



Includes

Rythm Bed, Serenade,
Vox Box and Re-Midi

VoxBox

A collection of four programs
Vox Sample, Vox Synth, Vox Beat and Perform

ProAction

3QD
DEVELOPMENTS LTD

APDL

VoxBox

© 3QD Developments Ltd 2026
All rights reserved.

No part of this product may be reproduced in whole or part by any means without written permission of the publisher. Unauthorised hiring, renting, lending, public performance or broadcasting of this product or its parts is prohibited.

While every care is taken, the publisher cannot be held responsible for any errors in this product, or for loss of data or consequential effects from the use of this package.

Software and Documentation supplied by
APDL
a part of
3QD Developments Ltd
Single User version downloaded from the free of charge software section of
www.riscos.com

Please note that any free of charge downloads from RISCOS.com are supplied without any technical support package as per the licence agreement.

This document last edited on 29/09/26



If you have found this application and this manual useful please consider making a donation to help cover the costs of cataloguing, documenting, updating and hosting the APDL applications. Scanning the QR code will take you to our donations page.

Contents

Introduction	5
Perform	6
Starting and quitting Perform	6
Setting up a programme	6
Saving a programme	
Playing a programme	7
The Panel window	8
Playing Samples	8
The Vox Box	9
Overview	9
VoxBeat	9
VoxSample	9
VoxSynth	10
Playing sounds	10
Non Musical uses	10
Common features	11
Getting Started	11
The Iconbar menu	11
Using the keyboard	11
Saving your voice	11
Vox Beat	12
The main window	12
No.	12
Name	12
Play/Compile	12
Waveform window	12
Red scroll bar	12
Type	12
Octave	12
Vol	12
Pitch	13
Reset	13
Loading and editing a sample	13
Sample types	13
Changing the pitch of a sample	13
Changing the volume of a sample	14
Allocating the right octave	14
Changing the names	14
Displaying information about samples	14
Saving your module	14
Removing a sample	14
Hints on using VoxBeat	14
Vox Sample	15
The main window	15
Sample type	15
Play/Compile	15
Waveform window	15
Red scroll bar	15
No of loop cycles	16
Pitch	16
Reset	16
Loading a sample	16
Marking the sample	16
Changing number of loop cycles and pitch	16
Changing the volume	17
Editing the amplitude and pitch envelopes	17

Vox Synth	18
The main window	18
No.	18
Time	18
Name	18
Play	18
Waveform window	18
Harmonics window	19
FM synthesis parameters	19
Creating new waveforms	19
Using harmonic synthesis	19
Drawing freehand	20
Using the FM synthesis parameters	21
Using Real-Time Linear Interpolation	22
Loading samples directly	22
Copying all or part of a voice	22
Editing the amplitude and pitch envelopes	23
Amplitude and Pitch Envelopes	23
Editing the envelope windows	23

Introduction

This package is a combination of four separate applications. Together they provide a very useful addition to the Rhapsody program, also available from APDL - ProAction.

The four programs are:

- **Perform** - an application enabling you to play Rhapsody and MIDI files as well as sound samples without the need to have Rhapsody present.
- **VoxBeat** - an application to turn sampled percussion sounds into a useful module that enables a user without access to a MIDI instrument to obtain realistic percussion accompaniment to his or her score.
- **VoxSample** - an application to convert sound samples into usable Archimedes voices for use with Rhapsody, Maestro or any other application using voices.
- **VoxSynth** - an application to synthesise Archimedes voices digitally.

This manual is split into sections dealing first with Perform, which is unrelated to the other three programs and then with the Vox series.

Perform

When you have built up a library of scores you may wish to play a number of them in sequence eg as movements of a sonata or tracks in an album. This can be done using the application Perform which will play up to 12 Rhapsody or MIDI files in sequence. Perform will also play samples such as those produced by Armadeus. If you have a sampling board, you could announce each track before it is played using this feature!

In addition, Perform can load voice modules. This solves the problem of needing to have the right voices loaded before Rhapsody will play.

This is achieved by setting up a 'Programme' (ie a concert programme, not a computer program) which can be saved as a special file and reloaded at any time. Items on the programme can easily be added, deleted and moved around. The programme may then be played either once or continuously, and because Perform is a standard multi-tasking application, other RISC OS applications may be run at the same time.

Starting and quitting Perform


Perform is a standard RISC OS application and is loaded by double clicking on its icon in the usual way.

Clicking with the MENU button on the iconbar icon leads to the iconbar menu.

The first item on the menu gives access to some information about the program.

The second, 'Play', is shaded because there is no current programme available. The third opens a small panel which provides you with 5 'tape-recorder' style icons including play, pause, stop etc.

Choosing the 'Quit' option, by clicking on it with the SELECT button, stops any playing that may be in progress and quits the program.



Setting up a programme

To set up a completely new programme, click on the iconbar icon using SELECT. This opens the main programme window. Alternatively, if you have opened the panel, you can open the programme window by clicking on the Perform icon in the far left of the Panel window.

Each piece on the programme has two (initially blank) boxes, the second of which holds the filename of the piece to be played. To fill in this field, open one or more directory displays which contain the scores you wish to play and 'drag' the file icon into the appropriate box. Using the SELECT button. When you release the mouse button the file name appears in the box and its small filetype icon appears on the right.

(Note that although only the leaf name shows in the box, the full filename is held by the programme. If the programme and its files are held on the same disc then the names are stored relative'.. In this case, if you copy the programme and files to another disc and maintain the same relationship between the files then the programme will run. If the files are not stored on the same disc as the programme then the full pathnames, including filing system and disc names, are stored. This means that you can construct programmes with pieces from any combination of hard and floppy discs, with the proviso that the score must be available when it is needed by the program. However, if you move any of the files the programme will be unable to find them. Note that dragging a file icon into the programme box does not load me selected file. This happens later, only when it is needed.)



