A Computer Model
for
Bar Percussion Instruments

Xavier Serra

Center for Computer Research in Music and Acoustics (CCRMA)
Department of Music, Stanford University
Stanford, California 94305

Abstract

This paper presents a computer model for bar percussion instruments that captures their main physical characteristics in a very flexible structure. The starting point is current knowledge of the acoustics of the instruments. Then a modular analysis/synthesis process is presented based on Discrete Fourier Transform techniques and filter methods. The result of the research is an algorithm implemented on the Systems Concepts Digital Synthesizer ("Samson Box") that integrates several synthesis strategies. This algorithm maps to the underlying physics of the instruments and its parameters are easily specified.

Introduction

A computer model of a musical instrument consists of an algorithm which can be implemented in a computer or an integrated digital circuit. It produces a digital signal that simulates the sound produced by the instrument, given as input a set of parameters specified by the user.

In this paper a model of bar percussion instruments is developed. We are aiming for a model that is general for the whole family of bar percussion instruments, and that means that it will have to be flexible enough to accomodate the different instruments. Another choice is that the model has to relate to the underlying physics, since it is on physical reality that our intuition is rooted. It also has to be computationally efficient, but without compromising its accuracy. Finally, the validity of the model will be based on the perceptual quality of the sound that it produces, and the analysis results or physical characteristics will be discarded whenever they are not perceptually relevant.

1. Acoustics of Bar Percussion Instruments

The discussion of this article concentrates on bar percussion instruments that use rectangular bars suspended at both ends. This family of instruments includes: marimbas, xylophones, vibraphones and glockenspiels (orchestral bells). The physical unit which is responsible for the sound production comprises three separate elements: bar, resonator (where applicable), and mallet.

1.1. Bar

A rectangular bar with free ends vibrates primarily transversely (by bending perpendiculary to its length) [7]. We will not consider less important vibrations such as longitudinal, torsional, or shear waves, since in the normal use of the mallet, striking the bar at right angles to its length, only the transversal modes are excited (Fig. 1).

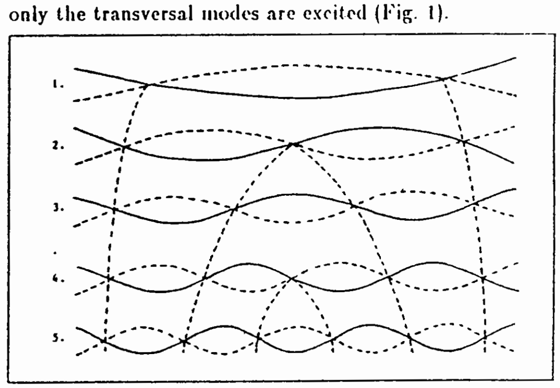

Figure 1. First five transversal modes of a bar with free ends.

The frequencies of the transverse modes of vibration for an ideal bar are described by the following formula [9]:

$$f_n = \frac{\pi v K}{8L^2} m^2$$

n =mode number $(1, 2, 3, \ldots)$, v =velocity of sound= $\sqrt{Y/\rho}$, Y =Young's modulus of elasticity, ρ = density, L = length, $m = 3.0112, 5, 7, \ldots, (2n+1)$, K = "radius of gyration" $= t/\sqrt{12}$, t = thickness.

As this formula shows, the frequencies of the transverse vibrations depend upon the length of the bar, upon the elasticity of the material from which it is made, and upon the thickness of the bar. The frequency ratios of the first few modes with respect to the first one (fundamental frequency) are: $1, 2.76, 5.70, 8.93$, etc., given any Y, ρ, L and t.

In the real world only some glockenspiels have bars with perfectly rectangular cross-section and will, in ideal conditions, behave according to the model. By contrast, the bars of marimbas, xylophones, vibraphones, and many glockenspiels have arch-shaped bottom surfaces for the purpose of tuning the upper partials and reducing the length required to get low frequencies. The shapes of these arches vary with manufacturers, and so do the frequencies of the different modes of vibration.

