4.5 Fractional Delay Operations with Allpass Filters

The previous sections of this chapter have concentrated on the FIR implementation techniques of fractional delay waveguide filters. In Sections 4.6 and 4.7 similar structures are developed with allpass FD filters. The idea of using allpass filters in waveguide models is not new—Jaffe and Smith (1983) tuned their string synthesizer with a first-order allpass filter. However, to our knowledge allpass filters have not been used for waveguide junctions before. The initial work in this field was published in Välimäki and Karjalainen (1995) and Välimäki (1995b). Very recently Cook (1995) used allpass filters for realizing finger holes for a waveguide flute model.

There are major differences between the FIR and allpass implementations of FDWFs. These differences are caused by the fact that a given impulse response of an FIR interpolating filter can be used for approximating both a fractional delay d and a *complementary fractional delay* $\delta = 1 - d$ (by flipping the impulse response). The complementary fractional delay (CFD) is needed in FDWFs because the aim is to use two digital delay lines to represent a waveguide and connect the junctions between the unit delays of these lines. For a particular FD junction, d is the delay from the last integral sampling point while δ is the delay to the next integral sampling point.

In contrast to an FIR filter, the coefficients of an allpass FD filter are completely different for fractional delays d and δ. For this reason it is not advantageous to use a one-multiplier scattering junction as a starting point for derivation of an allpass FD junction: each junction would require altogether four allpass filters, two for approximating z^{-d} and another two for $z^{-\delta}$. Furthermore, the approximation error in the transmission function of this kind of junction would be unacceptable (see Section 4.6.4).

4.5.1 Implementing a Variable Delay Line Using an Allpass Filter

In Chapter 3 it was shown that there are two techniques for the design of fractional delay filters: FIR and IIR filter approximation. While the use of a nonrecursive interpolating filter leads to a structure that is easier to understand, it is useful to examine the use of a recursive filter as well. As discussed in Section 3.4, the digital allpass filter is a good choice for an IIR FD filter. Now the implementation of a variable-length delay line using an allpass filter is described.

The use of an allpass filter is clarified using a first-order filter as an example. Higher order allpass structures are straightforward extensions. The length of a delay line can be varied using a first-order allpass filter as depicted in Fig. 4.30. It may be useful to use a

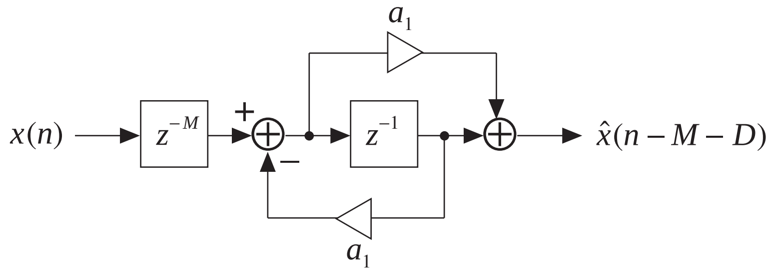

Fig. 4.30 Implementation of a variable-length digital delay line by means of a first-order allpass filter.

form of an allpass filter that has a minimum number of delay elements, since the delays can be shared with the delay line as shown in Fig. 4.30. This scheme can be extended for higher-order allpass filters.

The structure illustrated in Fig. 4.30 implements the following difference equation

$$y(n) = \hat{x}(n - M - D) = a_1\left[x(n - M) - y(n - 1)\right] + x(n - M - 1) \quad (4.36)$$

Note, however, that the structure in Fig. 4.30 is not a direct-form structure and thus it may not be apparent that the underlying difference equation is in fact Eq. (4.36).

The coefficients of the interpolating allpass filter can be computed applying one of the techniques presented in Section 3.4. The simplest and the most useful allpass design technique for audio signal processing is the maximally-flat group delay, or Thiran, design (see Section 3.4.3).

