Behavior Research Methods, Instruments, & Computers
1989, 21 (6), 623-626

Frequency synthesis with the Commodore Amiga for research on perception and memory of pitch

ANNABEL J. COHEN and MAREK MIESZKOWSKI
Dalhousie University, Halifax, Nova Scotia, Canada

The potential of the Commodore Amiga as a digital synthesizer for research and demonstration in psychoacoustics and memory is discussed. Economy, ease of use, flexibility, portability, and accuracy outweigh disadvantages of narrow bandwidth, narrow dynamic range, and storage limitations for many applications encountered in pilot research and education. The Amiga also bears serious consideration for psychoacoustic studies requiring frequencies below 4000 Hz and modest signal-to-noise ratio, as exemplified by an implementation for research in absolute judgment, similarity scaling, and sequential pattern tracking.

Over much of the audible spectrum, human sensitivity is impressive. For example, under ideal conditions, many listeners can discriminate a frequency of 1000 Hz from one of 1002 Hz (Green, 1976). In order to study this sensitivity and to use acoustic stimuli in other experiments in perception and memory, the signal must be specified with great accuracy. The present article outlines the potential of the Commodore Amiga as an audio synthesizer for psychoacoustic and auditory research and demonstration.

The Amiga Computer and Synthesizer

The Amiga 1000 microcomputer was first marketed in the fall of 1985. In contrast to other 68000-based microcomputers developed at about the same time, such as the MacIntosh and Atari, the Amiga boasted a four-channel 8-bit digital audio synthesizer with stereo output. The other machines typically had one synthesizer (D/A converter or analog signal) and monaural output. Commodore had a history of successful innovation in musical synthesis. Two years earlier, the company had provided the most sophisticated sound chip for a home computer, the SID (sound interface device) on the Commodore-64, which controlled three independent analog voices. With an additional voice and complete digital control of signals, the Amiga offered much more. For example, it provided a flexibility of waveform specification far greater than did analog function generators of comparable cost, which are typically limited to fixed waveforms (usually sine, square, triangle, and pulse).

This work was supported by a grant to A. J. Cohen from the Natural Sciences and Engineering Research Council of Canada. Programs were written by Ron Whitaker. Wayne Zelmer, Larry Henrickson, and Shawn Currie assisted with the technical measurements, some of which were made at the School of Human Communication Disorders at Dalhousie University. Inquiries concerning the availability of the program may be addressed to Annabel J. Cohen, Department of Psychology, Dalhousie University, Halifax, Nova Scotia B3H 4J1.

Characteristics of the audio signal. The Amiga hardware is designed to generate up to 28,867.29 samples/sec (or the sampling period is 1/28867.29 Hz = .000034641 sec = 34.641 μsec). The number is derived from the time to scan one video line; the direct memory access (DMA) can retrieve two audio samples during each horizontal video scan line. The Amiga's system clock period is .279365 μsec, (clock frequency is 3579545 Hz) and the minimum period value is derived by dividing 34.641 μsec per sample by .279365, the number of μsec per clock interval, resulting in a minimum of 124 timing intervals (clock periods) per sample (see Peck, Deyl, & Miner, 1985, chap. 5, pp. 5-9).

The maximum sampling rate of 28867 Hz allows frequencies up to 14433 Hz, assuming the theoretical Nyquist limit. However, to prevent aliasing distortion, the manufacturers provided a fixed filter set at 7000 Hz, (30 dB attenuation), but beginning to have an effect even at 4000 Hz, according to Amiga specifications (Peck et al., 1985, chap. 5, p. 26). The fixed filter reduces the effective frequency range of the synthesizer. More recently, the Amiga 500 was introduced as a lower cost version of the 1000, having identical audio facilities but with some relaxation of the filter. More specifically, the Amiga filter has two components; on the Amiga 500, one of these is connected to the power-on LED (see Finkel, 1987, schematic A500-3). By disconnecting the power-on LED through software, the filter attenuation for higher frequencies can be lowered by about 10 dB, as measured by sweeping a frequency source under both bypassed and original filter conditions. Because only one component of the filter can be bypassed in any case, the difference between the original and the altered filter is not dramatic.