We cannot use the stated equation to study the vibrational behavior of the non-uniform bars, and to accomodate the model to all the deviations found in real bars is a very difficult problem. There are some experimental estimates for the first few modes in different instruments, but the results are very different depending on the author, especially for the high registers. In the low registers the results are more consistent and we can use the following frequency ratios from Moore [6] to compare the different instruments:

Marimba	Xylophone	Vibraphone	Glockenspiel
1 : 4 : 10	1 : 3 : 6	1 : 4 : 10	1 : 2.72 : 5.3 : 8.6 : 12.7

In the marimba, xylophone and vibraphone only the first three modes of vibration are really distinct; in the glockenspiel there are a few more. It is important to mention that we are only talking about the frequency of the different modes, not their amplitude nor their evolution in time. These variables depend on other factors apart from the shape of the bar, for example, the material from which it is made, the mallet used, and the way it is struck.

1.2. Resonator

Most of the bar percussion instruments have a resonator underneath each bar. They are cylindrical tubes closed at the lower end and tuned to the fundamental frequency of the bar.

The air inside a closed tube vibrates at odd integral multiples (i.e., $1, 3, 5, 7, \ldots$) of the fundamental frequency. This fundamental frequency can be expressed as a function of the length of the tube by the formula

$$f = \frac{c}{4L}$$

$$\frac{\lambda}{4} = L$$

f =fundamental frequency ($= \frac{c}{\lambda}$), λ = wavelength ($= \frac{c}{f}$), c = speed of sound, L = length of tube.

This means that in the ideal case the length of the tube is one quarter of the wavelength of the fundamental frequency of the bar. Such a tube will resonate the odd integral partials of the bar. Only the xylophone has an upper partial with this characteristic; in the other cases only the fundamental will be emphasized by the tube, plus some energy that happens to be in a resonated frequency region due to deviations from the ideal case.

1.3. Mallet

The mallet is used to excite the system; it is practically the only means that the performer has to control the sound. It consists of a heavy head fastened to a light shaft.

Apart from the way of playing, the hardness and size of the head will determine the sound quality produced by the instrument. A head of soft material and large contact area with the bar will produce a tone with a strong fundamental and less prominent partial tones. A head of hard material and a smaller contact area with the bar will produce a tone with more and stronger partial tones.

2. Computer Model

The objective is to create a general computer model for the family of bar percussion instruments. We first choose the characteristics of the model we want. (1) It has to be computationally efficient, but without compromising its accuracy. (2) Its parameters have to map to the underlying physics of the instruments. (3) The final judgment of the model will be based on the quality of the sound it produces, and so we will exclude the physical characteristics and analysis results whenever they are not perceptually important.

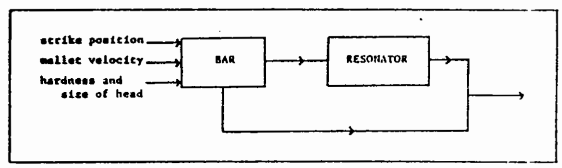

Figure 2. General model for bar percussion instruments.

The bar and the resonator are acoustically separate and will be modeled accordingly. The bar is a complex system, especially if we consider the deviations from the ideal case. For it, the model is based on sound analysis of instrumental bars, and it consists of a combination of sinusoids and a linear time-invariant filter driven by white noise. This model has proved very efficient and at the same time maps to the underlying physics of the bar. The resonator can be approximated by a simple Helmholtz resonator and modeled with a linear time-invariant filter similar to the one of the Karplus-Strong algorithm [8]—special case of the McIntyre-Woodhouse string model [5] . For the bar-resonator interaction the only control needed is the amount of the bar output going into the resonator. In the case of the mallet-bar interaction, there are three main variables to be specified: strike position, mallet velocity, and hardness and size of the mallet's head. Fig. 2 shows the basic structure of the model.

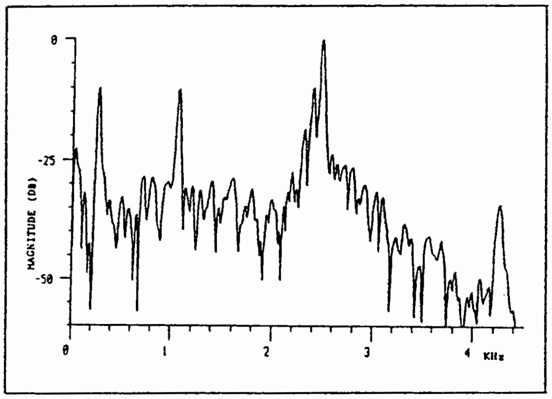

Figure 3. Magnitude spectrum of the attack portion of the middle C of a marimba bar.