To rearrange the order of pieces, use the same technique - ie 'drag' the filename from one box to another. If you use the SELECT button, the filename is moved from one box to the other. If you use the ADJUST button, the filename is copied, leaving the original in place. Note that if you drag a filename onto a box which already has a name, that name is overwritten.

To delete a piece, 'drag' the filename onto the 'bin' icon at the top of the window.

While you are constructing your program, you will notice that the first box always contains a tick. This means that this piece is automatically flagged for playing. If you want a performance to omit a certain piece (but not

to delete it from the programme entirely) toggle the tick on and off using the SELECT button. The 'loudspeaker' icon is non functional, it is a reminder of what this tick means.

Perform recognizes four types of file object:

- Rhapsody files (filetype &COO)
- Midi files (filetype &FD4)
- Sample files (filetype &D3C) - (signed linear format)
- Programme files (filetype &CB9)

In the last case, Perform will load and perform a whole new programme. Effectively this means that Perform can perform an unlimited sequence of pieces with, if you like, announcements in between!


Perform also accepts any other file and when its turn comes round it is *RUN. This means that it can load new voices and execute obey files etc. For example, if you need to load a voice module called OrchLib before a piece, simply include the file in the list at the appropriate place. If you need to load several voice modules at once, a more efficient way would be to list them all in an obey file, and then include the name of the obey file in the programme.

Since Rhapsody version 1.20 now recognises voices by name rather than by number, it does not matter in what order the voices are loaded. On the other hand, you cannot load more than 32 voices at once. It is for this reason that the !Perform directory contains a short utility called 'No Voices' which removes all the voices from 10 to 32 (ie it leaves the default voices of WaveSynth, StringLib and Percussion intact). It is a good idea to include this utility at the head of any programme or obey file which loads any new voices.

Saving a programme

Click with MENU anywhere on the main programme window to open the main menu.

Moving over the Save option leads to the Save window. You can save your programme in the usual way either by entering a full pathname and clicking on the OK icon or by dragging the file icon into a directory viewer.



Playing a programme

Before choosing Play on the main menu, check the available options by moving over the options item to the Options menu.

Most of these are the same as in Rhapsody. The only one which is different is 'Continuous'. When this option is ticked, the programme repeats continuously. Do not confuse this with the 'Repeats' option which enables repeated sections within a piece to be played, or not as the case may be.


Options are saved along with the programme.

When the options are set correctly, choose 'Play' from the main menu and away you go. Note that MIDI files can only be played on a MIDI instrument. If the MIDI option is not selected (or if a MIDI interface is not fitted) then the MIDI files in the programme are ignored.

If, while a programme is being played, you wish to skip to the next item on the programme, choose the 'Skip' option.

To stop playing at any time, open the main menu and choose the third item which will now say 'Stop play'.

When appropriate, 'Play' and 'Stop play' appear on the iconbar menu as well,



The Panel window

You can open the Panel window by choosing the appropriate option on the iconbar menu.



As has been stated previously, clicking on the Perform icon at the far left opens the main programme window. The other icons are, in order from left to right:

- Skip back to previous piece
- Play from the start
- Skip on to the next piece
- Stop playing
- Pause if playing (toggles on and off)

Playing Samples

Perform will also play sampled sounds. It expects the sample to be in linear signed format with the first byte indicating the sample rate in micro-seconds. This is the format produced by Armadeus. The expected filetype is &D3C.

The Vox Box

Overview

The Vox Box is a suite of three simple but powerful programs which you can use to create and edit relocatable voice modules for use in programs such as Maestro, Rhapsody, Studio 24+ or programs of your own.

A relocatable voice module is a self contained piece of machine code which actually calculates the numbers which must be sent to the Archimedes Digital to Analogue converter in order to produce a sound from the loudspeaker. Over 20,000 of these numbers must be generated every second for each voice, but owing to the exceptional speed of the ARM microprocessor, it is possible to have 8 voices sounding simultaneously and there is still time left over for the computer to run several programs as well!

Some voice modules provide only one voice (eg WaveSynth) while others may provide several different, but usually related voices (eg StringLib and Percussion as supplied in ROM). Some modules calculate the data mathematically, others may use data in the form of a sample which may have been captured originally from a real instrument. Modules of the former type are generally very economical with memory but sound rather synthetic and lifeless; those of the latter type can require large amounts of memory.

The three programs offered here provide you with the means to create a range of different types of module to suit your particular needs. Each program produces its own type of module and you can load modules back into the appropriate program for further editing later if you wish. In order to use VoxBeat and VoxSample effectively, you will need some sample data. A wide range of musical samples is provided on the sample disc to get you started. Other samples can be obtained from SoundTracker modules, or you can capture live samples using a sampling board and a suitable program such as Armadeus.

VoxBeat

VoxBeat is probably the simplest of the three programs and is used for playing sampled sounds. It is primarily designed to provide facilities for a drum kit. It uses only one Archimedes voice (you are limited to 32 voices in all) but it assigns a different sample to each of the 12 notes of one octave. So middle C could play a bass drum, D could play a snare drum, E a cymbal etc.

It is also possible to employ up to 7 VoxBeat modules at the same time, each assigned to different octaves, giving a possible total of 96 different samples, more than one for every note on a typical keyboard! Of course, samples occupy a lot of memory space so this is not very practical.

A totally different use for VoxBeat would be as a speech synthesizer. For example, 10 different notes could be used to say the digits from 0 to 9. The computer could then be used to speak lists of numbers.