Figure 1 represents a square wave of 600 Hz produced by the Amiga. By many standards, the results of this square wave test are quite respectable. Modest distortion (e.g., the overshooting) reflects the bandwidth limitation imposed by the filter. This limitation nevertheless per-

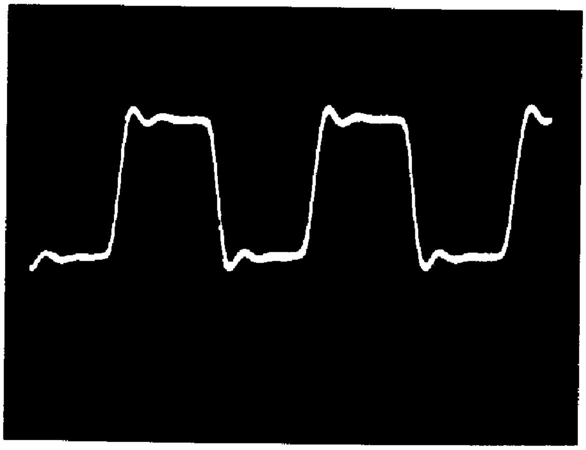

Figure 1. Oscilloscope representation of a square wave of 600 Hz.

mits examination of fundamental frequencies over the musical and much of the speech range. Figure 2 (panel a) shows a Fourier analysis (computed by Signal Technology Inc. Interactive Laboratory System ILS software) of a 250-msec 600-Hz complex tone consisting of the first six and the eighth harmonics. The harmonics (1, 2, 3, 4, 5, 6, and 8) in panel a are prominent as they should be, indicating the reliability of the method for the study of complex tone stimuli. (For comparison, a Fourier analysis of a 250-msec sine wave of the same fundamental produced by a voltage-controlled Tektronix TM503 oscillator is also shown in Figure 2, panel b).

The 8-bit D/A converter provides 256 amplitude steps, which determines a theoretical signal-to-noise (S/A) range of 48 dB, limiting studies to this dynamic range. More important however is the *actual* S/N ratio. Regression of the 256 output voltages on 256 DAC levels from −128 to 127 indicated modest nonlinear distortion for both output channels and an actual S/N level of around 7 bits (S/N = 42 dB). The Fourier analysis of an Amiga sine wave by the ILS program indicated a third harmonic 32 dB down from the fundamental, as can be seen in Figure 2 (panel c). This compares reasonably well with the analogue signal of panel b; nevertheless, the Amiga would not be adequate for studies of pure tones at high intensities, without other modifications (e.g., adding an analog filter over the range of the third harmonic for small ranges of stimuli; or, for a more general solution, adding digital dither; see Vanderkooy & Lipshitz, 1984).

It is possible to use the Amiga as a synthesizer within these constraints. Application depends upon a number of other factors discussed below.

Signal-generation algorithm. The present implementation for psychoacoustic research is written in the C programming language. Our algorithm maximizes accuracy at the expense of computer overhead by numeri-

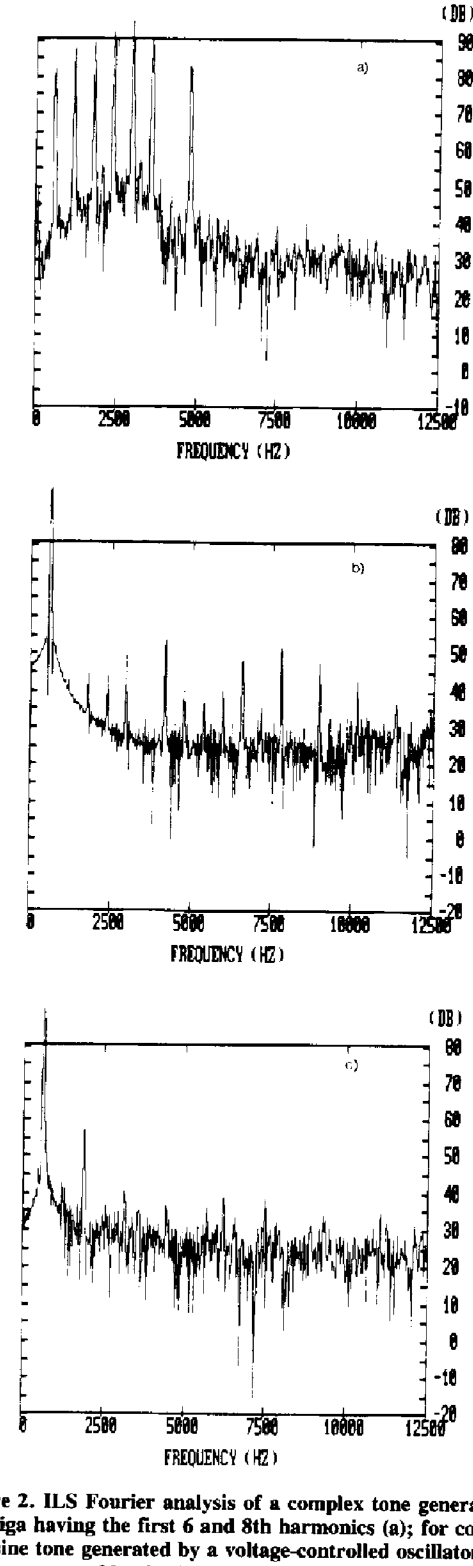

Figure 2. ILS Fourier analysis of a complex tone generated by the Amiga having the first 6 and 8th harmonics (a); for comparison, a sine tone generated by a voltage-controlled oscillator (b); a sine tone generated by the Amiga (c). All fundamentals are 600 Hz and of 250-msec duration.

