

# Vocal Harmonizer and Vocoder

Philip Decamp  
Uriel Klieger  
Andrew McPherson

December 10, 2003  
6.101  
Prof. Ron Roscoe

# Abstract

We designed and built a vocoder and vocal harmonizer. The vocoder allows one sound to be manipulated to sound like another, such as to make white noise sound like speech. The harmonizer detects the frequency of the operator singing into a microphone and creates a set of output waveforms which form a musically significant chord around that frequency. Used as an input to the vocoder and modulated by the original voice, the sum of these waveforms sounds like a set of separate voices, all in perfect harmony. A comprehensive set of controls allows the operator to choose from a wide selection of available chords.

## Table of Contents

1.	Project Overview .....	3
1.1	Vocoder [Uriel].....	3
1.2	Harmonizer [Uriel].....	4
2.	Circuit Detail.....	5
2.1	Vocoder.....	5
2.1.1	Overview [Philip].....	5
2.1.2	Filter Banks [Philip].....	6
2.1.3	Envelope Followers [Uriel].....	12
2.1.4	Voltage-Controlled Amplifiers [Uriel] .....	13
2.2	Harmonizer .....	15
2.2.1	Overview [Andrew] .....	15
2.2.2	Microphone Preamp [Andrew] .....	17
2.2.3	Filters and Frequency Detector [Andrew] .....	18
2.2.4	Frequency to Voltage Converter [Andrew] .....	20
2.2.5	Harmony Selectors [Andrew] .....	21
2.2.6	Harmony Keypad Control Logic [Philip] .....	24
2.2.7	Voltage to Frequency Converter [Andrew] .....	26
2.2.8	White Noise Generator [Philip] .....	27
2.2.9	Mixer [Andrew] .....	28
2.3	Power Supply [Uriel] .....	29
3.	Measurements .....	30
3.1	Filter Bank [Philip] .....	30
3.2	Harmonizer [Andrew].....	30
4.	Errors and Corrections .....	35
4.1	Filter Banks [Philip].....	35
4.2	Envelope Follower / VCA [Uriel].....	35
4.3	Harmonizer [Andrew].....	35
5.	Conclusions [All].....	37
6.	Acknowledgements.....	40
7.	References.....	40
8.	Appendices.....	41

## Catalog of Figures

1	Vocoder Block Diagram	Pg. 3	Uriel
2	Harmonizer Block Diagram	4	Uriel
3	Single Filter Response	7	Philip
4	Filter Bank Schematic	8	Philip
5	Calculated Total Frequency Response	12	Philip
6	Envelope Follower and VCA Schematic	13	Uriel
7	Microphone Preamp Schematic	17	Uriel
8	Microphone Filter Schematic	18	Uriel
9	Frequency Detector Schematic	18	Uriel
10	Frequency to Voltage Converter Schematic	20	Uriel
11	Harmony Selector Schematic	22	Uriel
12	Keypad Logic Schematic	25	Philip
13	Voltage to Frequency Converter Schematic	26	Uriel
14	White Noise Generator Schematic	27	Philip
15	Mixer Schematic	28	Uriel
16	Power Supply Schematic	29	Uriel
17	Accuracy of F-V and V-F harmony system	34	Andrew

# 1. Project Overview

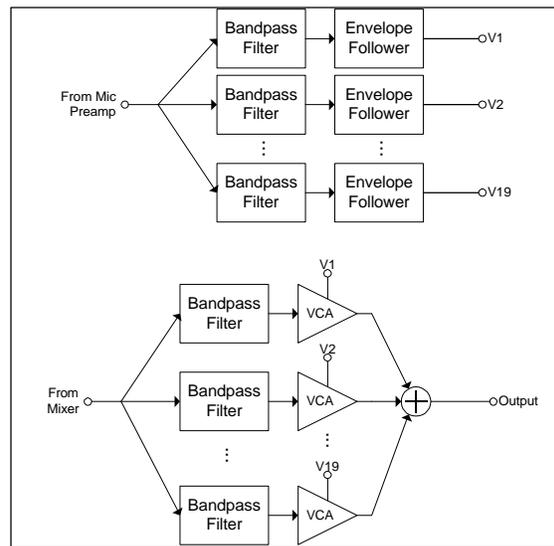
This project can be neatly divided into two parts: the vocoder and the harmonizer. The vocoder is a new implementation of an old idea, while the harmonizer is our own creative addition to the assembly.

## 1.1 Vocoder [Uriel]

The vocoder was invented in 1939 by Homer W. Dudley, a researcher at Bell Laboratories. Originally intended as a method of compression for phone signals, the unique sound of the vocoder has since found a very comfortable home in the effects libraries of popular music artists.

Our vocoder design actually follows the original design very closely. Two input signals are presented to the device: a control signal (usually a person’s voice), and an “instrument.” Both signals are then separated by an array of bandpass filters into nineteen different frequency ranges. We chose this number of ranges to ensure adequate coverage of the vocal range; more bands would give a better sounding output, but the tradeoff is that incredibly narrow filters are required.

The voice, being the control signal, then runs through an envelope follower, which generates a slowly-changing waveform that represents the strength of the voice in each frequency band. This output (V1 to V19 in the diagram) is passed as the control input to a set of nineteen voltage-



**Figure 1: Vocoder Block Diagram**

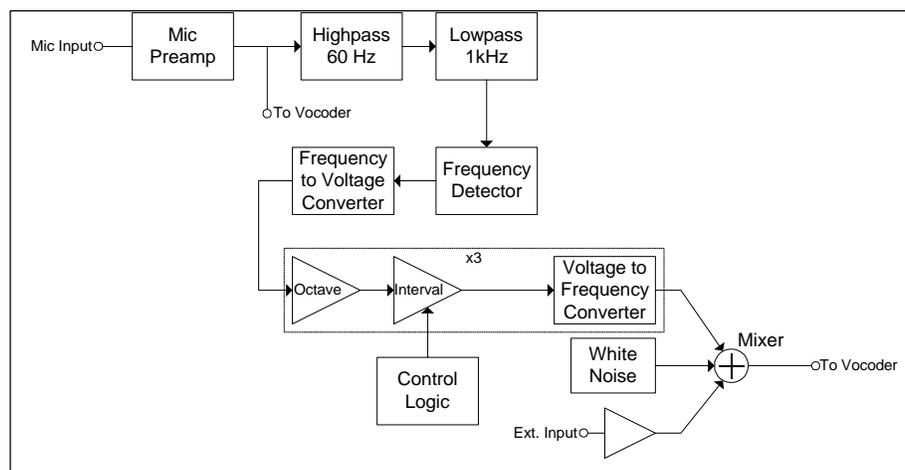
controlled amplifiers (VCAs). The VCAs set the gain of the instrument in each band based on the output of the envelope follower, essentially functioning as a multiplier.

The outputs of the VCAs are then mixed back into one signal, which is a purely synthesized version of the voice made by modulation of the instrument. The instrument can be pretty much anything; different inputs will result in a vastly different range of sounds at the output. The output signal is designed to be used with an external amplifier.

In Dudley's original design, the envelope follower output was then sent over the phone and used to synthesize voice from a noise generator at the receiving end. These slow-moving control signals required much less bandwidth to transmit than an actual voice.

## 1.2 Harmonizer [Uriel]

The harmonizer is our addition to the vocoder. Its purpose is to detect the frequency of the input signal, and generate a set of waveforms in musically useful related frequencies. It then mixes these generated signals with other possible vocoder instruments, such as a guitar, or white noise. A generator for white noise was built for



**Figure 2: Harmonizer Block Diagram**

this feature, as it is useful for production of fricative phonemes like *s* and *f*.

The microphone input is an XLR jack, differential input, which is amplified considerably in the mic preamp. This signal is then run through a small filter bank to filter out sounds outside of the vocal range. The bandpassed signal is passed through a circuit which detects its frequency and converts this to a voltage via a linear function.

This voltage is sent through a set of three harmony generators, which change the frequency of the input signal and output a pulse wave at a different frequency to the mixer stage. To do this, the harmony block amplifies or attenuates the voltage from the frequency to voltage converter, and then converts this voltage to the pulse wave in a voltage-to-frequency converter. Because of the linear frequency-voltage relationship, the proper change in voltage can raise the output a fifth, third, octave, etc. allowing the creation of chords. This block is controlled by a set of rotary switches and a button panel.

This mix of synthesized waveforms is sent as the instrument to the vocoder, which applies the voice characteristics, resulting in a set of harmonized voices. An optional connection not shown on the block diagram allows mixing in the original voice after the vocoder stage, for a total of four voices.

## 2. Circuit Detail

### 2.1 Vocoder

#### 2.1.1 Overview [Philip]

The vocoder takes two inputs: a voice or microphone input and an instrument input. Each input is fed into a large filterbank that separates it into 19 different frequency bands. The bands from the voice input are fed into an envelope follower,

which outputs a DC signal indicating the volume in that particular band. That is, an audio signal of a voice goes into the device and the vocoder generates 19 different signals that form a rough frequency spectrum of that input. The instrument input goes through identical filters as the voice input, but does not go through the envelope followers. For each filter there is a Voltage Controlled Amplifier (VCA) that multiplies one of the instrument frequency bands by the volume of the corresponding voice frequency band.

This schematic represents one of 19 such circuits in the vocoder, each one covering a different frequency range. The total output of the decoder is the sum of the outputs of all the VCAs. This output is essentially the instrument signal being dynamically equalized to imitate the frequency spectrum of the voice, which is how the vocoder can synthesize speech from arbitrary sounds.

While the voice input for the device is always a microphone signal, several possibilities are provided for the instrument input. These include a vocal harmonizer, which generates several tones from the user's voice and adds a chorus effect to the synthesized speech; a white noise generator, which equally covers all the frequency bands and is thus easily intelligible, and a jack that can be used for any kind of input the user wishes.

### 2.1.2 Filter Banks [Philip]

The vocoder has two separate, but identical filter banks. One is used for the voice input and the other for the instrument input.