Alternatively, a VoxBeat module could contain several special sound effects for use in a theatrical performance. These could then be activated at the touch of a key or at predetermined times. The possibilities are endless.

VoxSample

If you try using a simple sample player module such as those produced by VoxBeat or Armadeus to simulate the sound of an orchestral instrument or a piano you run into three serious problems.

In the first place, it is difficult to get the pitch right. Naturally, the sample will only play back correctly if it was sampled at the same rate as it is being played back. For various reasons, this is not always the case. so some way has to be found of standardising the pitch of a sample.

Secondly, a simple sample player cannot play a note which is longer than the original sample. This does not matter with samples which are inherently short such as drum sounds or harpsichord sounds, but it is a problem with orchestral and piano sounds which appear to be cut off in mid stream when playing long notes.

Thirdly, when a short note is played, a simple sample player either cuts the note off abruptly or insists on playing the whole sample through regardless.

VoxSample attempts to get over all these problems in the following way. It uses the fact that the characteristic sound of an instrument owes far more to the initial 'transient' as it is called, than the sound of the sustained note which follows it. A VoxSample sound is made up of two parts - first the transient is played, then one or more cycles of the sample are looped continuously to form the sustained section. This approach is extremely economical in memory as it is rarely necessary to use a sample of more than about 0.5s in length. Moreover, the sustained section is guaranteed to be at the correct pitch. Facilities are also provided to alter the pitch of the

initial transient and amplitude and pitch envelopes can also be added to the sample.

VoxSynth

Lastly we come to VoxSynth. This is a moderately complex program and in order to use it effectively, you will have to learn something about the theory of harmonic sound synthesis. In brief, however, we can summarise its purpose as follows. Each VoxSynth module provides 8 independent voices. Each voice is defined by 2 waveforms, each representing one complete cycle. These waveforms may be drawn by hand, synthesised from harmonic components or created mathematically using a form of FM synthesis similar to that used by many digital synthesizers. At the start of the note, the first waveform is used, but gradually the wave changes into the second as the note proceeds. The technical term for this is 'Real Time Linear Interpolation'.

VoxSynth will also load real samples and uses the linear interpolation technique to imitate the sampled sound as closely as possible using real time synthesis!

As with VoxSample, amplitude and pitch envelopes may be superimposed on the sound.

The advantage of VoxSynth is that the modules are extremely compact, taking up less than 5k for 8 voices. It is capable of producing a vast range of different sounds similar to those produced by digital synthesizers. It is also of great value in learning about what makes different sounds different, and is great fun to experiment with.

Playing sounds

All three programs have facilities for playing the sounds while you are editing them either using the computer keyboard, or, if you have a MIDI module, using a MIDI keyboard. It is now possible to buy simple keyboards with MIDI sockets for less than £100 and it can be argued that the combination of such an instrument and the Vox Box is equivalent to a digital synthesizer costing many times as much.

In addition, because the Vox Box programs edit the modules directly, it is possible to edit the sounds while they are being played! In other words, you can be playing your favourite piece using Rhapsody and fine tune the sounds to be exactly what sounds best!

Non Musical uses

These programs, particularly VoxSynth, are very useful for exploring sound and waveforms in a scientific context. It enables a physics teacher to bring the subject of waveforms to life in a controlled manner without using an expensive oscilloscope.

Common Features

Getting Started


All three programs are standard RISC OS utilities and are loaded by double-clicking on their icons in the usual way. A default sample module is also loaded.

If you wish to load a previously created module appropriate to the program, then drag the Voice Module icon onto the program icon on the icon bar. If the file is not recognized as a voice module, it is treated as a sample if appropriate,

The normal way to load a sample into a program is to open up the main window by clicking on the iconbar icon using the SELECT button, select the required voice and then drag the Voice Module icon onto the waveform window.

The Iconbar menu

The iconbar menu is the same for all three programs and looks like this:



The first item gives access to information about the program.

If you do not have a MIDI podule the next two items are shaded. If you do have a podule fitted clicking on 'Midi in' enables you to play the internal Archimedes voices from an external keyboard. All 8 voice channels are allocated to this task so that you can play polyphonically.

If your keyboard is touch sensitive, you can use this facility by clicking on 'Midi touch'. By default, it always plays at maximum volume.

The last item, 'Quit', is of course used to exit the program.

Using the keyboard

Providing that the program has the input focus, ie when the window header is cream coloured, you can use the computer keyboard as a simple monophonic piano. Two chromatic octaves are provided laid out as follows:

S	D	G	H	J	2	3	5	6	7					
Z	X	C	V	B	N	M	Q	W	E	R	T	Y	U	I

with Q as middle C. Pressing the SHIFT key at the same time moves the pitch up two octaves.

If the keys do not seem to work, and the title bar of the main window is grey, simply click anywhere in the main window to obtain the input focus. The title bar will turn cream and the keys will work.

Saving your voice


When you have edited a voice to your satisfaction. click on the main window using the Menu button. This brings up a menu which includes the items Info and Save. Note that VoxSynth is the exception as it has no Info option.

Moving over the 'Info' arrow reveals a box giving some information about the voice you are about to save, including its length. (Not applicable to VoxSynth as all modules are 5k in length.)

Moving over the 'Save' arrow pops up a standard Save Dialogue box with a default name. Alter this if you wish, then either drag the file icon into an appropriate directory display or click on the OK box.

VoxBeat

The main window



The important icons are, in order from the top left:

No.

This displays the number of the sample displayed (1-12) and the pitch which it is assigned to, C4 is middle C. Clicking with SELECT or ADJUST moves up or down through the samples.