2.1. Bar model

By looking at spectral analyses of the sounds produced by different bars (the resonator being removed), we can easily see the first few modes of vibration. Fig. 3 shows the spectrum of the attack portion of the middle C of a marimba bar. In it the first 3 modes are prominent, with frequencies of 264 Hz, 1056 Hz, and 2480 Hz. But there is more than that: there is a noisy component from 0 to 4500 Hz that is definitely perceptible and has to be taken into account. This is the situation at the attack of the note.

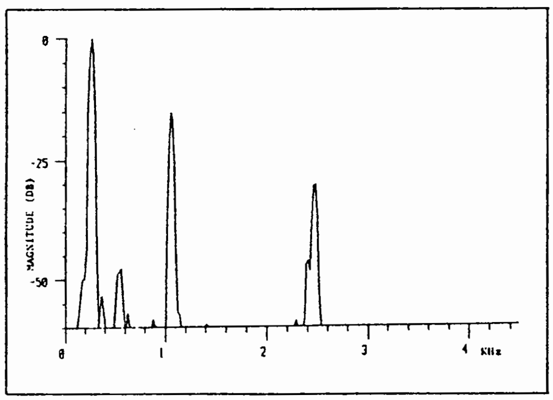

Figure 4. Magnitude spectrum of the same sound as Fig. 3 but taken 1/10 of a second later.

If we take a spectral analysis of the same sound but 1/10 of a second later (Fig. 4), we see that, practically, the only energy left corresponds to the three modes, the noise part has vanished.

We have observed that the bars of the different instruments behave in a similar fashion. They have a few prominent modes, very stable in frequency for the whole duration of their sound, plus a noisy component at the attack of the note. We aim at finding a technique to analyze and synthesize this kind of behavior and whose results can be manipulated to simulate different bars and mallet excitations. A possible choice is additive synthesis, using the phase vocoder [2] as the analysis-synthesis tool. Due to its identity properties, we are assured of a good resynthesis. However, this technique proved to be very expensive, and not flexible enough for our purposes. Neither the noise at the attack time nor the inharmonic partials lend themselves easily to the phase vocoder. Another choice is to use some kind of filter technique, since with filters it is easier to capture noisy signals. Linear prediction [4] is an example of this method. We also discarded this technique due to excessive complexity and lack of flexibility for further transformation. In addition the current linear prediction techniques give severe bandwidth errors for the resonant modes.

Our choice has been to analyze and model the prominent modes and the noise separately, using a different technique for each one. The partials are analyzed with a tracking phase vocoder and synthesized with additive synthesis. The noisy component can be approximated very efficiently by noise and modeled with a linear time-invariant filter driven by an exponentially decaying noise source.

2.1.1. Analysis/Synthesis of the prominent modes

We want a technique that finds the prominent partials and tracks them in time. We have used a tracking phase vocoder developed by Julius Smith at CCRMA. Smith's program was originally designed for the analysis and synthesis of piano tones but it has proven to be useful for many other sounds. This program follows the peaks in a series of magnitude spectra (computed using the Fast Fourier Transform) taken over the duration of a sound. It tracks the amplitude and frequency trajectories of the n most prominent peaks, where n is any integer specified by the user. The output is a frequency envelope plus an amplitude envelope for each partial; but in all the bar percussion instruments, the frequency trace is flat and can be left out. Another similar tracking phase vocoder was developed by Dolson [1].

For a good resynthesis of most of the instruments, it has been found that we only need the amplitude envelopes of the first three peaks and their frequency values. Then, to model the entire frequency range, the analyses of two or three notes per octave gives enough data. The amplitude envelopes can be simplified using a line-segment approximation technique [12] and they are interpolated to obtain the envelopes of the intermediate notes.

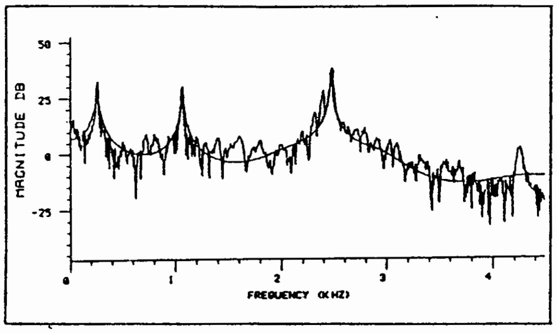

Figure 5. 12-order all-pole filter fitted to the magnitude spectrum of Fig. 3 with linear prediction.

