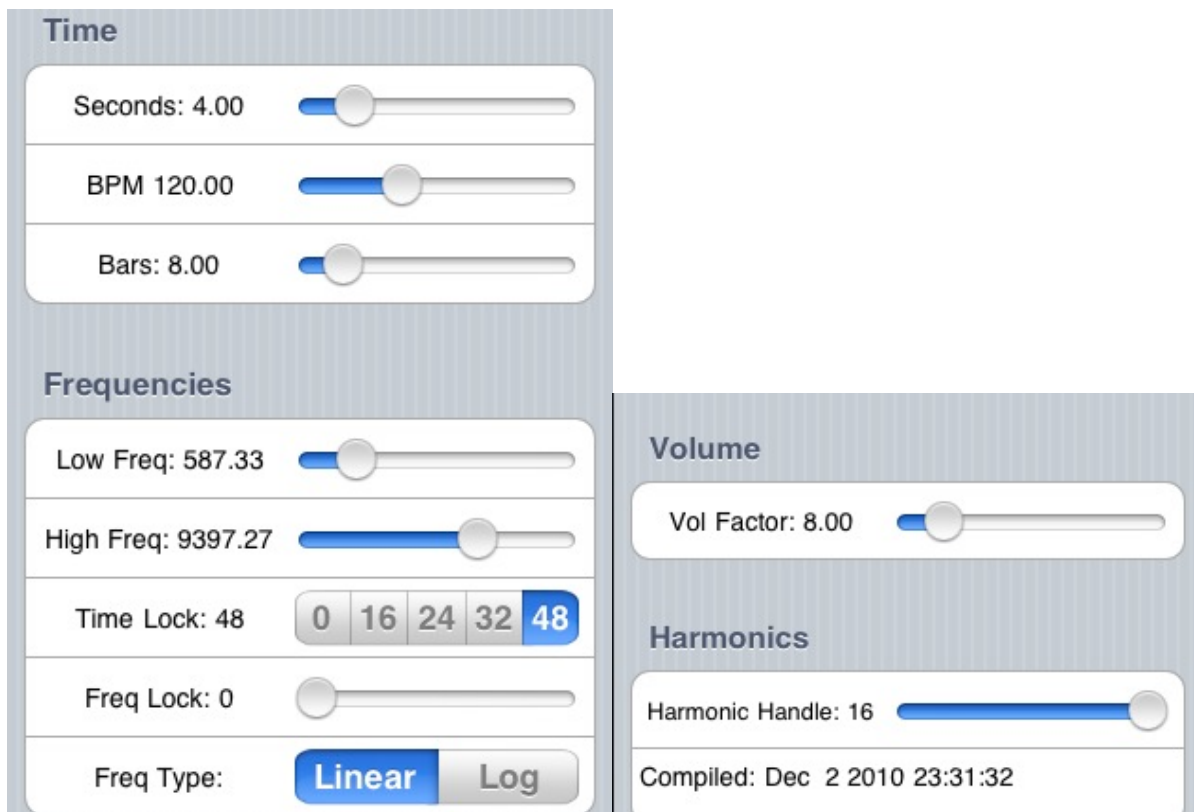


You can start and stop the audio with this button. Sometimes the audio stops by itself (sorry!) and tapping this twice will get it going again.

## CONFIGURE



The configure screen is a little longer than the phone screen, so you may need to scroll it!



Each control has a slider (or segmented chooser) to help set the value. Some sliders (**Seconds, Low Freq, High Freq**) are non-linear to favor low values. The **Seconds** slider constrains its values to half-seconds.

Tapping on the associated label lets you set the value explicitly! If the value you type is out of range, a message will tell you the range limit. Also, if you touch any slider for about a second and a half without moving it, it turns blue and gives you more accuracy. This is important on a small device like an iPhone!



Changing the settings forces a resynthesis of the sound when you are done, which may take a while.

Notes on each setting control:

### **Time section:**

These three sliders are linked: changing one changes the others (unless the values fall out of their range).

**Seconds:** the number of seconds in the loop, and also for any recordings. It can be set from **0.30** to **60** seconds. If you want to make an LP lock groove, you should set this to **1.8** seconds.

**BPM:** Beats per minute.

**Bars:** Number of bars in the whole sample.

### **Frequencies section:**

The frequency sliders don't let you put the low frequency above the high one, nor this high below the low. You can type in a note name, like **F3** as a frequency. The **High Freq** has an additional feature that you can set it to any number of higher octaves over the current low frequency by typing **@<octave>**, for instance, **@2**. You can set it to any harmonic of the low frequency by typing **:<harmonic>**, for instance **:6**. The lowest frequency is **0.25 Hz**, the highest is **21000 Hz**.