Name

There are two name icons. The first is the name of the complete Voice Module;

the second is the name of the particular sample displayed.

Play/Compile

Clicking on this icon plays the sample. If it has turned red, it compiles the sample before playing it. This has to be done if you have altered the sample in any way, eg by marking a different area.

Waveform window

This displays the complete sample. Various parts of it may be displayed in more detail using the red scroll bar below. Any part of it may be marked, and this area is shown in green. Unmarked areas are grey. For details see below.

Red scroll bar

This is used to move around the displayed sample. For details see below.

Type

VoxBeat can use samples in three different formats: signed linear, unsigned linear and logarithmic. Click on this icon with SELECT or ADJUST to select the appropriate type. If you do not know the type of sample, it is usually pretty obvious from the display which type to select.

Octave

The 12 samples are assigned to each of the 12 notes in an octave. This icon determines which octave is used. Use SELECT or ADJUST to display the desired octave in the range 1-7. Octave 4 is the one whose lowest note is middle C.

Vol

When each of the samples is compiled, you may choose a scaling factor from 25% to 400%. This will determine

the volume of the sample.

Pitch

This icon is used to transpose the sample when it is being compiled. The range is 24 semitones up or down in steps of 1/10th of a semitone. Click SELECT or ADJUST on the up and down arrows to increase or decrease the interval by one semitone. Click SELECT or ADJUST on the number itself to increase or decrease the pitch by 1/10th of a semitone.

Reset

The blob icon resets the pitch interval to zero.

The four grey icons at the bottom give useful information about the size of the original sample, the compiled sample, the module size and lastly, the amount of free space left for the module in the RMA.

Note: If you are in danger of running out of spare room in the RMA, save your module and exit VoxBeat. Allocate more memory to the RMA by entering the Task Manager and dragging the slider for Module area. Then reload VoxBeat and your partly edited module.

Loading and editing a sample

Before loading a sample, make sure that the window is displaying the required note in the icon at the top left hand corner of the main window. As mentioned above, each VoxBeat module can hold up to 12 samples, one for each note of the octave. To cycle through the samples, click on this icon with SELECT or ADJUST.

To load a sample, drag its file icon onto the waveform window. If VoxBeat does not recognise it as a sample, it will give you a warning but you can load any file as if it was a sample - eg SoundTracker modules.

At this stage, the whole of the sample is marked for use and is automatically compiled. By clicking on the 'Play' icon you can hear what it sounds like.

If you want to select a small part of it, click on the start of the desired section with the SELECT button and hold it down whilst dragging the outline box to cover the required area. When you release the SELECT button, the selected box is highlighted in green. If you use the ADJUST button instead of the SELECT button, you will find that you can make adjustments to whichever end of the marked section is nearer to the pointer when the button is pressed.

If the sample is a long one, you may want to examine some of it in more detail. To do this, click on one of the white arrows at the end of the red scroll bar and drag it towards the middle. You will see the sample open out. You can drag the scroll bar from side to side and if you click to the right or left of the shortened scroll bar, it moves to the right just like ordinary scroll bars.

Finally, you may wish to select a short section of a sample to save memory. One problem is that the sample will sound as if it has been cut off in the middle. To alleviate this, double click somewhere inside the marked (green) section. You will see that the right hand side of the sample is now tapered off and should give a smoother finish to the sound. To remove this effect, double click somewhere to the right of the marked section and click on Compile.

Sample types

VoxBeat can deal with three types of sample. A signed linear sample contains numbers in the range -128 to +127. This is the format used by most Archimedes sampler programs such as Armadeus. Unsigned linear samples using numbers in the range 0 to 255 are sometimes used on other computers. The Archimedes sound system actually uses samples in logarithmic format and this is the format in which VoxBeat compiles its samples. When you load a VoxBeat module back into VoxBeat, you will therefore find that all the samples are in this format.

When you click on the sample type icon to select the desired format, no calculation is done on the data; instead the program simply notes how the data is to be interpreted when it is subsequently compiled. The display, however, is updated to reflect the change.

If you do not know what format the sample is in, it is usually easy to tell by cycling through the three possibilities using the SELECT or ADJUST buttons.

Changing the pitch of a sample

Sometimes a sample will appear to play back at the wrong pitch. This is because VoxBeat assumes that the

original sample was recorded at the default rate of 20,833 Hz. If the sample sounds wrong, try transposing it up or down using the 'Pitch' group of icons. The pitch interval is given in semi-tones and tenths of a semi-tone. There are 12 semitones in an octave and the allowed range is -24 to +24 semitones. Note that the transposing only takes effect when the sample is compiled. The original data is not affected in any way.

Sometimes interesting effects may be obtained by transposing a correctly recorded sample. Note that the quality of a sample will be reduced when transposing down more than one octave.

Changing the volume of a sample

When putting together a collection of samples, it is important to get the balance between the different samples right. This can be achieved using the 'Vol' icon. By clicking SELECT or ADJUST you can select a scaling factor between 25% and 400%. As you do this, you will see the waveform display change accordingly. As usual, however, the scaling only actually takes place when the sample is compiled. Beware of introducing distortion when increasing the settings.

Allocating the right octave


The 'Octave' icon determines the base note of the octave to which the 12 samples are allocated. Octave 4 is the one whose lowest note is middle C. The range of octaves is 1 to 7.

Changing the names

Before you save your Voice Module, enter suitable names for each of the samples in the 'Sample' icon (eg BassDrum) and a suitable name for the complete module in the 'Name' icon (eg DrumKit1).

Displaying information about samples

If you click with the Menu button on the main window, a second menu appears which looks like this:



The first item on this menu gives access to a window which lists all the current samples and the memory they take up. This is useful if you want to economise on memory as it quickly identifies those samples which use a lot.