2.1.2. Analysis/Synthesis of the noise component

In order to isolate the noise component we have to subtract the prominent modes, these being modeled with the technique discussed in the previous section. The direct way is to band-stop filter each partial individually. This can be very tedious and gives a resulting sound with the frequency band around each peak completely removed. For our purpose we want a more continuous spectrum. The solution has been to use linear prediction to fit an all-pole filter to the magnitude spectrum of the sound (Fig. 5). This traces the prominent peaks but not the rest of the spectrum. Then by inverting the filter (converting it into an all-zero structure, Fig. 6) and passing the sound through it we obtain the noise remainder (Fig. 7).

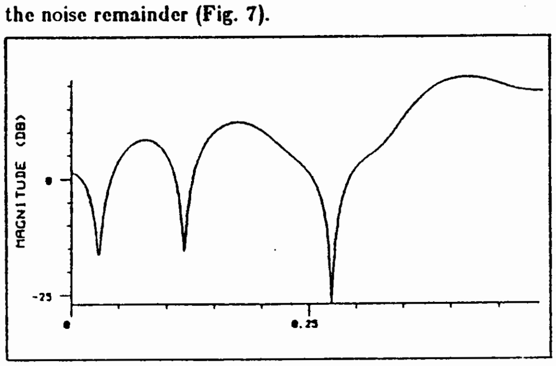

Figure 6. Frequency response of a 12-order all-zero filter obtained by inverting the all-pole filter of Fig. 5.

Once the noisy attack has been isolated, the task is to find a digital filter of low complexity that, given a certain input, will reproduce the noise remainder. This noise is different depending on the note, instrument and way of playing, and so we want a filter structure flexible enough to accomodate all the different "noises" while being easily manipulated.

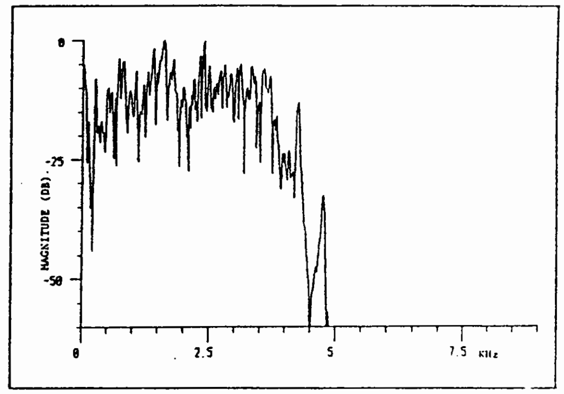

Figure 7. Spectrum of the noise remainder after removing the prominent peaks with the filter of Fig. 6.

We have simulated this part of the sound very accurately by using and exponentially decaying noise source to drive a filter, which has a frequency response that approximates the observed noise magnitude spectrum. Smith [11] discusses many methods of finding a filter structure that approximates a given spectrum. For example, linear prediction will fit an all-pole filter to a given spectrum, and by inverting the spectrum, the same technique will give an all-zero filter. To fit pole-zero filters is more complex, but Prony's method, equation error method, and Hankel's norm methods [11] can give good results.

For our requirements an exact fit is not needed. Flexibility is our main concern, and none of the methods that we have tried gives an easily changed filter structure. The most efficient and easiest way has been to design second-order pole-zero filters by hand and cascade them. Two such sections have proved to be sufficient in most cases.

The following formula is the transfer function of two second-order pole-zero sections, cascaded:

$$H(z) = g\frac{(1 + a_1 z^{-1} + a_2 z^{-2})}{(1 + b_1 z^{-1} + b_2 z^{-2})}\frac{(1 + c_1 z^{-1} + c_2 z^{-2})}{(1 + d_1 z^{-1} + d_2 z^{-2})}$$

where g is the gain, and a, b, c, and d are the filter coefficients.

Each section gives a peak and a notch in the frequency response of the overall filter. By controlling their placement, gain, and bandwidth, we can approximate any spectrum, and it is easily changed to fit another one. By adding sections we add resonances and notches.