Since the vocoder was intended to be used with a vocal harmonizer, it had to be precise enough to reproduce several, simultaneous notes without significantly obscuring

the tonality. For this reason, most of the filters only cover a fourth of an octave. 15 fourth-octave filters cover the majority of the vocal range from 250 to 3364 hertz. Two filters covering two-thirds of an octave are placed at 111 and 177 hertz to cover some of the lower sounds, while a full octave filter placed at 5187 hertz catches most of the higher harmonics. Altogether, each filterbank consists of 19 filters.

Another criteria in the filter design was to keep them as discrete as possible. For a vocoder, having the filters overlap causes the various frequencies of the speech synthesis to meld together and become less intelligible. Therefore, the filters need to fall off extremely quickly as the frequency moves away from the center frequency.

A third obstacle was the sheer volume of filters needed. Switch-capacitor filter ICs were originally considered because they are extremely easy to wire, but even the cheapest ICs would have cost hundreds of dollars. Elliptical filters could have been used to achieve extremely fast cutoff rates, but each filter would have required at least eight op-amps, making them prohibitively labor intensive as well as expensive.

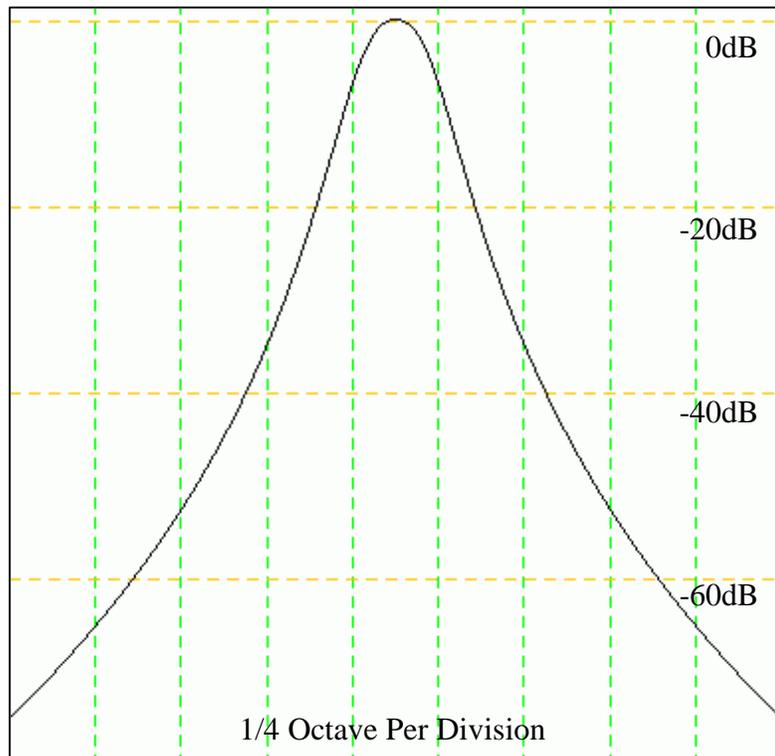
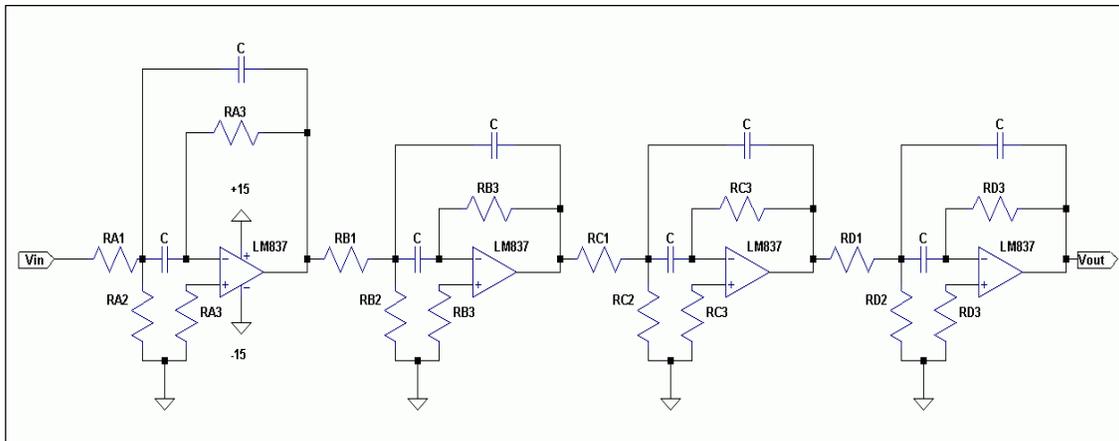


Figure 3: Single Filter Response



**Figure 4: Filter Bank Schematic**

The final design is a simple, fourth-order, "multiple-feedback" bandpass filter. The design was taken from Don Lancaster's *Active Filter Cookbook* on page 154. The transfer function for a single stage of this filter is straightforward to compute, albeit tedious. The result is  $-1/R1/C2 * s / (s^2 + 2/R3/C*s + (R1+R2)/R1/R2/R3/C^2)$ . By changing the values of R1, R2, R3, and C, it is possible to set the passband gain, Q, and center frequency of the filter.

But selecting a Q of about 8.5, the filters become extremely narrow and steep, thus fulfilling the requirements of being very discrete while only taking up a quarter-octave. The tradeoff for this is that the filters are very resonant, which results to a significant amount of ringing. To lessen this effect, the filters use staggered center frequencies. That is, the individual stages of each filter have slightly different center frequencies, which gives them a flatter passband and allows the filter to maintain a very fast cutoff while lowering the Q.

After the narrow resonant peak is passed, the rolloff of the filter is determined by its order. Fourth-order filters were used to achieve 24db of rolloff per octave. One of the drawbacks to using four stages is that, because they are all multiplied together, slight deviations in frequencies of each stage can lead to large deviations in the overall

amplitude. 1% resistors and, for the most part, 3% capacitors were used to keep the filters as exact as possible.

The resistor and capacitor values used for each filter and filter stage are shown below. Each block of data represents one filter. Each line in that block contains the data from a single stage or all the stages combined. Note that these are calculated, not measured, values.

Filter Stage	Center Freq(Hz)	Q	Gain	C (nF)	R1 (O)	R2 (O)	R3 (O)
Total	111	3.17	0.99				
A	97	6	1.39222	100	71500	1400	196000
B	111	2	1.39222	100	20500	4320	57600
C	128	6	1.39222	100	53600	1050	150000
D	111	2	1.39222	100	20500	4320	57600

Total	177	2.16	0.98				
A	143	6	1.6476	100	40200	953	133000
B	177	2	1.6476	100	11000	2800	35700
C	218	6	1.6476	100	26700	619	86600
D	177	2	1.6476	100	11000	2800	35700

Total	250	8.46	1.02				
A	237	6	1.2	100	34000	576	80600
B	264	6	1.2	100	30100	511	73200
C	237	6	1.2	100	34000	576	80600
D	264	6	1.2	100	30100	511	73200

Total	297	8.46	1.02				
A	282	6	1.2	100	28000	475	68100
B	314	6	1.2	100	25500	432	60400
C	282	6	1.2	100	28000	475	68100
D	314	6	1.2	100	25500	432	60400

Total	354	8.46	1.02				
A	335	6	1.2	100	23700	402	57600
B	373	6	1.2	100	21500	365	51100
C	335	6	1.2	100	23700	402	57600
D	373	6	1.2	100	21500	365	51100

Total	420	8.46	1.02				
A	398	6	1.2	100	20000	340	47500
B	444	6	1.2	100	17800	301	43200
C	398	6	1.2	100	20000	340	47500

D	444	6	1.2	100	17800	301	43200
---	-----	---	-----	-----	-------	-----	-------

Total	500	8.46	1.02				
A	474	6	1.2	47	35700	604	86600
B	528	6	1.2	47	32400	549	76800
C	474	6	1.2	47	35700	604	86600
D	528	6	1.2	47	32400	549	76800

Total	595	8.46	1.02				
A	563	6	1.2	33	43200	732	102000
B	628	6	1.2	33	38300	649	93100
C	563	6	1.2	33	43200	732	102000
D	628	6	1.2	33	38300	649	93100

Total	707	8.46	1.02				
A	670	6	1.2	33	35700	604	86600
B	746	6	1.2	33	32400	549	76800
C	670	6	1.2	33	35700	604	86600
D	746	6	1.2	33	32400	549	76800

Total	841	8.46	1.02				
A	797	6	1.2	33	30100	511	73200
B	888	6	1.2	33	27400	464	64900
C	797	6	1.2	33	30100	511	73200
D	888	6	1.2	33	27400	464	64900

Total	1000	8.46	1.02				
A	947	6	1.2	33	25500	432	60400
B	1056	6	1.2	33	22600	383	54900
C	947	6	1.2	33	25500	432	60400
D	1056	6	1.2	33	22600	383	54900

Total	1189	8.46	1.02				
A	1127	6	1.2	22	32400	549	76800
B	1255	6	1.2	22	28700	487	69800
C	1127	6	1.2	22	32400	549	76800
D	1255	6	1.2	22	28700	487	69800

Total	1414	8.46	1.02				
A	1340	6	1.2	22	26700	453	64900
B	1493	6	1.2	22	24300	412	57600
C	1340	6	1.2	22	26700	453	64900
D	1493	6	1.2	22	24300	412	57600

Total	1682	8.46	1.02				
A	1593	6	1.2	22	22600	383	54900
B	1775	6	1.2	22	20500	348	48700