Saving your module

This is done in the usual way using the window which appears when you move over the arrow to the right of the 'Save' item on the menu.

Removing a sample

To remove a sample completely from memory, select the last item on this menu. The sample is then marked as <empty>.


Hints on using VoxBeat

Some ideas on using VoxBeat were given in the 'Overview' section.

When building a drum kit, it is a good idea to have related drums such as snares and toms on the white keys, with cymbals and blocks on the black. The bass drum will probably have to be boosted using the 'Vol' option as bass frequencies contain relatively little energy. Cymbals, on the other hand, may have to be cut.

VoxSample

The main window



The important icons are, in order from the top left:

Sample type

VoxSample can use samples in three different formats: signed linear, unsigned linear and logarithmic. Click on this icon with SELECT or ADJUST to select the appropriate type. If you do not know the type of sample, it is usually pretty obvious from the display which type to select and if you are not sure then playing the sound will confirm which is correct.

Voice name

This is the name of the Voice Module. This icon is writable.

Play/Compile

Clicking on this icon plays the sample. If it has turned red, it will compile the sample before playing it. This has to be done if you have altered the sample in any way eg by marking a different area.

Waveform window

This displays the complete sample. Various parts of it may be displayed in more detail using the red scroll bar below. Any part of it may be marked, and this area is shown in green. Unmarked areas are grey. The looped section is shown in cream. For details see below.

Red scroll bar

This is used to move around the displayed sample. For details see below.

No of loop cycles

This group of icons determines the size of the looped section. Clicking with SELECT or ADJUST on the up and down arrows increases or decreases the number by 10. Clicking on the number itself increases or decreases it by 1. You can select any number of cycles from 1 up to the number of cycles which the sample contains.

Pitch

This icon is used to transpose the sample when it is being compiled. The range is 24 semitones up or down in steps of 1/10th of a semitone. Click SELECT or ADJUST on the up and down arrows to increase or decrease the interval by one semitone. Click SELECT or ADJUST on the number itself to increase or decrease the pitch by 1/10th of a semitone. For more details on this, see below.

Reset

The blob icon resets the pitch interval to the default state which was determined when the sample was loaded.

At the bottom of the window are three large icons used for setting the amplitude and pitch envelopes. For more details, see the section on envelopes in the VoxSynth chapter. This is common to both VoxSample and VoxSynth.

Loading a sample

If you wish to load a new sample, open the directory display containing the file you wish to load. If the main VoxSample window is open, 'drag' the file icon anywhere inside the window. Alternatively, drag the file icon onto the VoxSample icon on the iconbar. This automatically opens the main window.

Note that VoxSample accepts any kind of file whether it is a sampled sound or not. If you load a BASIC program into VoxSample, you are not likely to produce any musical sounds from it but SoundTracker modules, for example, contain useful samples which you can extract and use to create new sounds.

If the file you are loading is in fact a VoxSample module, then VoxSample automatically extracts the envelope and loop parameters and effectively returns you to the point at which the module was created (with the difference that any data originally outside the markers is lost). When a new sample is loaded, VoxSample makes an attempt to determine the pitch of the sample. It does this by trying to identify 10 recognisable cycles somewhere in the middle of the sample. If it is successful, it marks the last complete cycle in the sample for looping and remembers the transpose interval that will be required to standardise the pitch. For economy and efficiency, the VoxSample module expects the data to have a cycle length of 64 bytes. The note which has this cycle length at the default sample rate of 20833Hz is the A flat below middle C. This explains why, when loading a correctly sampled note at a pitch of middle C, VoxSample indicates that it has to be transposed by 3.8 semitones. It is really assuming that the note is always an A flat!

In order to be able to determine the pitch effectively, the sample must have some recognisable cycles in the middle. If the sample does not have a recognisable pitch, then you get the following message:

Sample has indeterminate pitch: using default value.

Even if the sample you have loaded does have a pitch, all is not lost as you can adjust the pitch and the length of the loop cycle anyway.

Marking the sample

The first thing to do is to mark the section of the complete sample which you actually want to use. This is done by clicking on the start of the desired section with the SELECT button and holding it down whilst dragging the outline box to cover the required area. When you release the SELECT button, the selected box is highlighted in green. If you use the ADJUST button instead of the SELECT button, you will find that you can make adjustments to whichever end of the marked section is nearer to the pointer when the button is pressed.

If the sample is a long one, you may want to examine some of it in more detail. To do this, click on one of the white arrows at the end of the red scroll bar and drag it towards the middle. You will see the sample open out. You can drag the scroll bar from side to side and if you click to the right of the scroll bar, it moves to the right just like ordinary scroll bars.

Changing number of loop cycles and pitch

The next thing to do is to set the number of loop cycles and the pitch. To begin with, set the number of loop cycles to 1 and click on the 'Play/Compile' icon to hear what it sounds like. If the program was successful in

finding a suitable pitch automatically, it should not sound too bad. If the pitch is wrong, however, the transient section at the beginning will be at the wrong pitch. The looped section will, however, always be at the right pitch because this is determined by the rate at which it is looped - not by the sample data the loop contains.

If necessary, therefore, bring the pitch into line by adjusting the 'Pitch' group of icons.