To control the center-frequency and shape of each resonance or notch, it is convenient to use the filter's polar representation [10]. For a single second-order section we get

$$H_i(z) = g\frac{1 - 2R_{n,i}\cos(w_{n,i})z^{-1} + R_{n,i}^2 z^{-2}}{1 - 2R_{r,i}\cos(w_{r,i})z^{-1} + R_{r,i}^2 z^{-2}}$$

i = section number, r = resonance, n =notch, g = gain, R = pole or zero radius, w =resonance or notch frequeucy in radians $= 2\pi f/f_s$, f = center frequency in Hz, f_s = sampling rate in Hz.

By specifying $w_{r,i}$ we control the center frequency of the resonance, and with $R_{r,i}$ the gain at the peak and the bandwidth of the resonance. To control the notch we use $w_{n,i}$ and $R_{n,i}$.

Fig. 8 shows the frequency response of a fourth-order filter, constructed by cascading two second-order pole-zero sections. It fits approximately the noise spectrum of Fig. 7.

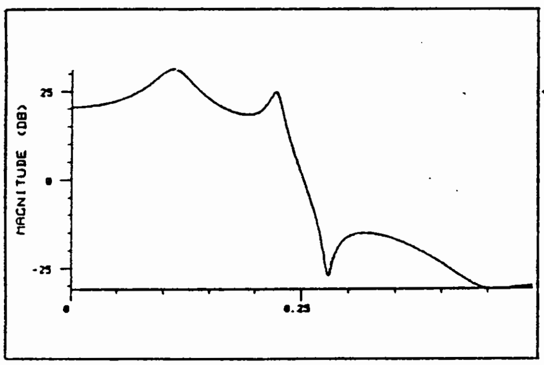

Figure 8. Frequency response of a fourth-order filter that approximates the spectrum of Fig. 7. resonances: $f_{r,1} = 2000Hz$, $R_{r,1} = .9$, $f_{r,2} = 4000Hz$, $R_{r,2} = .064$; notches: $f_{n,1} = 5000Hz$, $R_{n,1} = .98$, $f_{n,2} = 8000Hz$, $R_{n,2} = .8$.

2.2. Resonator model

A tube, closed at one end, can be modeled by the structure of Fig. 9. It consists of two delay lines, each one supporting one quarter of a period of the input waveform propagating in opposite directions, a lowpass filter, $H(z)$, to simulate the reflection of the closed termination, and a multiplication by -1, to simulate the wave inversion that takes place at the open end of the tube.

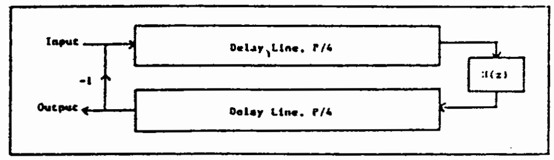

Figure 9. Block diagram for the resonator model. P is the period.

Using the simplification of the Karplus-Strong algorithm we can convert the structure of Fig. 9 into the structure of Fig. 10.

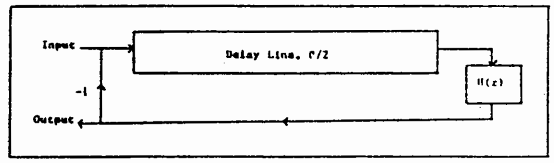

Figure 10. Block diagram for the simplified resonator model.

The structure in Fig. 10 is practically the same as the Karplus-Strong model for the plucked string, the only difference is that in our case the delay line is half as long in order the get only the odd partials. We can use all the extensions that Jaffe and Smith [3] did to the plucked string algorithm. Important extensions to implement are: the tuning, decay time shortening, and decay time stretching. The resonator has to be tuned to the exact same frequency as the fundamental of the bar, and the structure of Fig. 10 gives a crude tuning, especially at high frequencies. Jaffe and Smith use an allpass filter in the feed-back loop to make up for the difference between the frequency given by that structure and the desired frequency. It is also important to have control over the decay time for a more realistic simulation as well as for musical flexibility. By controlling the coefficients of the lowpass filter $H(z)$, Jaffe and Smith are able to specify the decay time, or, in our case, the resonating time.

2.3. Mallet-bar-resonator coupling

By putting together all the sections, the structure of Fig. 11 is obtained. To control the mallet-bar interaction we need to change the amplitude envelopes of the sinusoidal generators (OSC) and the amplitude of the noise generator. Mallet velocity, position of the strike, and type of head (hardness and size) can all be simulated in this way. The bar-resonator connection is controlled just by specifying the amount of the bar output that goes into the resonator.