cally specifying the entire wave. The chosen number of system clock points/sample is 171, and the corresponding sampling frequency is 20933.0117, allowing for the generation of any frequency within the bandlimited range as previously described. The duration of a sample is 47.771 μsec. A 500-msec note will take 10,466 bytes of computer memory. One megabyte of memory permits a library of 95 notes (994,270 bytes) in the experiment. The number of bytes per tone is limited, since less than 400,000 bytes are accessible by the sound-generating chip at any one time. With the above sampling rate, generation of a tone of up to approximately 20-sec duration is possible. Thus, within certain restrictions, for any given experiment, a library of tones that vary widely in number, waveform, and duration can be created.

Many psychoacoustic researchers use noise stimuli. Although not needed for our current applications, high quality noise can also be produced by our algorithm, which creates each sample of the wave independently, again subject to the memory-access limitation mentioned above. As Eggermont and Smith (1984) demonstrated with 32K samples (an order of magnitude less than that available to the Amiga sound chip), noise generated in software by a uniform distribution of random numbers can easily exceed specifications produced by commercially available analog noise or pseudonoise generators often used in psychoacoustic research. As yet, however, we have not made such tests on noise generated by the Amiga.

In order to avoid onset and offset transients, it is necessary to provide a smooth envelope. To accomplish this, the wave is segmented into rise, sustain, and decay portions (see Figure 3). Rise and decay functions are computed by adjusting the amplitude by a linear function over the rise and decay durations.

Complex tones are specified through additive synthesis. The greater the number of components, the greater the time to initially generate the numerical specification of the wave. The initial time to generate a 250-msec sine tone is approximately 20 sec, but a complex tone may take 4 min. Once generated, the numbers may be stored on disk and quickly read in, prior to the experiment.

Combination of channels. The Amiga permits the combination of four channels, allowing for modulation and for binaural stereo control. Combined channels indicated good coherence as shown by the Lissajous figure (Figure 4) for a 1000-Hz wave. Phase also can be controlled.

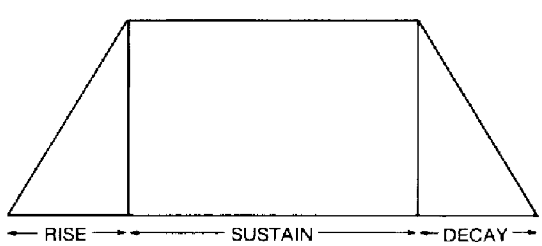

Figure 3. Representation of the window function, showing independent portions for rise, sustain, and decay (x-axis time, y-axis, amplitude).

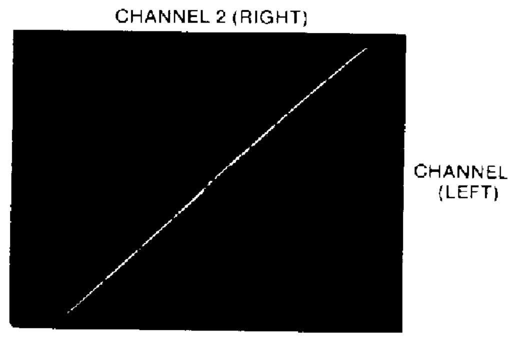

Figure 4. Oscilloscope Lissajous Figure for two 1000-Hz sine waves, each generated by a separate Amiga channel. Perfect phase coherence is indicated by a straight line.

Applications

Programs have been written for studies of absolute judgment, discrimination, and similarity scaling of tones and of melodies. Software has also been developed for studies of sequential tracking of tone sequences.[1] Responses and reaction times are recorded on-line.

In a typical study of absolute judgment (Cohen, 1988), the experimenter first defines the tones to be used. In the parameter file, one line per tone specifies fundamental frequency, relative intensity of each of the first 10 harmonics, sustain and rise/decay duration, overall intensity, and output channel. Additional parameters are also set, such as number of sessions and type of feedback. The algorithm produces all stimuli initially for storage on disk. These are loaded quickly into the program before the experiment is conducted. The data (response and reaction time) are collected on-line from the keyboard for later analysis.

For discrimination studies, pairs of stimuli are presented on each trial: One pair is identical, the other differs, and the subject must choose which of the pairs is different (e.g., see Cohen, 1989). In the fixed condition, the same discriminanda are examined over a block of trials; in the roving condition, different discriminanda appear on each trial. A pitch-discrimination paradigm in which tones are interpolated between the standard and comparison also has been successfully executed (Frankland & Cohen, 1989). For other research on similarity scaling, pairs of stimuli are presented, and the subject responds on a numeric scale.