You can now adjust the number of loop cycles to find the best sounding compromise. If you use 1 loop cycle, when the note enters the looping phase, every cycle suddenly becomes identical. This is often quite noticeable. Suddenly the note sounds synthetic and lifeless. If you loop 2 or more cycles, the note retains more of the character of the original sample but you run into another serious problem. When using 2 loop cycles, you will find that a second note appears when using 3, 4 or more cycles. As you increase the number of loop cycles the frequency of this harmonic decreases until it eventually becomes an audible pulsing of the note. In general it is best to use a number of cycles which is a power of 2 (eg. 2, 4, 8, 16 etc.) In this way the harmonic is a least related to the fundamental. Whether you go for a short loop or a long one depends very much on the sample you are using and the type of music which you want to use it with. It should be remembered that it is the transient at the beginning which gives a note most of its character; the loop phase is not, in practice, going to be called upon that often.

Obviously, whenever you change the number of loop cycles, the cream section of the waveform window reflects this change. Slightly less obviously, if you change the pitch of the sample, the length of the loop section also changes. This is because, when you change the pitch, you are changing the length of what is deemed to be one cycle.

Changing the volume

The volume of the sample can also be adjusted. The volume scaling factor can be varied from 25% to 400%. Whenever you click on this icon the waveform window reflects the alteration but the scaling only takes effect when the voice is compiled.


Whenever you make any alterations (except changing the number of loop cycles) the voice module has to be re-compiled. Test your new voice by clicking on the 'Play' icon or using the keyboard.

Editing the amplitude and pitch envelopes

For details on this, see the appropriate section on page 23.

VoxSynth.

The main window



The important icons are, in order from the top left:

No.

Each VoxSynth module contains eight voices numbered 1 to 8, and each voice has two waveforms, A and B, associated with it. This icon is therefore divided into two pans; the first pan is the voice number (1-8) and the second pan is the waveform (A or B). To cycle through the voices or the waveforms, click on either pan of the icon using the SELECT or ADJUST buttons.

Time

When the note is played, waveform A gradually turns into waveform B. This icon determines the time taken for this process to occur from 8s to 1/32s. If you do not want to use waveform B, you should set this icon to 'off'.

Name

There are two name icons. The first is the name of the complete Voice Module;

the second is the name of the particular voice displayed.

Play

When you click on this icon, the program sounds a note with a pitch of middle C. Unlike the other two programs, VoxSynth compiles the module immediately any changes are made so this icon does not change.

Waveform window

This is a green 'Oscilloscope' type display showing the synthesized waveform. You can draw in this window freehand using the SELECT or ADJUST buttons. For more detail see below.

Harmonics window

This has a cream background and shows a frequency spectrum containing the first 32 harmonics of the displayed waveform. You can add harmonics to this window using the SELECT button and remove them using the ADJUST button. For more detail see below.

FM synthesis parameters

Below these windows is a panel containing eight icons and an OK icon which can be used to synthesize complex waves. For more details on these icons, see below.

At the bottom of the window are three large icons used for setting the amplitude and pitch envelopes. For more details, see the section on envelopes later in this chapter.

Creating new waveforms

There are three ways of creating and modifying a waveform: you can draw it freehand, create it by adding harmonics or calculate it using a form of FM synthesis. Together, these three methods combine to provide exceptional power and flexibility in designing your own new waveforms. At all times, the oscilloscope display shows the shape of the current waveform and the frequency spectrum display analyses it (using a mathematical process known as the Fast Fourier Transform or FFT) into its harmonic components. The significance of this window will become clearer as we experiment with different synthesized waveforms. We shall, in fact, begin with harmonic synthesis as it is the most fundamental approach.

Using harmonic synthesis

Mathematically speaking, the simplest possible wave is a sine wave. If you have just loaded VoxSynth, you will see that voice 1 is a simple sine wave and its harmonic spectrum consists of one frequency only - called the fundamental.

Voice 2 has two frequency components, the fundamental plus one an octave higher. The waveform has one extra 'wiggle' in it.

Voice 3 contains the fundamental and the third harmonic. (The fundamental is itself the first harmonic.) This time the waveform has 3 'wiggles' and, if you have a good ear, you may be able to hear that the note sounded actually consists of middle C plus another note an octave and a fifth higher (G').

Have a look at voices 4 to 8 and hear what they sound like. None of them are very interesting as they contain just the fundamental plus one other harmonic.

You may have noticed that the frequency spectrum window is divided up unequally. Each division is, in fact one octave. The first octave contains one harmonic, the second, two, the third, four, the fourth, eight and so on. This sequence of frequencies is known as the harmonic series.

Go back to voice 1 and try adding some more harmonics by clicking on the frequency spectrum window with the SELECT button. If you hold the SELECT button down and drag it across the spectrum, you can add a whole range of harmonics at once. You will quickly appreciate that adding harmonics alters the quality of the note (ie how it sounds, not how good it is!)

You can alter the relative strengths of the harmonics by clicking on it with the SELECT button. If you click on one of the harmonics using the ADJUST button, that harmonic is removed. Similarly, if you drag the ADJUST button across a range of harmonics, they are all removed. Note that if you remove the largest harmonic, then all those that are left are scaled up automatically! v because it is only the proportions of the harmonics that matter. This explains why, if you try to delete all the harmonics, other ones appear to spring up at random. They were simply too small to (j;display at the resolution required by the largest harmonic.

A detailed treatment of the harmonic analysis of various instruments is out of place in this manual but here is a table of harmonic content of a few orchestral instruments with the intensity of each harmonic on a scale of 0-10. (When entering these into the frequency spectrum window, always begin with the maximum harmonic, ie. 10, so as to prevent the program from automatically scaling the others up.)

Harmonic	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
Instrument															
Flute	10	5													
Violin	10	4	2	1											
Oboe	2	3	4	8	10	2	1								
Clarinet	10	0	5	0	2	1	2	4	5	6	1	1			
Trumpet	10	9	8	7	6	5	4	3	2	1					
Piano	10	6	2	2	1	0	5	2	2	2	2	2	2		

Drawing freehand


If you drag the SELECT button from left to right across the waveform window, a new waveform is drawn which joins up the points you have selected fairly smoothly. You can also use the SELECT button to edit individual points.