C	1593	6	1.2	22	22600	383	54900
D	1775	6	1.2	22	20500	348	48700

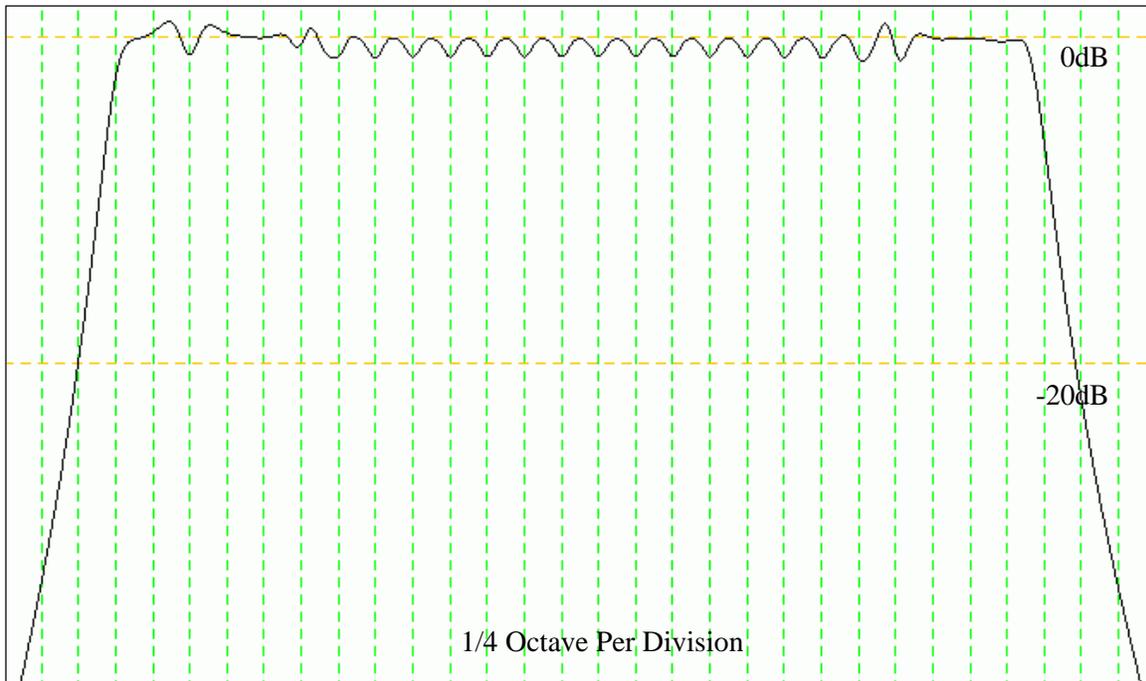
Total	2000	8.46	1.02				
A	1895	6	1.2	22	19100	324	45300
B	2111	6	1.2	22	16900	287	41200
C	1895	6	1.2	22	19100	324	45300
D	2111	6	1.2	22	16900	287	41200

Total	2378	8.46	1.02				
A	2253	6	1.2	10	35700	604	84500
B	2511	6	1.2	10	31600	536	76800
C	2253	6	1.2	10	35700	604	84500
D	2511	6	1.2	10	31600	536	76800

Total	2828	8.46	1.02				
A	2679	6	1.2	10	29400	499	71500
B	2986	6	1.2	10	26700	453	63400
C	2679	6	1.2	10	29400	499	71500
D	2986	6	1.2	10	26700	453	63400

Total	3364	8.46	1.02				
A	3186	6	1.2	10	24900	422	60400
B	3551	6	1.2	10	22600	383	53600
C	3186	6	1.2	10	24900	422	60400
D	3551	6	1.2	10	22600	383	53600

Total	5000	1.58	0.99				
A	3900	6	2	4.7	26100	750	105000
B	5657	2	2	4.7	6040	2000	23700
C	4757	2	2	4.7	7150	2370	28700
D	6900	6	1.8	4.7	16500	422	59000



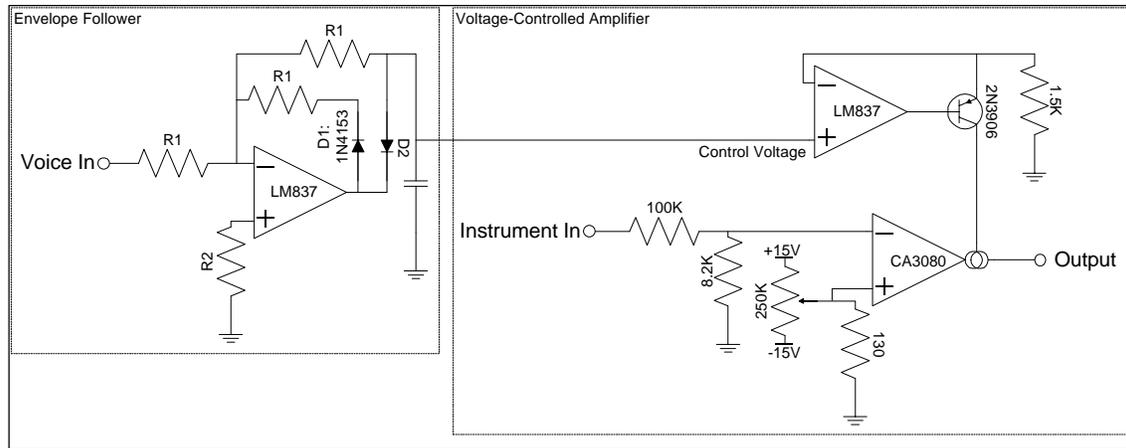
**Figure 5: Calculated Total Frequency Response**

### 2.1.3 Envelope Followers [Uriel]

Once the inputs have been filtered, the vocoder needs to modulate the instrument with the voice. This is handled by passing the voice through an envelope follower, and using this new wave as the control input to a voltage-controlled amplifier. The VCA works, for all intents and purposes, like a signal multiplier, setting the gain on the instrument input using the envelope follower output.

The envelope follower is essentially a simple diode/RC peak detector, except the diode has been replaced with an “ideal diode” constructed with an op-amp. The VCA control input requires a negative voltage, so this needs to be considered when discussing the ideal diode setup.

Referring to the schematic, when the voice in is positive, D2 is activated, and C is charged to a negative voltage using the op-amp as a current sink. When the voice in goes negative, D1 turns on and D2 turns off, giving the capacitor a discharge path



**Figure 6: Envelope Follower and VCA Schematic**

through R1 into the virtual ground at the inverting terminal of the op-amp. Because of this, the envelope follower has an RC time constant of  $R_1C$ . The values for R1 and C were chosen for a time constant approximately 50 times the period of the center of each band, in order to have good peak detection with as little ripple as possible.

Two diodes were used in the envelope follower in order to keep the op-amp from saturating at the rails when the input voltage switches polarity.

#### 2.1.4 Voltage-Controlled Amplifiers [Uriel]

The VCAs have two components: a current source, and an operational transconductance amplifier (CA3080). The current source is a 2N3906 driven by an op-amp. When the non-inverting input to the LM837 goes negative, the op-amp output goes negative, allowing current to flow through the transistor to the CA3080, and putting a negative voltage (due to emitter drop) on the inverting input, completing the feedback loop. When the input is positive, the op-amp output is positive, and the transistor doesn't conduct. The 1.5K resistor allows us to set the range of current output based on control inputs. We determined control inputs up to -13 volts were allowed, but we tried not to exceed -10V.

The CA3080 OTA has an adjustable gain, based on a biasing input current (from the 2N3906 in this design). The instrument input to the CA3080 is first sent through a voltage divider (approximately 12x attenuation) because the OTA requires a small input voltage to maintain linearity, in the range of 100mV. The OTA also has a potentiometer installed to adjust input offset, which was causing output noise.

The CA3080 outputs a current. We simply connected all 19 VCAs in parallel, using a resistor to ground (7.5k) to convert the current sum into a voltage. This resistor can be adjusted to control the total gain of all the VCA stages.

The following is a table of values used for R1, R2, and C:

Center Freq ( $f_0$ )	C (uF)	R1	R2
111	0.82	180K	91K
177	0.82	110K	56K
250	0.22	330K	160K
297	0.22	300K	150K
354	0.22	270K	130K
420	0.22	220K	110K
500	0.22	180K	91K
595	0.22	160K	75K
707	0.22	180K	91K
841	0.22	110K	75K
1000	0.22	91K	43K
1189	0.22	82K	43K
1414	0.22	68K	33K
1682	0.22	56K	27K
2000	0.22	47K	24K
2378	0.22	39K	18K
2828	0.22	33K	16K
3364	0.22	27K	13K
5187	0.22	18K	9.1K

This design for the VCA was based on a schematic in *Musical Applications of Microprocessors*, by Hal Chamberlin.

## 2.2 Harmonizer

### 2.2.1 Overview [Andrew]

On many popular recordings, the singer sounds as if he is singing along with himself, often a third, fifth, or even an octave above the main melody. Twenty years ago, this would have been done by recording on multi-track tape and going back to record another vocal part alongside the initial performance. These days, much of it is done electronically, with circuits designed to transform the voice to sound at a different pitch. Automatic harmony generation is very tricky to do well, and nearly impossible to do in analog. Nonetheless, we were able to construct a circuit that could generate chords to an input pitch with reasonable accuracy and reliability.

Most harmony generators work digitally, using an FFT on the input signal, multiplying the result in the frequency domain, and then converting back with an inverse FFT to obtain a pitch higher or lower than the input that still has the same timbre as the input waveform. This option was not available to us. The only good way to directly manipulate the pitch of a waveform in analog is to use a multiplier or ring modulator, where multiplying a wave of one frequency by another results in sum and difference frequencies at the output.

Unfortunately, this is not musically useful for two reasons. The first is that pitch is perceived logarithmically: doubling the frequency raises the pitch by an octave, multiplying the frequency by  $3/2$  results in a perfect fifth,  $4/3$  a perfect fourth, and so on. The interval generated by adding a constant to a frequency will vary (for example, adding 400Hz to a 400Hz wave will raise its pitch by an octave, but will only add a fifth to an

800Hz wave). The second problem is that any real-world waveform is full of harmonics, and shifting the frequency up by a constant will result in the harmonics no longer being multiples of the fundamental frequency, which completely distorts the sound.

As such, the system we implemented effectively cheats by not trying to alter the input at all but instead synthesizing completely new waves at the correct frequency, and relying on the vocoder to transform their spectra to match that of a voice.

The signal from the microphone, after being amplified by the preamp, goes through a 2nd-order high-pass filter at 60Hz and a 2nd-order low-pass filter at 1kHz to remove unwanted noise and harmonics, leaving only the fundamental frequency of the note being sung. (60Hz is around the C two octaves below Middle C, and 1kHz is near the C two octaves above it). The filtered signal then goes into a frequency detector which generates a clean square wave (of indeterminate duty cycle) corresponding to the positive peaks of the input waveform. The square wave drives a frequency-to-voltage converter, whose output is a DC voltage linearly proportional to the input frequency.