There are two locks in there as well: the **Time Lock** constrains the time to that many divisions of the whole sound. The segmented control makes it easy to do this, but the label lets you set it from **0** to **128**.

**Frequency Lock** actually constrains the distance from the inner ring to integral multiples of that lock. What that frequency works out to depends on the **High** and **Low Frequency** settings and the **Frequency Type** setting. The exact value, in Hertz for the **Linear** type, and in cents for the **Logarithmic** type, is reported in the label. These are for reference only; you can't go in and specify the **Frequency Lock** as these kinds of values.

**Frequency Type** sets whether the sonogram orders its frequencies **linearly (Linear)** - good for building harmonic sounds, or **logarithmically (Log)**, which is more like a conventional piano keyboard.

### **Volume Section:**

There is no way to know in advance how loud your spectrogram will turn out to be, so this setting scales all the amplitudes by this amount as it sums them up into a sound. Setting **Volume** to zero will tell **Tondo** to reset the volume to the optimal value for the current sonogram.

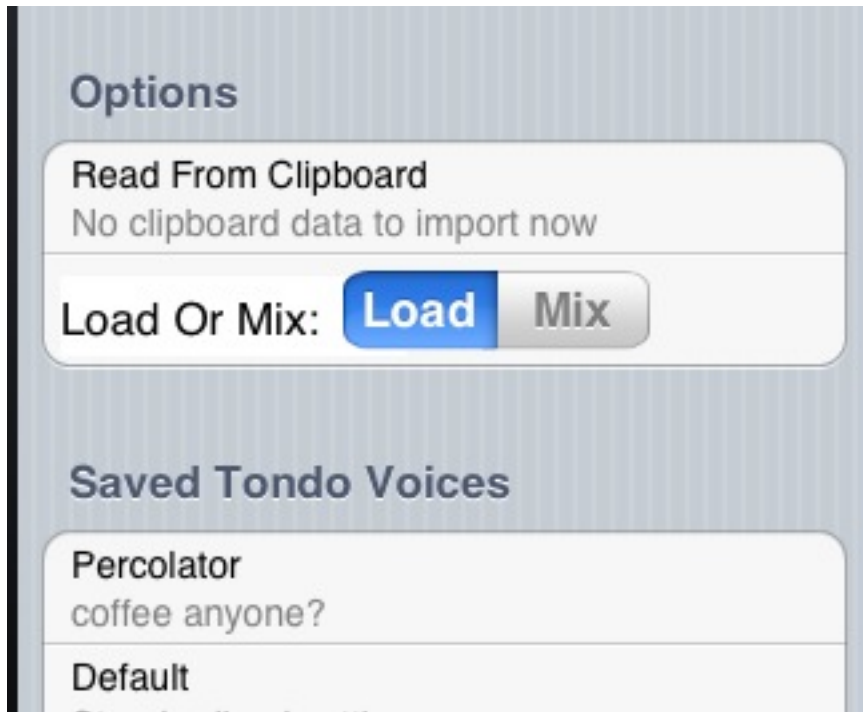
### **Harmonics section:**

The **Harmonics** tool uses this setting to determine which harmonic is being tracked when you are using it to draw sound. It ranges from **1** to **16**.

### **Compilation Date:**

The compilation date is mostly so I can keep track of which development version I'm running!

## LOAD



### Options section:

#### Read from Clipboard:

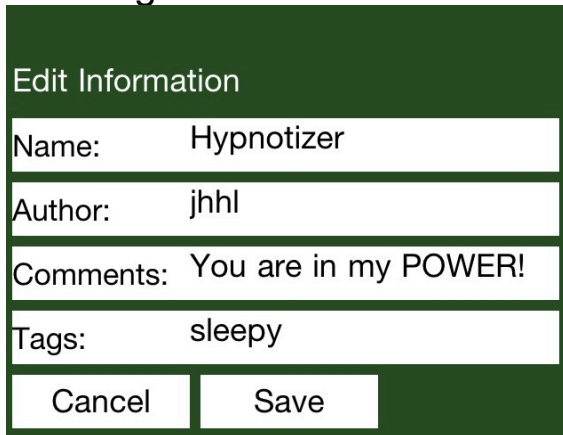
If a validly exported **Tondo** Voice is in the iPhone clipboard, you can import it with this button. Otherwise you see the error message: **No clipboard data to import now**. The export format is in UTF-8 and is very straightforward so you can alter a lot of your data easily. Details are found in a further appendix.