4.5.2 Implementing Interpolation and Deinterpolation with Allpass Filters

Figure 4.31 illustrates the first-order allpass filter structure that realizes an output at a fractional point of a delay line. This system is comparable to that presented in Fig. 4.30. The principal difference between these structures is that in Fig. 4.31 we assume that the delay line continues after the fractional point and the allpass filter merely approximates the signal value at one point along that delay line—not at its end point, like in Fig. 4.30. When the filter coefficient a_1 has been chosen in an appropriate way, the output $y(n)$ of this structure approximates the value of signal $x(n)$ at point $M + D$ along the delay line, that is

$$y(n) = \hat{x}(n - M - D) \quad (4.37)$$

The structure shown in Fig. 4.31 directly implements the following difference equation:

$$y(n) = a_1\left[x(n - M) - y(n - 1)\right] + x(n - M - 1) \quad (4.38)$$

This time we use the first-order allpass structure with two unit delay elements and a minimum number of multiplications. This choice is appropriate since one of the unit delays can be shared with the delay line to be processed and the other unit delay is reserved for the recursive part of the filter. The allpass structure with a minimum number of delay elements would not be more efficient since the recursion prevents the sharing of the unit delay with the delay line. For different structures for the implementation of first-order allpass filters, see Mitra and Hirano (1974).

Next we develop a recursive version of deinterpolation using an allpass filter. The

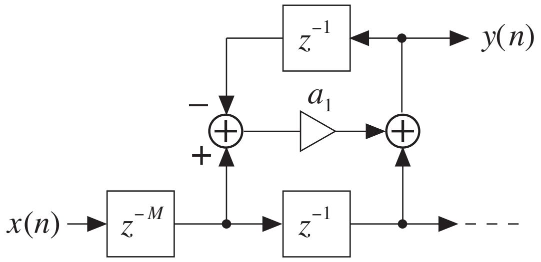

Fig. 4.31 A first-order allpass interpolator is used for the implementation of an output at a fractional point of a delay line.

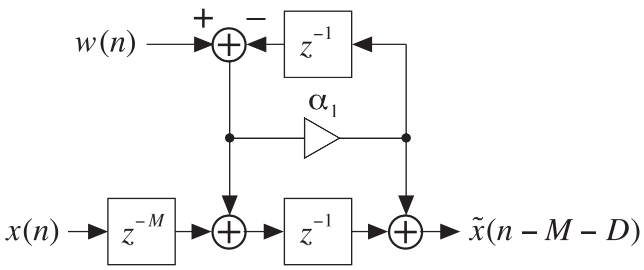

Fig. 4.32 A first-order allpass deinterpolator implements an input at a fractional point of a delay line.

allpass deinterpolation scheme is clarified using a first-order allpass filter as an example. In Fig. 4.32, an implementation of a fractional delay input to a delay line is shown. This allpass filter structure is the *transpose* of that in Fig. 4.31, i.e., an allpass deinterpolator may be realized using the transpose structure of the allpass interpolator. This is an analogy to the FIR filter structures studied earlier.

The transfer functions of the first-order allpass interpolator and deinterpolator, respectively, are

$$A_{\mathrm{i}}(z) = \frac{a_1 + z^{-1}}{1 + a_1 z^{-1}} \tag{4.39}$$

$$A_{\mathrm{d}}(z) = \frac{\alpha_1 + z^{-1}}{1 + \alpha_1 z^{-1}} \tag{4.40}$$

Note that the filter coefficients a_1 and α_1 are not equal in these structures and they do not, in general, have a simple relation. In allpass interpolation the coefficient should be chosen so that the filter approximates the delay $\lfloor D \rfloor + d$, where d is the fractional delay. The coefficient α_1 of the allpass deinterpolator, however, should be chosen so that the allpass filter approximates the delay $\lfloor D \rfloor + \delta$, where δ is the complementary fractional delay.