The output of the F-V converter goes into three parallel synthesis blocks (for a total of three different voices), where in each block the voltage can be multiplied by 2 or divided by 2, and then passed through a gain stage whose gain can be selected among 8 choices by a multiplexer. Finally, the voltage is put into a voltage-to-frequency converter calibrated identically to the F-V converter, such that if the output of the F-V were hooked directly to the input of the V-F, the input and output frequencies would be the same. The 8 gain choices correspond to the ratios in frequency of common musical intervals, so that a multiplication in voltage there corresponds to a multiplication in frequency at the output,

generating a voice a third, fourth, fifth, or sixth above the input pitch. Optionally multiplying or dividing the voltage by 2 allows a selection of three octaves.

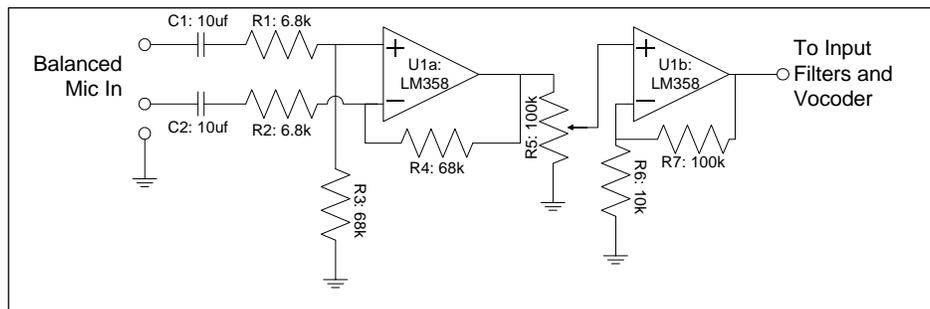
Finally, the outputs of the voltage-to-frequency converters (each of which is a pulse wave containing a wide variety of harmonics) are mixed along with white noise and an optional external input, and the result fed into the vocoder to be modulated. The level of each voice, along with the level of the noise and the external input, can be controlled independently with knobs on the front panel.

### 2.2.2 Microphone Preamp [Andrew]

The circuit is designed to use a balanced microphone, the standard for all professional recording. (Balanced microphones reduce noise by having a differential signal which is separate from the ground wire). U1a is configured as a differential amplifier. The gain of the inverting signal, with no signal at the noninverting input, is  $-R4/R2 = -10$ . The gain at the noninverting input with no signal at the inverting input is  $R3/(R3+R1)*(1+R4/R2) = (10/11)*11 = 10$ . Thus the output of U1a is 10 times the differential input.

Common-mode rejection is mainly determined by the tolerances of R1-R4. A high common-mode rejection reduces noise, which appears on both input signals.

Professional microphone preamps will be calibrated for minimum noise. In our case, the



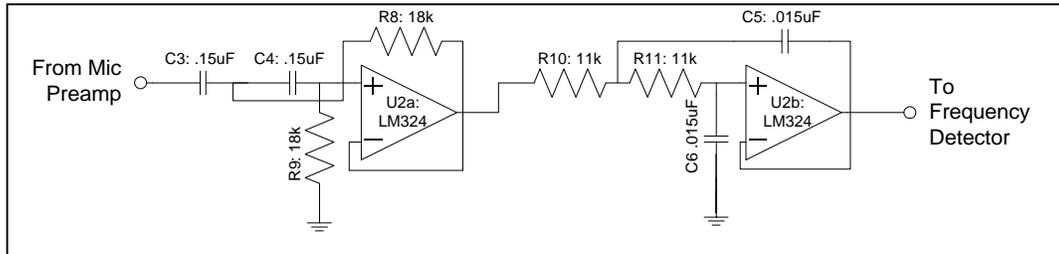
**Figure 7: Microphone Preamp Schematic**

SNR provided by 5% carbon resistors was sufficient, given the likelihood of noise to be picked up by other long wires or generated elsewhere in the project.

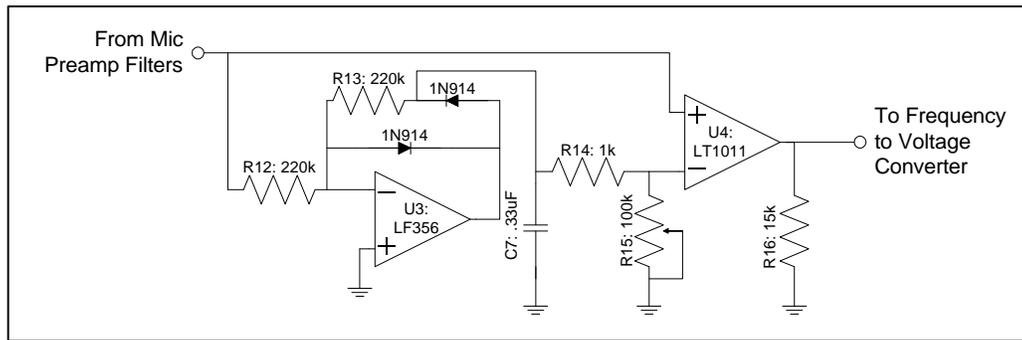
U1b is a standard noninverting amplifier with a gain of 11, for a maximum total gain of 110, or 41db. The potentiometer R5 (mounted on the front panel) was inserted between U1a and U1b so that only one pot would be needed to adjust the now single-ended signal, and so that a strong signal could be reduced in amplitude before it caused U1b to rail.

### 2.2.3 Filters and Frequency Detector [Andrew]

The frequency detector takes the output of the microphone preamp as its input. The filters are standard second-order Sallen-Key filters, obtained from Horowitz & Hill. U2a is a HPF at 60Hz, which in practice has a -3db point of  $1/(2 * \pi * C3 * R8) = 58.9\text{Hz}$  after selecting the nearest 5% component values. The LPF at 1kHz (U2b) works out to  $1/(2 * \pi * C5 * R10) = 965\text{Hz}$ . The precise values are not critical; both values were chosen



**Figure 8: Microphone Filter Schematic**



**Figure 9: Frequency Detector Schematic**

as a ballpark estimate of the extreme high and low ranges of the human voice, and the frequency of a waveform outside this range could still be detected even with mild attenuation. The point is to reduce the amplitude of the spurious noise and harmonics relative to the fundamental pitch. Two sections of an LM324 quad op-amp were used, because it seemed likely that further filtering might need to be added later.

The filtered waveform then passes into U3, which acts as an ideal diode with no forward voltage drop in a peak detector circuit. The improved version of the circuit from class is used which prevents U3 from swinging all the way to the rail on negative inputs. Since R13 provides a path from output to ground even on the negative half-wave, R12 and R13 had to be made very large to avoid reducing the RC time-constant formed with C7, and still provide a gain of (-)1.

Of course, the RC time constant is also affected by the series combination of R14 and R15, which form an (adjustable) voltage divider at the input of the comparator U4. (Perhaps R14 = 10k and R15 = 1 Meg would have been a better choice to reduce this influence, but a smaller time constant seemed to work well). The other input to U4 is the filtered wave before it enters the peak detector. The purpose of this setup is that on the positive peak of each cycle, C7 will be charging up and thus following the input voltage. Since the voltage divider of R14 and R15 reduces its amplitude by somewhere around 1%-10%, the positive input to U4 will be greater than the negative input and U4 will swing high (the collector at the output is connected directly to +15, with a pull-down resistor). The rest of the time, the input signal will be below its peak-detected version and the output will be low. The result will be a square wave whose frequency is precisely equal to the input waveform's frequency. This circuit was taken from *Musical*

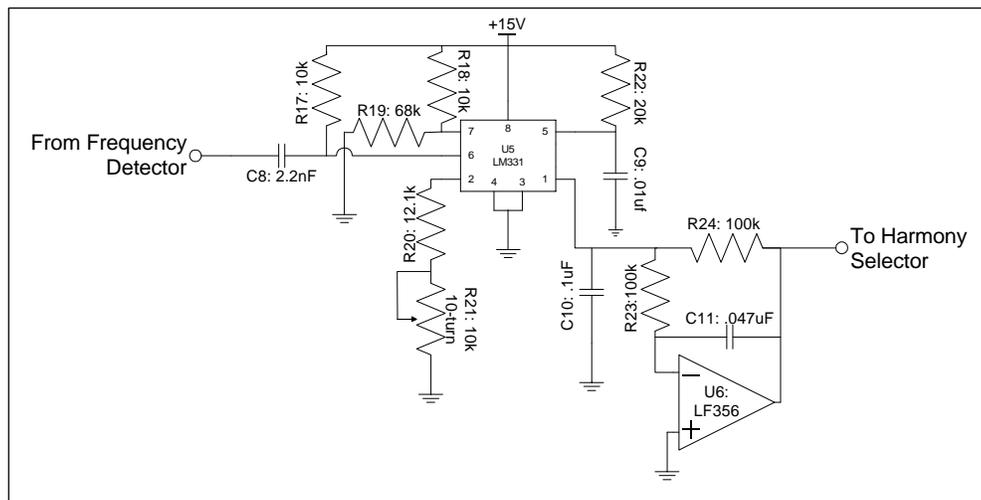
*Applications of Microprocessors* by Hal Chamberlin and works well under a wide variety of input waveforms. However, as described in Chamberlin, certain waveforms such as those with two equal positive peaks will cause the circuit to mistrigger. These situations did not typically come up with vocal input signals.

#### 2.2.4 Frequency to Voltage Converter [Andrew]

The frequency-to-voltage converter uses a National Semiconductor LM331AN chip (U5) in a configuration taken from the datasheet. C8 and R17 differentiate the input signal, to produce positive and negative pulses on the positive and negative edges of a square wave (the LM331 triggers on the positive edge only, making the duty cycle of the input wave irrelevant). R18 and R19 provide an external reference for an internal current source. According to the data sheet, the output voltage

$$V_{out} = -f_{in} * 2.09V * R24 / (R20 + R21) * R22 * C9.$$

The values of R20, R21, and R24 were left as in the standard configuration, but C9 and R22 were chosen to provide around 3kHz full scale, which is high enough for the highest harmony of the highest input pitch. Allowing a higher full-scale output magnifies



**Figure 10: Frequency to Voltage Converter Schematic**

the effects of offsets at lower frequencies, which will need to have correspondingly lower output voltages. All resistors were 1% metal film, not because of their tolerance but for their temperature coefficient, which is  $^{\circ}\text{C}$  versus up to 2500PPM/ $^{\circ}\text{C}$  for carbon film. A high temperature coefficient will cause the circuit to drift out of tune as it warms up or the ambient temperature changes. Similarly, C9 is a polystyrene capacitor, also chosen for its low temperature coefficient.