#### Load/Mix

This button lets you load saved **Tondo** voices on top of each other. The settings are reset to the latest load, and the canvases are added onto whatever is already there. This can make for a lot of useful presets before starting to draw, or after (if mixing in a blank canvas).

## Saved Tondo Voices section:

Here are all the saved voices and some presets as well. The most recent is on top. You can edit the list by touching the **Edit** button, which will allow you to delete a voice, or you can edit its information by touching its name.




The image shows a screenshot of a mobile application's 'Edit Information' dialog box. The dialog has a dark green header with the title 'Edit Information'. Below the header are four text input fields, each with a label on the left and a value on the right. The fields are: 'Name:' with the value 'Hypnotizer', 'Author:' with the value 'jhh', 'Comments:' with the value 'You are in my POWER!', and 'Tags:' with the value 'sleepy'. At the bottom of the dialog are two buttons: 'Cancel' on the left and 'Save' on the right.

Edit Information	
Name:	Hypnotizer
Author:	jhh
Comments:	You are in my POWER!
Tags:	sleepy
Cancel	Save

These fields are pretty self-explanatory. *Tags* are not yet being used, but will be, so put some tags in if you like, separated by spaces. You can always change them later.

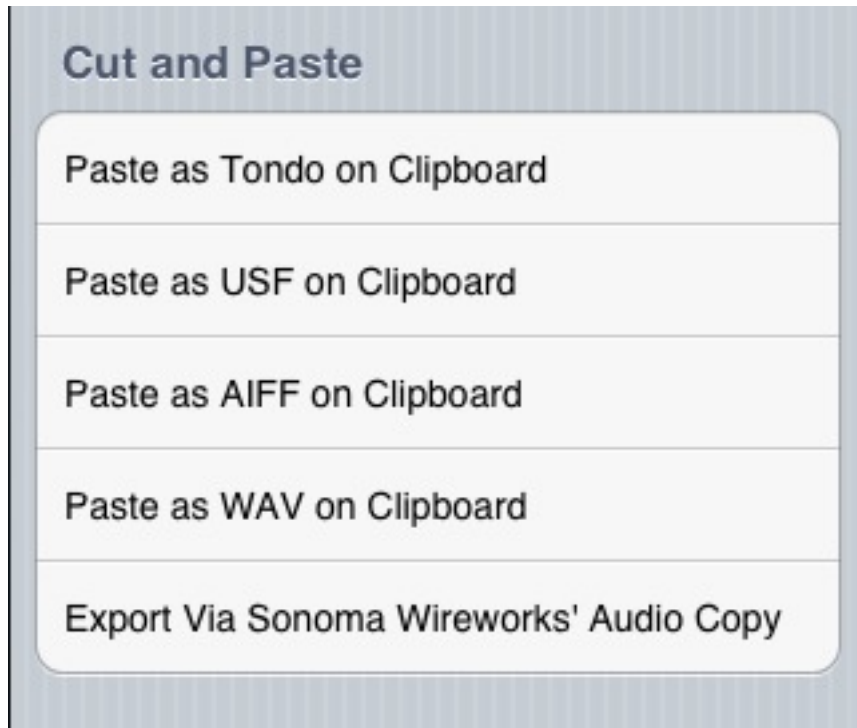
## SAVE



Here you are provided with a number of ways to save the current **Tondo** voice. You can save the parameters that create the file as a **Tondo File** or **Universal Sonogram File**, or save the sound that it generates.




Besides saving things as files, you can also put the same information on the Clipboard. That way, it's easier for other programs to pick up that information:



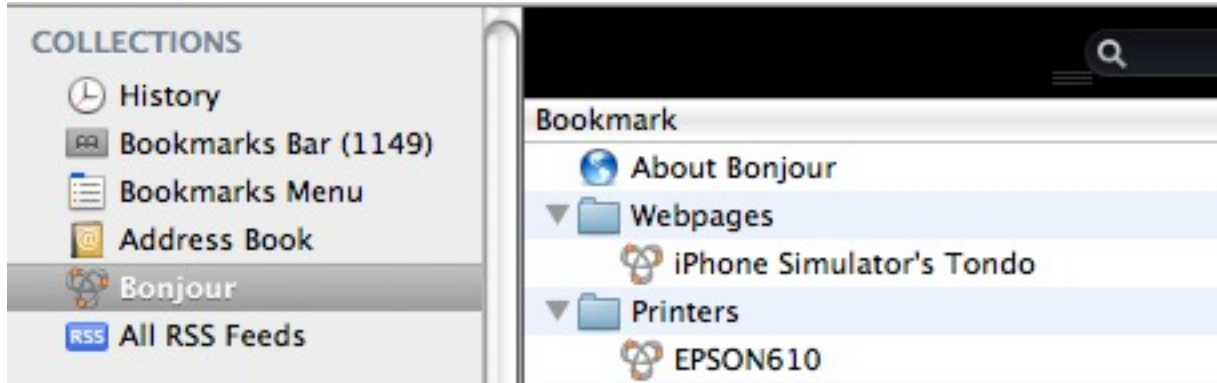
When saving, you are prompted for the same fields as you see in the “**Edit**” section of the **Load Tondo Voices** screen.

After you have saved the voice as a file, you can export that file in a number of ways. **Tondo** can run as an *HTML server*, which shows a special web page that lets you access and delete these files. **Tondo** announces the web page via Bonjour, so you can find it on your local WiFi connection. You set this up in the lower section of the **Save** page.



Your device must be connected to the internet via WiFi, so **Tondo** can pick a port for its HTTP server within that local network. **Tondo** will then automatically show up in a locally connected Web browsers list of Bonjour servers.

If there's a hole in the network's firewall, you can probably access the phone through that hole. To make that more predictable, you can specify the server port that **Tondo** will try to claim.



Or you can just type in one of the URLs that appears in that Server connection status area.

Tondo's web page looks like this:

A screenshot of a web page with a light green background. At the top left is a circular logo with a stylized 'T' and a green dot. To the right of the logo is the main heading: 

# Tondo Audio Files on iPhone Simulator's Tondo

. Below the heading is a paragraph: "The following audio files are in the **Tondo Audio Files** folder." Underneath is a table with columns for Name, Type, Size, Date, and an empty column. The table contains one row: 'chumpy cheep', 'tondo', '10.3', '2010-12-07 15:24:54 +0000', and a 'delete' button. Below the table are several links: [Tondo instructions](#), [Tondo Instructons as a PDF](#), [The official Tondo website](#), [Henry Lowengard's iPhone App website](#), and [Henry Lowengard's iPhone App blog](#).

Not only can you access these files, but you can also get a nice PDF copy of these instructions to print out!

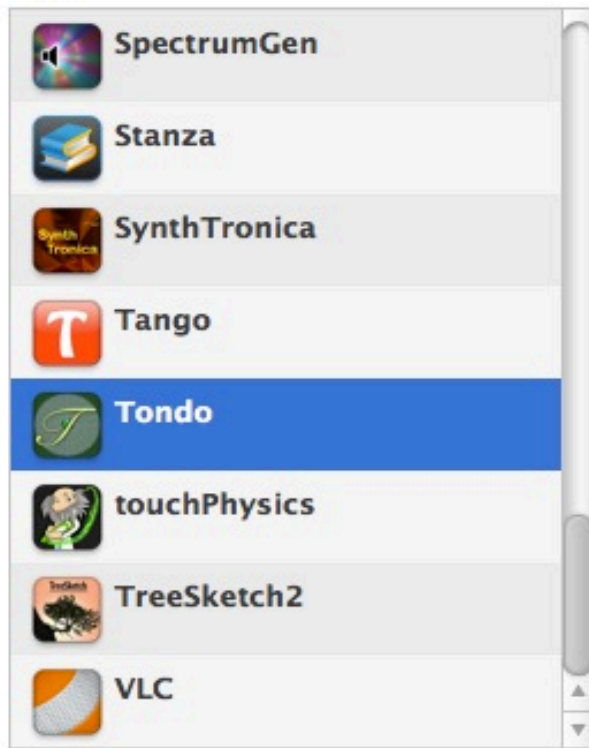


Tondo also exposes these files in the iTunes Document interface. You can find them by clicking the Apps tab when you select your device in iTunes. Look for Tondo in the list of apps:

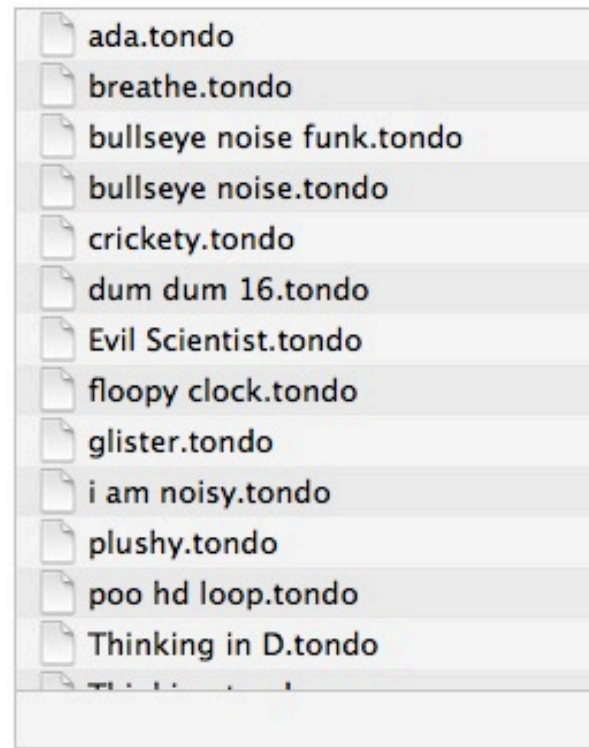
## File Sharing

The apps listed below can transfer documents between your iPad and this

### Apps



### Tondo Documents



It's a good way to export a bunch of sound files.

## **INFO**

You're reading it!

## Appendix I: Tondo voice export data format:

Here is the version 1.0 **Tondo** voice file export format. If the file format changes, this format will still be valid. Most of the fields presented can be in any order or omitted (not **version**). The values for the fields are validated when read. All the names are in UTF-8, so many alphabets and signs can be used in the comments. All the non-header lines are in the form **keyword: data** ;, and because of the semicolon, you can put several on a single line.

```
Tondo FV1.0 tondosities;
version: 101;
name: lickity;
author: jhhl;
comment: mmmm good;
tags: pointilistic;
date: 2010-11-17 10:10:55 PM;
length: 2.000000;
lowestFrequency: 127.999985;
highestFrequency: 8229.133789;
timeLock: 0.000000;
freqLock: 0.000000;
freqScaleType: Log;
volume: 1.000000;
volFactor: 8.000000;
harmonicHandle: 1.000000;
canvas:
```

```
MXic7d1RbsMgDADQnTH3P8w0TZo0Nbs0hGDj9766j3YGbAOptH19DTnG3g6s8Vnp
KngAAAAANLykBsKeVLwegHsTIVT3VknQAsAAAAAAOjmsTpMpMAAAAAAAKjFk3GI
bWqNagAAAABsxlUXuJm2AwAAwCth7o5hAgEAAM45tAMAAAAAAGzOF0IAAAAAAPCP
P8IPbVIYAEqx9QMAAABDPFyAhojFETEmWCJLMWSJE3iHyuaXTICdHM0fOBN/iuJH
CAAAAFdxCwYA4BnnRQAoLuNhIGPMADDPhztj/9tmb722dshCtdLtlmSRmRlyJUI6
Tl6NfQ6EJEWBUfoIAGWU2vTqDLbOSJlIGkFuajikjmWxclxBHsetlBoAAD+cC6GL
UgE0gl0sXUhZlJr1w8v8AwAMcZwC4JHdAQAA1kp3Jk8X8J+8kbMrOfnAlLxmjkqz
/EAsuhLvkC/ANXq6SZqOc02gaYYLAEGE2zvDBURF0hAAA0gW5QIRJY4hWwwCqEoL
azEzEMHqSlz9+4E29RmEhWB3cryHWQIA4DZxDp9jkcQZx6eOHQYxm3+gBQAAbGGv
+8Veo4nPfCdm8QCABscEuINKAx5oDDCVEgNa9Afgct/AeyQz;
```

- **The top line is a header.** Tondo looks for this in streams presented to it as a prerequisite for processing it. “tondosities” is a “Search engine hapax legomenon” - a string of letters so rare that searching

for them will turn up **Tondo** voices saved all over the Internet in many places.

- **version**: is another place for version details to be put.
  - **name**: The name of the voice.
  - **author**: The author
  - **comment**: any notes ore extensive comments.
  - **tags**: Some space separated tags. These will be used in the future!
  - **date**: date in YYYY-MM-DD HH:MM:SS PM format
  - **length**: in seconds
  - **lowestFrequency**: in Hz
  - **highestFrequency**: in Hz
  - **timeLock**: in integral units meaning how many equal parts of the loop to clamp to. 0 means no lock;
  - **freqLock**: in integral units, meaning how many frequency lines to clamp to. ;
  - **freqScaleType**: Linear or Log;
  - **volume**: (actually unused);
  - **volFactor**: the amount to divide the whole sound by to prevent clipping;
  - **harmonicHandle**: the setting for the **Harmonic Tool**'s "handle";
- canvas**: here is the actual sonogram data: it's base 64 encoded and when decoded, there is a 1 character format tag ('1') followed by the zipped (with zlib) matrix of unsigned bytes of amplitude information. But watch it: it's a trapezoid nestled in a rectangle that is 128 x 972. The lowest frequency has only 168 characters of information, and it grows linearly from there to 966 or so..