In a paradigm for examining serial order recall, the same method for generating tones is used; up to nine different sets of nine tones can be examined in the same study. Up to nine different serial orders can also be specified; for example, order 12345678 specifies an ascending order of the eight frequencies, whereas 12348765 specifies an ascending (1234) descending (8765) pattern. In these studies, the subject listens to a sequence of tones in the scale and is required to type the order in which he or she believes the tones were presented. Again, a variety of feedback conditions may be specified. Also, in one condition, the sequence may be recycled until the subject is sure of the serial ordering. In this case, the number of stimulus repetitions provides an additional dependent measure.

For our implementation, high-quality headphones are connected directly to the Amiga audio outputs. For free-field studies, the signal can be fed to an amplifier and speaker system. The quality of the resultant signal depends on the quality of these transducers. The Amiga's internal speaker is adequate for providing feedback during programming and experiment development, but it is not recommended for serious psychoacoustic studies.

Advantages

Within the frequency range of 30–4000 Hz and for a 42-dB dynamic range, the accuracy of the signal is excellent for many purposes. Since waveform is entirely arbitrary, preliminary research on frequency, intensity, waveform, and phase relations can be conducted confidently on this machine. Once results are obtained, and future research directions are clear, the right choice of more sophisticated equipment can be wisely made.

The Amiga also serves as a mobile psychoacoustic station, transportable to different subject populations rather than confined to one laboratory setting. For elderly or very young subjects, this can be a great asset—at least in exploratory work in which paradigms suitable for college students must be modified for application to special populations. As is the case with any common computer, software developed by one researcher can be shared by other Amiga users without the purchase of additional hardware. Researchers using the same sound generator may more easily compare experimental conditions. For educational purposes, the Amiga permits classroom demonstration and allows for individual experimental stations at relatively low cost.

Conclusion

Over a respectable frequency and intensity range, Amiga computers provide an accurate signal using algorithms written in C. Most of the present limitations described are built into the machine. As it stands, therefore, the Amiga is not an ideal device for all serious psychoacoustic applications. However, in the initial stages of certain projects, it likely cannot be beaten for cost and flexibility. It is practical for pilot experiments, the results of which may determine the need for more sophisticated instrumentation. With its additional graphics (see Anstis, 1986; Tanner, Jolicoeur, Cowan, Booth, & Fishman, 1989) and data handling capabilities, it is also a good choice for a teaching laboratory in perception.

REFERENCES

Anstis, S. (1986). Visual stimuli on the Commodore Amiga: A tutorial. *Behavior Research Methods, Instruments, & Computers*, **18**, 535-541.

Cohen, A. J. (1988). Identification of a microtonal scale of complex tones. *Bulletin of the Psychonomic Society*, **26**(6, Abstract No. 191), 501.

Cohen, A. J. (1989). Discrimination of complex tones. *Journal of the Acoustical Society of America*, **85**(Suppl. 1), S143.

Eggermont, J. J., & Smith, G. M. (1984). White noise analysis of nonlinear systems with application to the auditory system. *Canadian Acoustics*, **15**, 3-16.

Finkel, S. (Ed.) (1987). *Commodore Amiga A500/A2000 technical reference manual*. Westchester, PA: Commodore Business Machines.

Frankland, B. W., & Cohen, A. J. (1989). Interference effects of tonality and contour on short-term memory for tones. In A. J. Cohen (Ed.), *Proceedings of the Annual Meeting of the Canadian Acoustical Association: Industry, University, Government; Halifax, 1989* (pp. 117-122). Halifax: Canadian Acoustical Association.

Green, D. M. (1976). *An introduction to hearing*. Hillsdale, NJ: Erlbaum.

Peck, R., Deyl, S., & Miner, J. (1985). *Amiga hardware manual*. Westchester, PA: Commodore Business Machines.

Tanner, P. P., Jolicoeur, P., Cowan, W. B., Booth, K., & Fishman, F. D. (1989). Antialiasing: A technique for smoothing jagged lines on a computer graphics image—an implementation on the Amiga. *Behavior Research Methods, Instruments, & Computers*, **21**, 59-66.

Vanderkooy, J., & Lipshitz, S. P. (1984). Resolution below the least significant bit in digital systems with dither. *Journal of the Audio Engineering Society*, **32**, 106-113.

NOTE

1. The Amiga has also been used for investigations of auditory/visual integration using BASIC programs for applications that do not require the acoustic accuracy afforded by the signal processing under discussion.

(Manuscript received January 27, 1989;
revision accepted for publication November 6, 1989.)