C10, C11, R23, R24, and U6 form a 2-pole low-pass filter at around 16Hz to reduce the ripple on the output voltage. An LF356 is used for U6 because of its low input bias current which reduces offset problems. The linearity of this circuit is 0.01% or better with good components. While such accuracy may seem unnecessary, slight variations in pitch are very perceptible compared to slight variations in amplitude (a chromatic step is only a 6% change in pitch).

### 2.2.5 Harmony Selectors [Andrew]

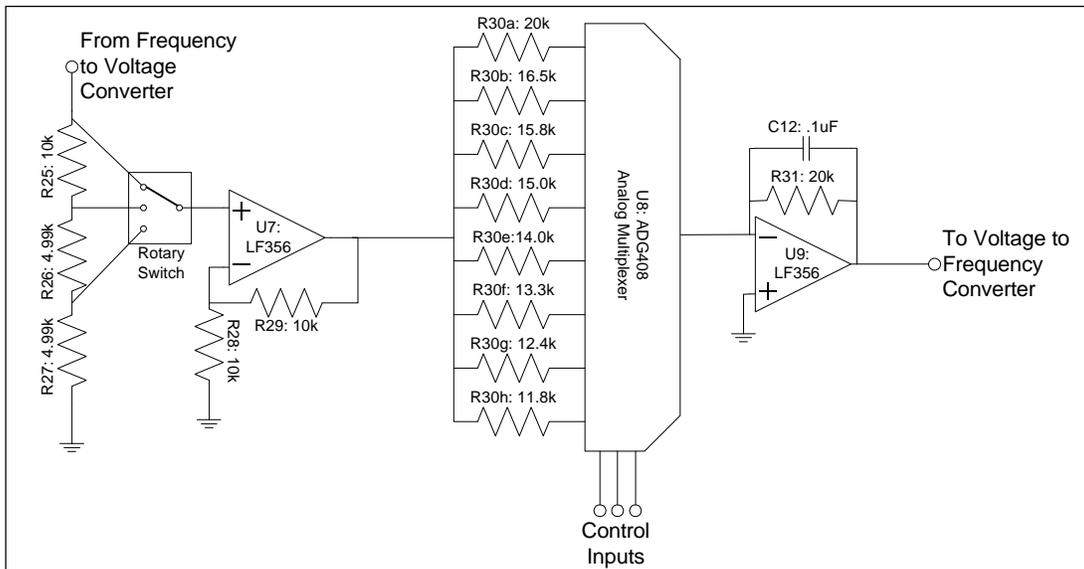
There are actually three identical harmony and synthesis blocks in our project, which allows independent control of three different voices. (Only one schematic is shown). Thus, one voice can be set to a unison, one to a third, and one to a fifth to form a complete triad, among other possibilities. The harmony selector is effectively a series of gain stages with selectable gain according to a switch on the front panel or other input signal. R25, R26, and R27 form a ladder network voltage divider, where the output voltage between R25 and R26 is half the input voltage, and the output between R26 and R27 is one fourth the input voltage. U7, also an LF356 for low offset, is configured as a noninverting amplifier with a gain of 2, so the output can be selected as twice, equal to, or half of the input voltage. This corresponds to an octave up, a unison, or an octave

down. U7 also isolates R25, R26, and R27 from the input impedance of the amp formed by U9, R30, and R31.

U9 is configured as an inverting amplifier with a feedback resistor R31 = 20k. However, the input resistor R30 is selected from 8 choices by U8, an Analog Devices ADG408BN analog multiplexer. U8 takes in three digital control signals and connects one of 8 analog inputs to its output. The gain of the system is thus  $-R31/R30$ , which can be selected according to the following table:

Input	R30	Gain	Interval
000	20k	1	Unison
001	16.5k	1.2	Minor Third
010	15.8k	1.25	Major Third
011	15.0k	1.33	Perfect Fourth
100	14.0k	1.41	Diminished Fifth
101	13.3k	1.50	Perfect Fifth
110	12.4k	1.60	Minor Sixth
111	11.8k	1.67	Major Sixth

The gain is the ideal from the definition of the interval; the actual values deviate slightly due to component tolerances and closest available values (all R30a-R30h are 1% metal



**Figure 11: Harmony Selector Schematic**

film). Actual gains are, of course, negative as U9 is inverting. This is in fact necessary as the V-F expects a positive voltage but the F-V outputs a negative one. (In the F-V case, the precision circuit is used. The normal version of the V-F actually worked better than the precision one). C12 adds another low-pass filter (at 80Hz this time) in an attempt to further reduce ripple.

The digital control signals come from the logic circuitry, which is driven from a keypad of nine buttons. The buttons are laid out in a grid according to type of chord and the relation of the input pitch to it. Chords can be major, minor, or diminished, and the note being sung could be either the root, the third, or the fifth of the chord. Thus there are nine different possible ways to generate the harmony to make sure it is within the correct chord structure-- a harmony generator that always generated a major third and perfect fifth to any pitch would be useless in any real-world situation.

The buttons are laid out with the rows denoting position of the note within the chord and the columns denoting the type of chord. The overall layout is:

	<b>Major</b>	<b>Minor</b>	<b>Diminished</b>
<b>Fifth</b>	A	B	C
<b>Third</b>	D	E	F
<b>Root</b>	G	H	I

So, for example, if the singer was singing C against a background harmony of C major, he would push button G which would generate a major third and perfect fifth above his voice. But if he was singing the same C against a prevailing harmony of A minor, he would push button E which would generate a first-inversion A minor chord of C, E, and A-- a major third and major sixth. Choosing the correct button for any note requires a basic knowledge of music theory, and to use it in any practical context would require a bit of planning to figure out ahead of time which button was the most appropriate for

each harmony. This device is in many ways an instrument of its own, and like any instrument, it requires a bit of practice to play properly.

The digital logic actually generates a 6-digit sequence for any button. The first three digits are the identity of the “third” interval, which could be a major third (as in buttons E and G), a minor third (as in buttons D, F, H, and I), a perfect fourth (in second-inversion chords like those of buttons A and B), or even a diminished fifth / augmented fourth (for button C). The identity of the fifth can be diminished (for button I), perfect (G, H), a minor sixth (B, D), or a major sixth (A, C, E, F). Three-pole, three-position rotary switches on the front panel take in all six inputs and produce three outputs, which are either the three digits corresponding to the third, the three digits corresponding to the fifth, or 000, which is a unison and never changes with the harmony. The three voices can be adjusted independently to unison, third, or fifth, although the buttons control the harmony of all three simultaneously.

#### 2.2.6 Harmony Keypad Control Logic [Philip]

This circuit utilizes three CMOS ICs containing digital logic gates in order to generate a different, 6-bit digital code for each of the 9 buttons. A register is used to store this code until the next button is pressed. If you follow the precarious logic gates and transistors that run along the bottom of the schematic, it can be seen that pressing any button will cause the voltage on the register's clock pin to rise.

Because each button causes a code to be generated while simultaneously signaling the register to trigger, a very small lowpass filter (R1 and C1) has been placed in front of the clock input. This filter slightly delays the voltage on the clock pin and ensure that all the logic propagation will be finished before the register is triggered.

The BJT transistors in this circuit are all being used logic gates. The voltage at their collectors should either be 15 volts or close to 0. Transistors were used for a few logic gates in order to save space and not waste ICs.

The six digit codes generated by the buttons are described by the following boolean equations:

$$\begin{aligned}
 D0 &= C \\
 D1 &= \neg(A+B+E+G) \\
 D2 &= \neg(C + E + G) \\
 D3 &= 1 \\
 D4 &= \neg(G + H + I) \\
 D5 &= \neg(B + D + I) \\
 CLK &= (A + B + C + D + E + F + G + H + I)
 \end{aligned}$$