0 : 168.000000	1 : 174.281250	2 : 180.562500	3 : 186.843750
4 : 193.125000	5 : 199.406250	6 : 205.687500	7 : 211.968750
8 : 218.250000	9 : 224.531250	10 : 230.812500	11 : 237.093750
12 : 243.375000	13 : 249.656250	14 : 255.937500	15 : 262.218750
16 : 268.500000	17 : 274.781250	18 : 281.062500	19 : 287.343750
20 : 293.625000	21 : 299.906250	22 : 306.187500	23 : 312.468750
24 : 318.750000	25 : 325.031250	26 : 331.312500	27 : 337.593750
28 : 343.875000	29 : 350.156250	30 : 356.437500	31 : 362.718750
32 : 369.000000	33 : 375.281250	34 : 381.562500	35 : 387.843750
36 : 394.125000	37 : 400.406250	38 : 406.687500	39 : 412.968750
40 : 419.250000	41 : 425.531250	42 : 431.812500	43 : 438.093750
44 : 444.375000	45 : 450.656250	46 : 456.937500	47 : 463.218750

48 : 469.500000	49 : 475.781250	50 : 482.062500	51 : 488.343750
52 : 494.625000	53 : 500.906250	54 : 507.187500	55 : 513.468750
56 : 519.750000	57 : 526.031250	58 : 532.312500	59 : 538.593750
60 : 544.875000	61 : 551.156250	62 : 557.437500	63 : 563.718750
64 : 570.000000	65 : 576.281250	66 : 582.562500	67 : 588.843750
68 : 595.125000	69 : 601.406250	70 : 607.687500	71 : 613.968750
72 : 620.250000	73 : 626.531250	74 : 632.812500	75 : 639.093750
76 : 645.375000	77 : 651.656250	78 : 657.937500	79 : 664.218750
80 : 670.500000	81 : 676.781250	82 : 683.062500	83 : 689.343750
84 : 695.625000	85 : 701.906250	86 : 708.187500	87 : 714.468750
88 : 720.750000	89 : 727.031250	90 : 733.312500	91 : 739.593750
92 : 745.875000	93 : 752.156250	94 : 758.437500	95 : 764.718750
96 : 771.000000	97 : 777.281250	98 : 783.562500	99 : 789.843750
100 : 796.125000	101 : 802.406250	102 : 808.687500	103 : 814.968750
104 : 821.250000	105 : 827.531250	106 : 833.812500	107 : 840.093750
108 : 846.375000	109 : 852.656250	110 : 858.937500	111 : 865.218750
112 : 871.500000	113 : 877.781250	114 : 884.062500	115 : 890.343750
116 : 896.625000	117 : 902.906250	118 : 909.187500	119 : 915.468750
120 : 921.750000	121 : 928.031250	122 : 934.312500	123 : 940.593750
124 : 946.875000	125 : 953.156250	126 : 959.437500	127 : 965.718750

## Appendix II: Universal Sonogram voice export data format:

I made this up as a more generalized format for importing and exporting audio sonograms. Contact me at [USF@jhhl.net](mailto:USF@jhhl.net) for more details, which are pretty flexible at this point.

```
UniversalSonogramFormat 1.0;
version: 1000;
source: Tondo 1.0.2;
name: empty;
author: jhhl;
comment: nothing in here;
tags: locked 16;
date: 2011-11-04 11:38:02 PM;
minTime: 0.000;
maxTime: 12.000;
lowestFrequency: 220.000;
highestFrequency: 2604.519;
frequencyScaleType: Log;
canvasPixelType: UI8;
channels: 1;
timePixels: 972;
freqPixels: 128;
canvas:
```

```
eJztwTEBAAAawqD1T20Hb6AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA  
AAAA  
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA  
AAAAAA  
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAADgN  
+YPAAE=  
;
```

the main difference is that the sonogram data is rectangular, and it describes itself a little better. I may support an unzipped version of the sonogram data, which will be bigger, but a lot easier to generate on the fly!