Using the ADJUST button, you can insert straight line segments by 'rubber banding'. To do this press the ADJUST button, drag out the desired line and then release the ADJUST button.

It is extremely difficult to produce a really musical voice freehand because very subtle changes in the shape of a wave can introduce new harmonics which quite alter the quality of the voice. It is, however, of great interest to experiment with different wave shapes to see a) what they sound like, and b) what their harmonic content is.


Early in the nineteenth century, a French mathematician called Fourier discovered an amazing fact. You have already found out that by adding harmonics (ie sine waves of different frequencies) you can generate a huge number of different squiggly curves. Fourier discovered that any repetitive waveform can be made up in this way. When you draw a wave in the waveform window VoxSynth calculates all the amplitudes of all the sine waves (or harmonics) which have to be added together to reproduce the squiggles you have drawn. It is the amplitudes (or the first 32 of them at any rate) which are displayed in the harmonics window.

To try this out draw a square wave in the waveform window looking like this:



You will immediately notice an interesting thing about its harmonics - all the even harmonics are missing! Click on the 'Play' icon once or twice and remember the sound. (Just like the old BEEB eh?)

As always, the data is stored in the module in the form of a waveform. In order to get the program to recalculate the waveform using the harmonics which it has just calculated, click on one of the empty harmonics using the ADJUST button. This will simply delete an already empty harmonic. The reconstructed wave will look something like this:



Play the new sound. Do you notice any difference? You will probably find that the tone has lost some of its edge. This is because VoxSynth only calculates and uses the first 32 harmonics.

Can you reconstruct any wave at all? Draw a random squiggle in the waveform window and repeat the experiment. You will probably find that the reconstructed wave looks very different - but that it sounds much the same!!

The reason for this is that, in addition to calculating the relative strengths of the harmonic components of the wave, the FFT algorithm also calculates their 'phase'. The phase of a harmonic determines at what point in the cycle the wave starts relative to the others. The human ear is not very sensitive to phase so VoxSynth throws this information away when it reconstructs the wave. It is for this reason that all waves created by harmonic synthesis are symmetrical about the centre point. The reconstructed wave therefore looks different, but sounds the same.

Using the FM synthesis parameters

Below the waveform and spectrum windows is a group of eight icons called the parameter list. This feature of VoxSynth lets you create new waveforms using a technique employed by many expensive synthesizers known as Harmonic Synthesis by Frequency Modulation. You start off with a basic sine wave, generated by Operator 1. Assuming that you have just started VoxSynth, or you have just loaded the default module ('VoiceMod'), click on the 'Op1: freq.' icon once or twice. As you do so you will notice that the waveform changes its frequency. As usual, SELECT increases the frequency, ADJUST decreases it. Set the operator 1 frequency back to 1. Having set the basic frequency, we can now modulate it using the two icons on the right called 'Mod1: freq' and 'level'. Use the SELECT button to increase the value in the level box to about 30. As you do so, you will notice that the waveform gradually changes and new harmonics appear. The sound also gets richer and more interesting. One thing to note is that, with a modulator frequency of 1, the frequency spectrum is continuous - ie there are no missing frequencies.

Now increase the Modulator 1 frequency to 2. Immediately all the even harmonics disappear. With modulator frequency set to 3, all the third harmonics are missing.

Other interesting effects may be obtained when the basic (Op1) frequency is set to 2. When the modulator frequency is set to 1, all the harmonics are present, including the fundamental, as you would expect. When the modulator frequency is set to 2, only the even harmonics are present. This means that there is no fundamental and the note sounds an octave higher.

A particularly interesting effect is produced when the operator frequency is set to 2 and the modulator frequency to 5 (at level 10). The resulting frequency spectrum clearly has no fundamental but the note sounds like middle C all the same! The reason for this is that the wave contains strong components or harmonics 2 and 3. The difference in the frequencies of these two components is the frequency of the fundamental. The ear 'reconstructs' the missing fundamental from these two components;

Two independent operators, each with their own modulators are provided. The two waves are summed together in the proportions determined by the Operator 'amp.' icons. You will notice that these icons are coupled together so that increasing one automatically decreases the other. It is a good idea to create two

separate tones first, then combine them. You can use one operator to produce the basic tone, then the second to give it just the right amount of 'edge'.

One last thing. Whenever the waveform has been calculated using FM synthesis, the parameter icons are highlighted in red. As soon as you click on either the waveform or the frequency spectrum windows, the parameter icons return to their normal colour to indicate that they are no longer 'active'.

Using Real-Time Linear Interpolation

This is a very grand title for a very simple idea. Very few musical sounds consist of one unchanging waveform. Even an organ pipe (the easiest instrument to synthesize) has a transient at the beginning before it settles down to a continuous wave. A piano too starts with a very complex wave containing many high harmonics, but after a short while, the high harmonics die away leaving only a few behind. Sounds such as these can be synthesised by defining an initial waveform (A) and a second waveform (B), together with a time interval in which the first turns into the second. A good way to do this is to define a rich sounding waveform in wave A: copy it into wave B (see the later section on copying all or part of a wave); then, delete or reduce some of the higher harmonics. In this way the sound does not alter its character too drastically.

Note: Because of the amount of calculation required it may not be possible to play eight notes simultaneously with an interpolated voice. When interpolation is not needed make sure that the 'time' icon is set to 'Off'.