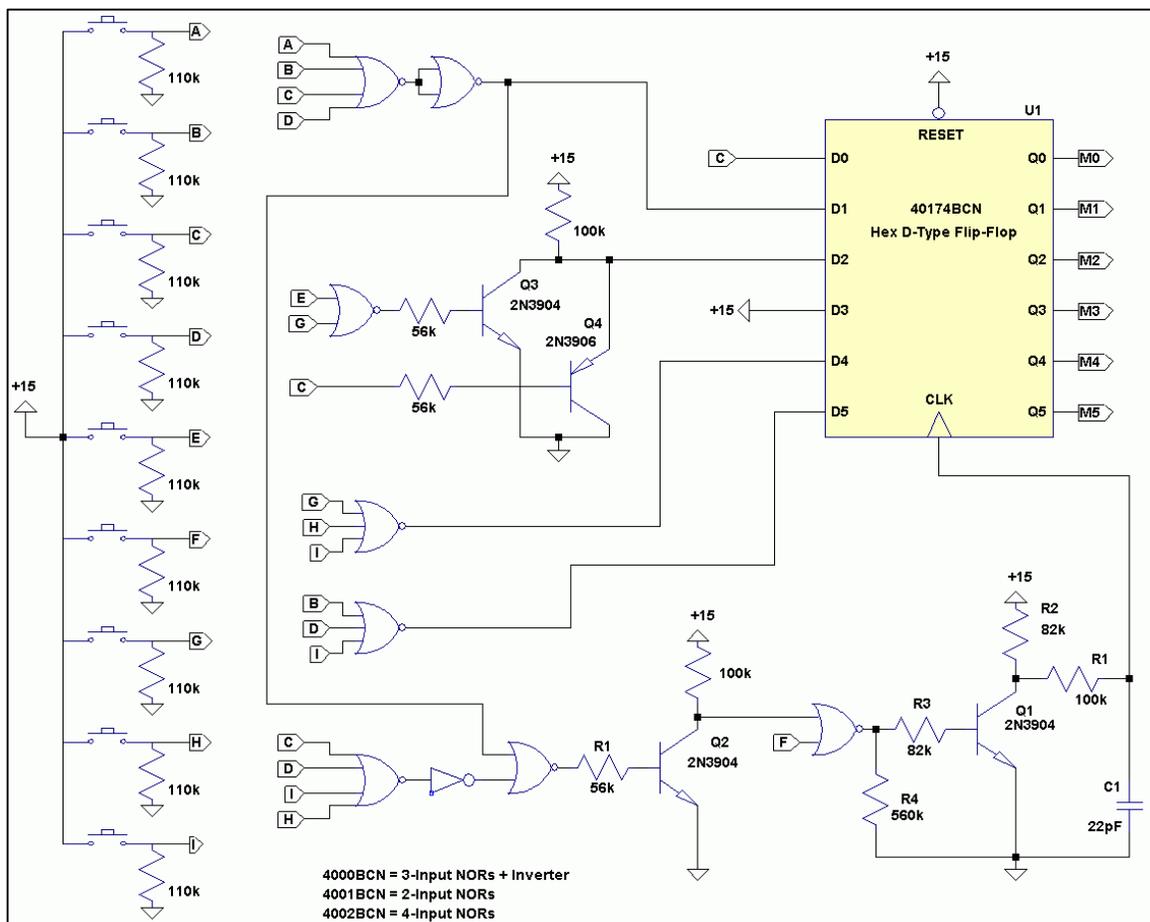
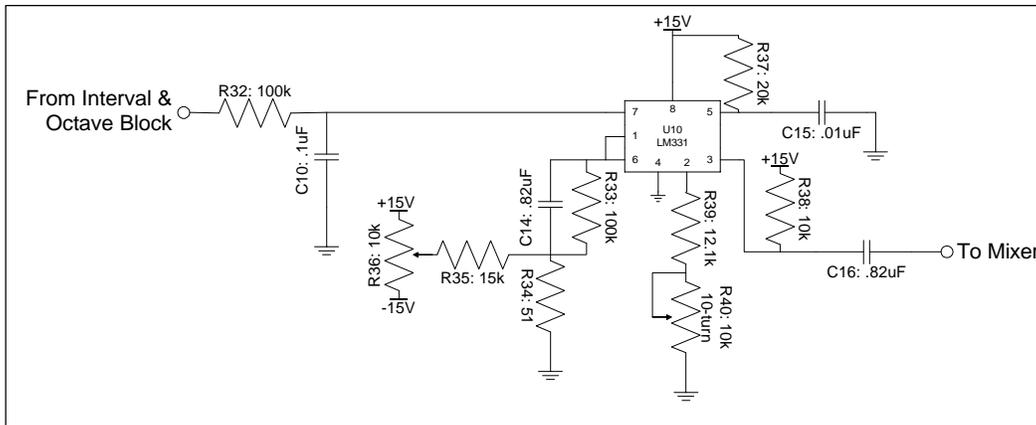


Figure 12: Keypad Logic Schematic



**Figure 13: Voltage to Frequency Converter Schematic**

### 2.2.7 Voltage to Frequency Converter [Andrew]

U10, like U5, is an LM331 chip. This flexible part can be used as either an F-V or V-F converter. The V-F configuration is also taken from the datasheet, but R37 and C15 are changed to match the values in the F-V converter. A 10k trim was selected for R40 (and R21) instead of 5k, because of availability. R35 was reduced from 22k in the datasheet to 15k here to allow a greater range of offset adjustment. Like the F-V converter, resistors are 1% metal film, and C15 is polystyrene.

The output of the V-F is given by

$$f_{\text{out}} = V_{\text{in}} / 2.09V * (R39 + R40) / R33 * 1 / (R37 * C15).$$

R40 was calibrated such that  $f_{\text{out}} = f_{\text{in}}$  for unity gain on the harmony selector. This made sure the unison was precisely in tune regardless of component tolerances, although other intervals might vary. The offset adjustment is required so that the output frequency is a precise multiple of the input voltage, without any constant added. That is, 0V in should produce 0Hz out.

The output of the V-F is a pulse wave with a very low duty cycle. C16 removes any DC component from the output to avoid riling the mixer.

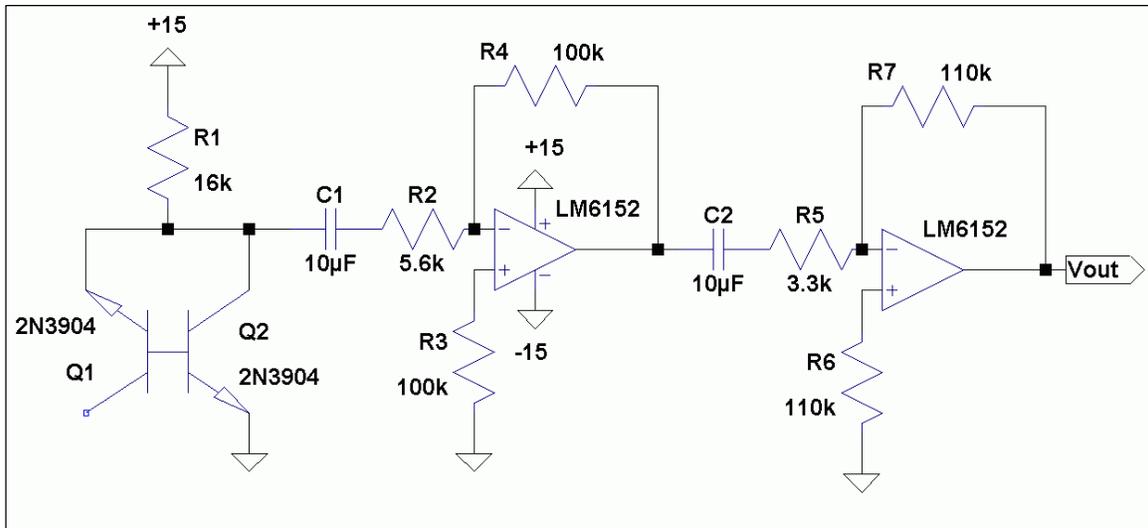


Figure 14: White Noise Generator Schematic

### 2.2.8 White Noise Generator [Philip]

The circuit shown in Figure 14 generates noise by utilizing the "avalanche noise" created by applying a strong reverse bias on PN junction. This is often done with a zener diode, but in this circuit is accomplished with the collector-to-base junction of Q1. In the circuit below, 14.4 volts is applied across Q1, due to the 0.6 volt drop across the base and emitter of Q2. This exceeds the maximum breakdown voltage of 6 volts, as listed in the 2N3904 data sheet, and an extremely noisy current is forced through Q1.

The current going through Q1 also goes through the base of Q2, causing a much larger current to go through the collector of Q2. This amplifies the current going through R1 and thus amplifies the voltage noise at the collector of Q2. If Q2 is removed and the base of Q1 is grounded, the amplitude of the noise drops by a factor of five.

Even with Q2, the voltage of the noise is only around 25mVpp. The two op-amps in series amplify the noise by a factor of 595 to give an output of 15Vpp. Since the upper range of audio is 20kHz, an op-amp with a gain bandwidth product of 11.9 MHz ( $595 * 20\text{kHz}$ )

20000) is needed. At 75 MHz, a single LM6152 would have been adequate. However, since they come in pairs, the other one was used in the design.

This design was modified from a "Hardware Random Bit Generator" found at <http://willware.net:8080/hw-rng.html>.

### 2.2.9 Mixer [Andrew]

The mixer is a straightforward inverting amplifier formed by U11b. The gain of each of the harmony inputs is -1 ( $R_{49}/R_{41}$ ,  $R_{49}/R_{42}$ ,  $R_{49}/R_{43}$ ). These signals already have up to 15V of swing, and do not need any additional gain. The white noise output was somewhat lower, so its gain is set to  $-R_{49}/R_{44} = -4.5$ . R50-R53 allow independent level adjustment of each of these channels. U11a provides extra gain to an external input signal, whose level is adjusted by R45. U11a is a noninverting amplifier with a gain of  $1+R_{47}/R_{48} = 11$ , and setting  $R_{46} = 10k$  gives this signal another gain of 10 in the mixer. The reason for this high gain is that the external signal may be something with low output like a microphone or instrument.