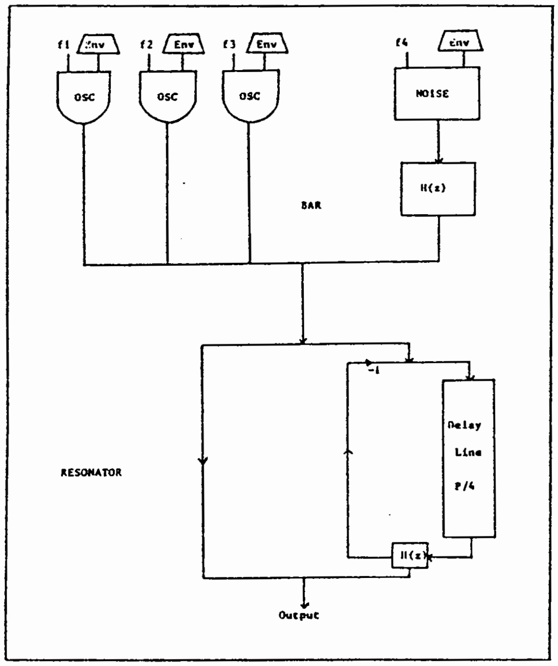

Figure 11. Block diagram of the overall model in a MUSIC V-like notation.

3. Conclusion

A computer model for bar percussion instruments has been presented which, at low cost, captures the most musically important aspects of the real instruments. Different methods were integrated into a unity, each one capturing a specific aspect of the sound.

By combining Discrete Fourier Transform methods (phase vocoder) and filter techniques (ex. linear prediction), we can analyze and model sounds in situations when each one of the techniques separately cannot give good results. This is especially appropriate for sounds that have a deterministic component plus a non-deterministic one.

Acknowledgments

Most of the tools used in this research were developed and implemented at CCRMA by Julius Smith, to whom I am most grateful for his teachings and guidance.

References

[1] M. B. Dolson, "A Tracking Phase Vocoder and its Use in the Analysis of Ensemble Sounds," Ph.D. Dissertation, California Institute of Technology, 1983.

[2] J. W. Gordon and J. Strawn, "An Introduction to the Phase Vocoder," In J. Strawn, ed., *Digital Audio Signal Processing: An Anthology*. William Kaufmann, Inc., Los Altos, California, 1985.

[3] D. Jaffe and J. O. Smith, "Extensions of the Karplus-Strong Plucked String Algorithm," *Computer Music J.*, vol. 7, no. 2, pp. 56–69, 1983.

[4] J. Makhoul, "Linear Prediction: A Tutorial Review," *Proc. IEEE*, vol. 63, pp. 561–580, April 1975.

[5] M. E. McIntyre and J. Woodhouse, "On the Fundamentals of Bowed String Dynamics," *Acustica*, vol. 43, no. 2, pp. 93–108, 1960.

[6] J. L. Moore, "Acoustics of Bar Percussion Instruments," Ph.D. Dissertation, Music Department, The Ohio State University, 1970.

[7] P. M. Morse, *Vibration and Sound*, published by the American Institute of Physics for the Acoustical Society of America, 1976 (1st ed. 1936, 2nd ed. 1948).

[8] K. Karplus and A. Strong, "Digital Synthesis of Plucked-String and Drum Timbres," *Computer Music J.*, vol. 7, no. 2, pp. 43–55, 1983.

[9] T. D. Rossing, "Acoustics of Percussion Instruments," Part I, The Instrumentalist, May 1976.

[10] J. O. Smith, "Introduction to Digital Filter Theory," In J. Strawn, ed., *Digital Audio Signal Processing: An Anthology*. William Kaufmann, Inc., Los Altos, California, 1985.

[11] J. O. Smith, "Methods for Digital Filter Design and System Identification with Application to the Violin," Ph.D. Dissertation, Elec. Eng. Dept., Stanford University, June 1983.

[12] J. Strawn, "Approximation and Syntactic Analysis of Amplitude and Frequency Functions for Digital Sound Synthesis," *Computer Music J.*, vol. 4, no. 3, pp. 3–24, 1980.