Loading samples directly

If you want to use VoxSynth to analyse the harmonic content of a real wave, or if you want to try to synthesize a real sound, you can load a sample wave directly into the waveform window from a sampling program such as Amadeus. If the sample has a definite pitch which VoxSynth can recognise, it selects a single cycle from near the beginning of the sample and puts it into waveform A, and a second cycle from near the middle which it puts into waveform B. In this way, VoxSynth can make a passable imitation of many types of sampled sounds. If VoxSynth cannot determine the pitch, it uses a default value.


If the file is less than 2k in length, it is obviously not going to contain a sufficient number of cycles to make sense of. VoxSynth therefore assumes that the file contains exactly one cycle and loads it directly into whichever waveform (A or B) is currently displayed. In this way, you can select individual cycles from a sampled waveform (using your sampling program) and then load them into VoxSynth to synthesize.

Copying all or part of a voice

In creating a new set of 8 voices, you may well want to copy either the waveform, or one of the envelopes from one voice to another. Alternatively, you may want to copy the whole voice so that you can make a small change (eg give it a little vibrato etc.) To do this, press the Menu button on the main VoxSynth window. This brings up the menu shown.

If you select the item 'Copy A -> B', the wave A is copied across into wave B in the same voice (together with its associated FM parameters). Similarly for the next item except in reverse.

If you move over the arrow to the right of any of the last four items, you come to the 'Copy' window. The two boxes contain the number of the voice you are currently editing. Use the SELECT and ADJUST buttons in the usual way to set the numbers to the appropriate values, then click on 'OK' and the voice is copied from the voice in the first box and the voice in the second box. Note that the voice name is not copied.



- **Copy voice** - copies the whole voice including FM and time parameters, amplitude and pitch envelopes.
- **Copy wave** - copies the wave data, FM and time parameters.
- **Copy amp** - copies the amplitude envelope only.
- **Copy pitch** - copies the pitch envelope only.

Editing the amplitude and pitch envelopes


For details on this, see the section below.

Amplitude and Pitch Envelopes

In addition to having characteristic harmonic content, most musical sounds have a recognisable shape to them. For example, a piano has a sharp attack, then the note dies away slowly. A pipe organ may take a little while to reach its maximum amplitude which then remains constant. In addition, some instruments have a tremolo effect (ie a repetitive variation in amplitude) while others may have a vibrato (a repetitive variation in frequency). All of these effects may be applied to VoxSynth and VoxSample voices using the amplitude and pitch envelopes.

Editing the envelope windows

The upper envelope window controls the amplitude of the voice, the lower one, the pitch. Beneath both windows is a time scale marked in 100ths of a second, 10ths of a second and seconds (in red). This time scale controls both the above windows but it is completely independent of the other parts of the main window (eg the waveform window and its scroll bar).



You may drag the time scale (and therefore the envelope windows in and out using the SELECT or ADJUST buttons.

Each envelope is divided into two phases, the sustain phase and the release phase. The release phase continues after the note is finished and usually determines how fast the note decays.

The two phases are separated by a black line. Each phase may contain up to six nodes (red or orange blobs) which are added by clicking with the SELECT button and removed by clicking on them with the ADJUST button. They can also be dragged around using the SELECT button.

Some of these nodes may be red and connected with red lines, others orange and connected with blue lines. Those in red form part of a loop - ie this part of the envelope is repeated over and over again until the end of the phase or the end of the note. Any section of the envelope may be made part of the loop but all the nodes following that section, up to the end of the phase, automatically become part of the loop and turn red. A section may be toggled from one state to the other by clicking on it with the ADJUST button, (Be careful not to click too close to a node or this will delete the node instead!)

When a loop section is repeated, the amplitude or pitch is not reset each time. Instead, it simply carries on from where it was. This means that if the end of the loop is higher than the beginning the amplitude or pitch will steadily rise. If the end is lower than the beginning, the amplitude or pitch will steadily fall. In order to construct a tremolo or vibrato at constant volume or pitch, you must make sure that the end of the loop is at exactly the same height as the beginning.

Owing to the way the envelope is stored in increments rather than actual values, you may find it impossible to place the final node in the right place. For example, if the pitch increments by 3 units for 15 centi-seconds it will change by 45 units in total. If you want to decrease it back to zero in 10 centi-seconds, you will find that it is impossible because 45 does not divide by 10. You will, however, find that it is possible to reduce it to zero in 9 centi-seconds because 45 divided by 9 is 5. In practice, you do not know the actual values of the increments

being used because their precise values are not very important. A little trial and error is in order here!

Normally, the amplitude envelope should rise to its maximum value in 1 centi- second (this effectively means instantaneously) and should have one continuous horizontal loop. The release phase should be one looped section sloping down at a suitable rate or one unlooped section sloping down to zero. The pitch envelope should normally be two horizontal looped sections.

Finally, note that the position of the amplitude release phase node need not be in the same place as the pitch release phase node. This means that the period of any tremolo (amplitude variation) need not be the same as the period of vibrato (pitch variation) and is a consequence of the fact that the two envelopes are totally independent. On the other hand, when the note is actually released, both envelopes will skip to the start of their respective release phases simultaneously.

In summary:

- To add a node, click with the SELECT button
- To adjust a node, drag it with the SELECT button
- To remove a node, click on it with the ADJUST button
- To mark a section as looped, or to turn looping off. click on it with ADJUST

Here are just a few of the other products you will be able to
download from RISCOS.com free of charge
(You can make a donation to help support future product releases but this is entirely voluntary)



This Documentation is © 3QD Developments Ltd 2026
Some of the products shown above might not be available immediately.
Not all products will run on all versions of RISC OS, check the product information page for details.

Other Products list last updated 29/2/26