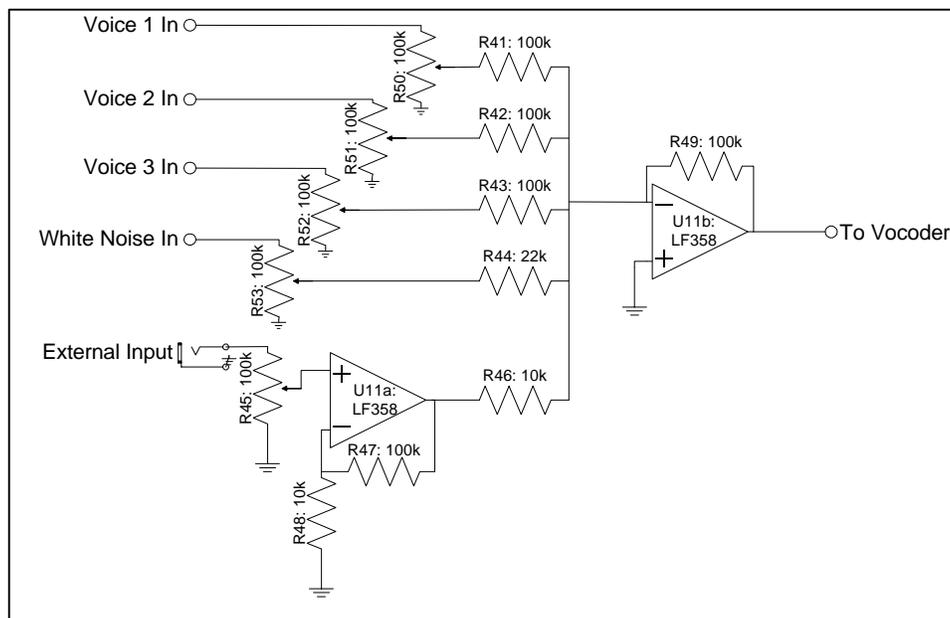


Figure 15: Mixer Schematic

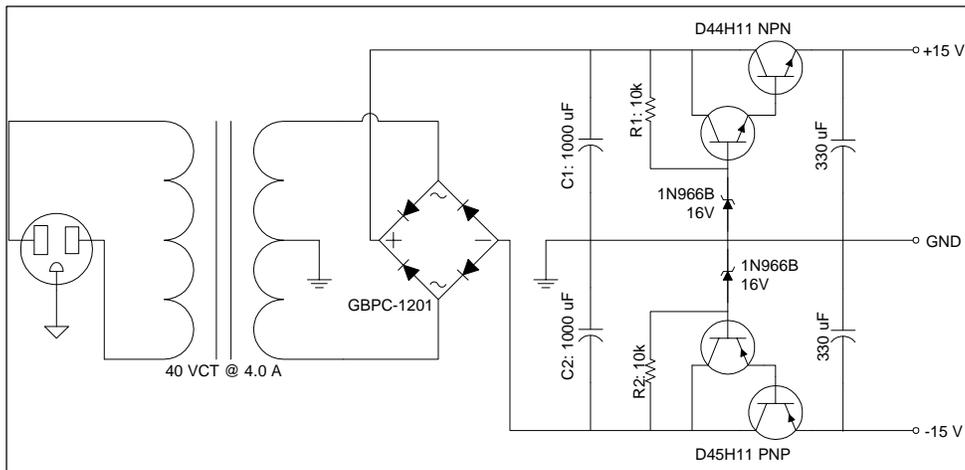
The output of the mixer is sent to the vocoder to be modulated by the voice signal. This allows the harmonic-rich pulse waves of the harmony to be shaped into waves that sound more like the human voice, to create the illusion of multiple voices singing at once.

### 2.3 Power Supply [Uriel]

Based on the number of components in our project, and their listed power requirements, we assumed that our project would require around 1 Amp of current. Therefore, the power supply was designed with current draw in mind. The basic design of the circuit is a series-pass, emitter-follower DC regulated power supply, using a 16V 1N966B zener diode as the regulator.

Our transformer was a 40 VAC center-tapped transformer capable of sustaining 4A. This was rectified by a GBPC-1201, 12A bridge rectifier into two full-wave signals at +/- 20 V<sub>RMS</sub>, which charged up C1 and C2, creating approximately +/- 26 V DC.

A small amount of current from the capacitors crosses R1 and R2 to supply the transistor base and zener diode. We used a Darlington pair to reduce the required base current; with a single transistor on each rail, there was a danger of the resistor exploding



**Figure 16: Power Supply Schematic**

due to the amount of power dissipating through it. The Darlington pair allowed us to use a larger, lower wattage resistor.

Each transistor was mounted onto a 2x2x1” heatsink. Under test loads, the power supply was able to sustain 1 Amp for several minutes.

### 3. Measurements

#### 3.1 Filter Bank [Philip]

The following set of measurements was taken from the filter bank:

Calculated			Bank 1			Bank 2		
Center Freq	Q	Gain@f <sub>0</sub>	Center Freq	Q	Gain@f <sub>0</sub>	Center Freq	Q	Gain@f <sub>0</sub>
111	3.17	0.99	111	2.85	0.862	118	3.10	0.902
177	2.16	0.98	189	2.17	0.761	188	2.14	0.795
250	8.46	1.02	260	8.12	0.772	261	8.42	0.789
297	8.46	1.02	309	7.92	0.738	310	8.16	0.786
354	8.46	1.02	366	8.32	0.724	366	8.13	0.761
420	8.46	1.02	438	7.96	0.724	444	7.93	0.733
500	8.46	1.02	494	8.23	0.888	498	8.59	0.950
595	8.46	1.02	590	9.08	0.959	588	8.28	0.837
707	8.46	1.02	709	8.65	0.880	709	8.65	0.871
841	8.46	1.02	837	8.21	0.817	834	7.38	0.701
1000	8.46	1.02	1008	8.76	0.899	1004	7.61	0.755
1189	8.46	1.02	1184	8.00	0.798	1173	8.38	0.857
1414	8.46	1.02	1409	8.44	0.857	1414	9.30	0.959
1682	8.46	1.02	1676	8.06	0.803	1668	8.42	0.868
2000	8.46	1.02	2006	7.71	0.733	2013	8.42	0.829
2378	8.46	1.02	2402	8.52	0.843	2395	9.74	0.871
2828	8.46	1.02	2870	8.20	0.846	2865	8.23	0.814
3364	8.46	1.02	3350	8.11	0.814	3385	8.06	0.783
5187	1.58	0.99	5324	1.67	0.960	5249	1.55	0.836

#### 3.2 Harmonizer [Andrew]

Much of the performance of the harmonizer is subjective, but some measurements can be taken. The exact gain of the microphone preamp or the exact -3db point of the filters are generally unimportant to the performance of the circuit, but the linearity and

calibration of the frequency-to-voltage and voltage-to-frequency converters are critical, as are the accuracy of the gains in the harmony selector. Measurements of the performance of the F-V converter are given in the following table:

Frequency In (Hz)	Voltage out (V DC)	Ripple (mV AC)
60	-.147	30.6
70	-.171	26.3
80	-.197	23.0
90	-.221	20.4
100	-.246	18.4
120	-.296	15.3
140	-.345	13.1
160	-.395	11.5
180	-.445	10.2
200	-.494	9.17
230	-.569	7.96
250	-.618	7.31
280	-.692	6.51
300	-.742	6.07
350	-.866	5.18
400	-.990	4.50
450	-1.11	3.98
500	-1.24	3.57
600	-1.49	2.95
700	-1.73	2.51
800	-1.98	2.17
900	-2.23	1.90
1000	-2.48	1.68
1200	-2.97	1.35
1400	-3.47	1.10
1600	-3.96	0.89
1800	-4.46	0.73
2000	-4.96	0.61
2500	-6.20	0.32
3000	-7.44	0.06

The relationship between  $V_{out}$  and  $F_{in}$  correlates to  $V = -2.48 \times 10^{-3} * f + 2.28 \times 10^{-3}$ .

This means that 1kHz = -2.48V, and that there is about 2.3mV of positive offset in the output voltage, which shows up in the measurement of 100Hz = -.246V. The correlation of this data is -.9999994, indicating very high linearity. Overall the performance of this circuit is excellent. Some undesirable effects will be discussed in the error analysis.

The voltage-to-frequency converters are of similarly high performance, although the results indicate some inaccuracies in calibration:

Vin (V DC)	Fout 1 (Hz)	Fout 2 (Hz)	Fout 3 (Hz)
.1	38.4	36.7	40.3
.2	78.9	77.6	80.8
.3	119.2	118.2	121.2
.4	159.6	159.3	161.6
.5	200.0	200.1	202.0
.6	240.3	204.9	242.3
.7	280.5	281.6	282.6
.8	320.9	322.4	323.0
.9	361.2	363.2	363.3
1.0	401.5	403.9	403.7
1.2	482.1	485.5	484.4
1.5	602.7	607.5	605.1
1.7	683.3	689.0	685.8
2.0	804.2	811.4	806.8
2.5	1005.8	1015.2	1008.5
3.0	1207.5	1219.3	1210.5
4.0	1611.0	1627.6	1614.3
5.0	2014.7	2035.8	2018.2
6.0	2417.2	2443.0	2421.1
7.0	2820.6	2850.9	2824.7

The best fit lines for each are:

$$\begin{aligned}
 \text{Fout1} &= 403.2\text{V} - 1.80 & \text{Correlation} &= .99999996 \\
 \text{Fout2} &= 407.8\text{V} - 3.99 & \text{Correlation} &= .99999996 \\
 \text{Fout3} &= 403.5\text{V} + 0.67 & \text{Correlation} &= .99999997
 \end{aligned}$$

Taking the inverse of the F-V converter, the ideal V-F converter would have

$$F_{\text{out,ideal}} = (\text{V} - 2.28 \cdot 10^{-3}) / 2.48 \cdot 10^{-3} = 403.2\text{V} - 0.92$$

However, it is important to note that due to offsets elsewhere in the system, and variations in the resistor values of the harmony selector, what is correctly calibrated when the F-V and V-F and hooked directly together may not be correct when placed in the full circuit. As such, it was important to measure the performance of each voice when set to

unison relative to the input frequency. At all frequencies, the output should be the same as the input but the results show some variation:

<b>Fin (Hz)</b>	<b>Fout1 (Hz)</b>	<b>Fout2 (Hz)</b>	<b>Fout3 (Hz)</b>
60	60.0	60.0	60.2
80	80.2	80.0	81.5
100	101	100.0	102.0
200	201.5	201.0	203.1
300	301.5	301.9	303.2
400	401.4	402.7	403.3
500	501.2	503.5	503.4
700	700.9	705.1	703.5
1000	1000.3	1007.5	1003.7
1500	1499.7	1511.6	1504.1
2000	1999.1	2015.8	2004.7
2500	2498.3	2519.7	2505.0
3000	2997.7	3023.9	3005.3

The results reveal that the first voice is calibrated very well, the third fairly well, and the second is in need of a much better calibration. Calibration is a difficult procedure because it involves adjusting two interrelated variables: gain and offset. Listening at a particular frequency, changing either the gain or the offset will change the output frequency, so it is difficult to hear when both are right without repeatedly changing frequencies and readjusting until the result approaches the ideal. From these results, it appears that a smaller gain and higher offset might bring the second voice into better calibration.

Lastly, the accuracy of the intervals for each voice was measured. These largely depended on the exact values of the resistors in the R30 bank. An A 440 was used as the test frequency, and the results are here compared to the “correct” frequencies:

Interval	Ideal Freq (Hz)	Fout1 (Hz)	Fout2 (Hz)	Fout3 (Hz)
Unison	440.0	441.4	443.0	443.8
Minor Third	528.0	534.6	538.1	537.7
Major Third	550.0	558.6	562.2	562.6
Perfect Fourth	586.7	585.3	592.5	591.8
Diminished Fifth	622.3	628.0	630.2	631.5
Perfect Fifth	660.0	664.6	669.4	666.1
Minor Sixth	704.0	711.5	716.9	710.6
Major Sixth	33.3	747.0	755.0	750.7
Up One Octave	880.0	882.0	888.8	884.5
Down One Octave	220.0	221.0	221.1	222.4

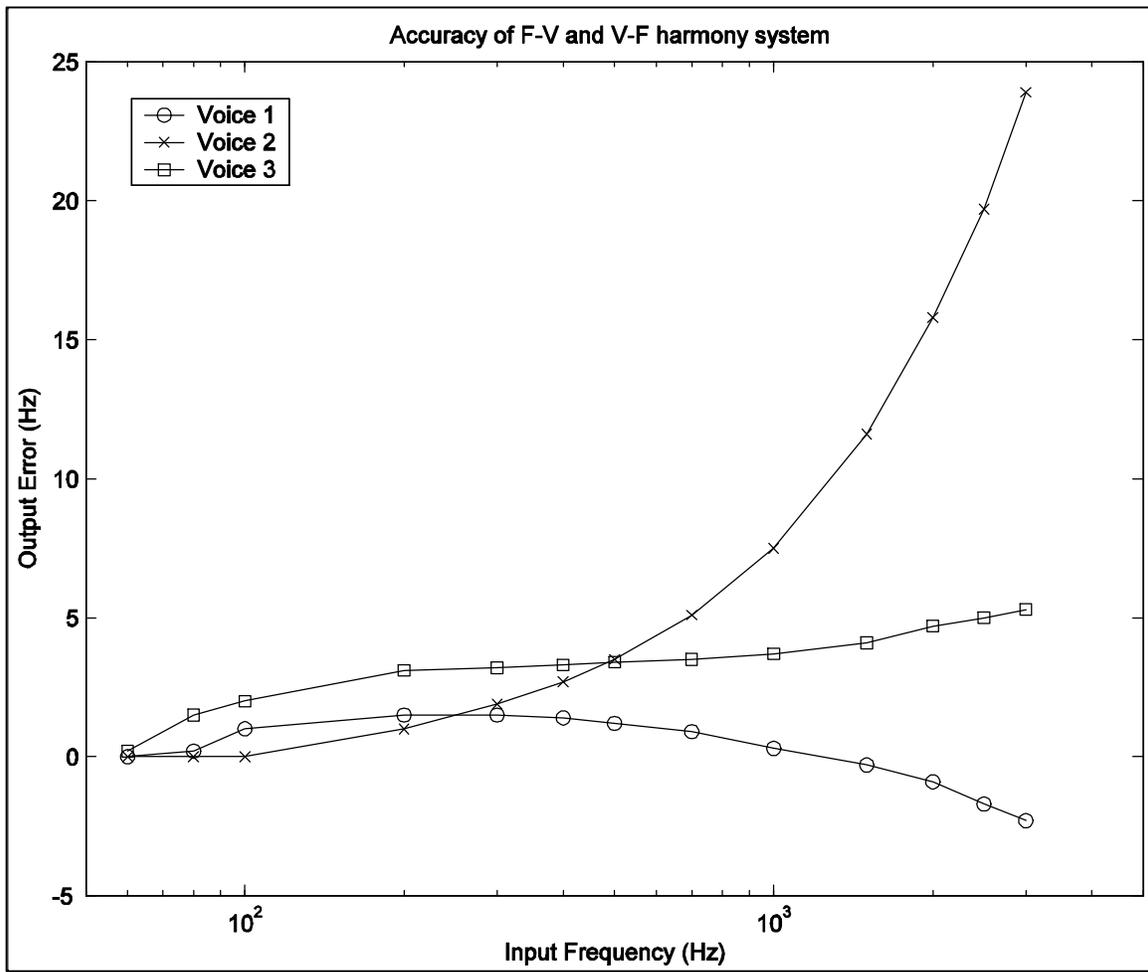


Figure 17

## 4. Errors and Corrections

### 4.1 Filter Banks [Philip]

The filters were designed at the beginning of the project and a lot of time was spent constructing them and testing them. Their outputs could have been attenuated to be more even, but other than that, they did not cause any problems.

### 4.2 Envelope Follower / VCA [Uriel]

The original VCA design we used was incredibly noisy, and had distortions which resembled exceeding the slew rate of an op-amp. We later learned that the voltage divider at the input was particularly important: too much signal would rail the OTA, too little signal would be noisy (and therefore useless in our project). Much more control was obtained by redesigning the current source.

The other major problem was DC offset on the VCA outputs, which was narrowed down to an offset current at the OTA inputs. A potentiometer was installed on the non-inverting input of each CA3080 in order to adjust this noise down to zero. While the noise did not go down to zero, it was certainly lessened.

### 4.3 Harmonizer [Andrew]

Are the measured results musically significant? In a word, yes. While the performance of the converters are very linear, and the results are close numerically, the differences are still enough for the output to sound out of tune. A trained listener can hear pitch differences down to a difference of 1Hz, and by the time the output deviates 20Hz from ideal as in some of the outputs of V-F converter number 2, the difference should be apparent to anyone. Subjectively, some of the difference is masked by the

inherent vibrato on any singer's voice that ensures that the pitch is not constant to begin with.

In any case, it is probably impossible to get results accurate to less than 1Hz, but better calibration would greatly improve performance, particularly of the second voice. The results also show that 1% precision resistors are not exact enough for the interval generators. Partly, it is because the ideal values are not always available (i.e. there is no 12.0k resistor in the 1% series, so 11.8k was used). Partly it is because a 1% variation of pitch is still several Hz and a noticeable change. Given more time to adjust and calibrate, the correct thing to do would be to add small trim resistors in series (or large resistors in parallel) with each of the R30 interval resistors. These trims would be somewhere on the order of 20-200 $\Omega$  (or 1-10M) and would correct for the slight variations in each of the individual resistors.

Another equally significant problem comes from the ripple present on the output of the F-V converter. Despite going through three low-pass filter stages, it is still significant at lower frequencies. In fact, due to the nature of the low-pass filter, the ripple gets worse as the frequency drops. This is a significant problem since the DC voltage also drops as the frequency drops, meaning the ripple as a portion of the total output goes up quadratically with decreasing frequency. The way this manifests itself at the output is by a beating sound caused by the interaction between the ripple at one frequency and the output of the V-F at a slightly different frequency. With the octave switch set to drop the pitch one octave, this also led to a spurious locking-in of the output pitch at to the input at low frequency rather than being an octave below as it should have been. Overall, the

ripple problem indicates that this circuit works better with higher pitches, and thus the device is probably better used with female voices.

Extra low-pass filtering would eliminate some of the ripple, but the tradeoff would be an increased settling time for the DC voltage, which could not respond to sudden changes very quickly. Finding the middle ground between these two design goals is difficult, but in retrospect, we perhaps should have gone more towards the reduced ripple side. This could easily be accomplished by increasing C10, C11, and C12.

Finally, the last remaining problem was in the frequency detection. Given the enormous variety of possible input waveforms, and the presence of ambient noise in the signal, the frequency detector would occasionally miss a cycle and fail to output the pulse it was supposed to. This led to a scratching, popping sound in the output of the voices that could be distracting. It might be worth experimenting with different frequency detection schemes to get a cleaner output.

## 5. Conclusions [All]

The filters worked fairly well and the ringing from all the resonance was not particularly noticeable. We had originally planned to build add a high pass and low pass filter to cover the rest of the spectrum, but did not quite have the time to design and build the four extra filters. Building 38 filters took a ridiculous amount of time and cost a lot of money to order the necessary components. The filter banks may have been overkill for the application.

The VCAs seem to be more trouble than they're worth. In the future, we would try to use newer OTAs such as the LM13700, which has a number of features to increase linearity, and hopefully decrease noise. Otherwise, a better system needs to be devised to

remove the DC current offset. It might be easier and better just to replace the VCAs altogether with analog multipliers, as that is their function anyway.

Overall, the harmony section worked fairly well also, but there were a number of things that could have been done to improve its accuracy. Most of these just required a large amount of time devoted to very fine tweaking, which did not seem like time well-spent while we were still finalizing aspects of the design. A redesign of the frequency detector might result in better performance, although the design used here is already the third generation, and performs better than solutions based on heavy low-pass filtering, auto-level correction, and Schmitt triggers.

The F-V and V-F converters, though, worked remarkably well, and the scheme for choosing intervals in a harmonically useful fashion was a success. The intervals can easily be tuned with trim resistors, but the basic concept was sound. The microphone preamp performed without a problem, as did the filters on the signal. The mixer also worked as expected. By and large, the block was a success, with the main required changes being minor tweaks of calibration.

The only other problem was that the output waveforms, while intelligible, did not sound much like the singer's voice even after being put through the vocoder. Some of this might be related to the performance of the vocoder, but it may also be that the output waves were too harmonically rich. A low-pass filter at the output of each voice would take care of that problem.

The project was certainly very impressive, but temperamental. Some more work on this project might easily turn it from a lab to a sellable product. Unfortunately, the

market is probably saturated at this point by digital electronics that perform the same function at a fraction of the cost.

## 6. Acknowledgements

We would like to thank Joe Sousa for helping us time and time again with the bugs that plagued our design; Ron Roscoe for his supply of reference books, parts, and know-how; all the staff of the EECS lab desk, for giving us way too many breadboards and CA3080s; and finally, Rahul Agrawal and Cang Kim Truong for standing around telling us how cool our project is.

## 7. References

Chamberlin, Hal. *Musical Applications of Microprocessors*. Hayden Books. 1985.

Lancaster, Don. *Active Filter Cookbook*. Synergetics Press, USA. 1975

Lenk, John D. *Simplified Design of Filter Circuits*. Butterworth-Heinemann. 1999

*Hardware Random Bit Generator*. <<http://willware.net:8080/hw-rng.html>>. 10 Dec 2003. (Appended)

## 8. Appendices

Included in this section:

Datasheet: LM331 Precision Voltage-to-Frequency Converter

Datasheet: ADG408/ADG409 High Performance Analog Multiplexers

Website: *Hardware Random Bit Generator*  
from <http://willware.net:8080/hw-rng.html>