4.6 Interpolated Waveguide Model: Allpass Filter Approach

Here we derive the allpass filter counterpart for the interpolated scattering junction of two tubes.

4.6.1 Derivation of the Allpass Fractional Delay Junction

Figure 4.33a illustrates the ideal interpolated scattering junction where the fractional delay elements z^{-d} and $z^{-\delta}$ are explicitly shown. Figure 4.33b shows a modified version of this junction. Now the four FD elements have been pushed through the adders and branch nodes. This is legal because linear time-invariant systems commute.

In Fig. 4.33b, there are altogether eight FD elements. The FD elements z^{-d} and $z^{-\delta}$ can be combined into a single unit delay element (since $d + \delta = 1$), both in the upper and in the lower line, as presented in Fig. 4.33c. The two FD elements z^{-d} on the left of Fig. 4.33b can be combined into a single block z^{-2d}, and it is possible to do the same for the two CFD blocks on the right.

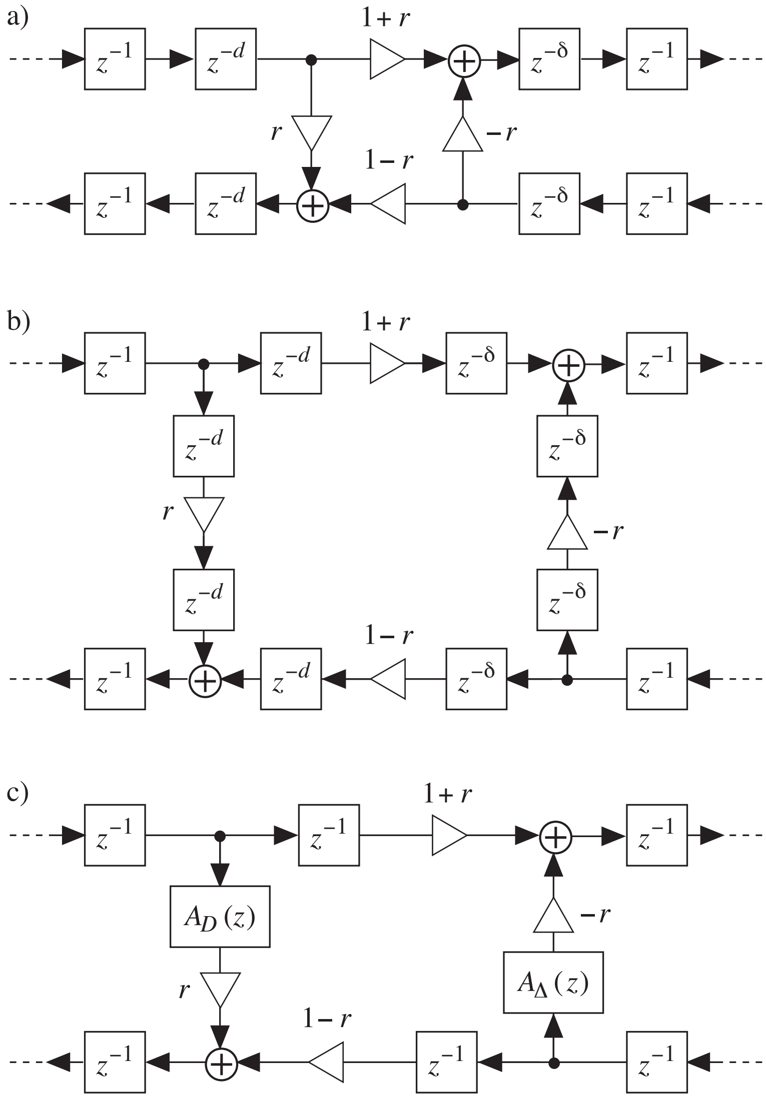

Fig. 4.33 Derivation of the FD scattering junction employing commutation. (a) The original FD scattering junction. (b) The FD elements have been pushed through the branch nodes and adders. (c) The structure that is implemented with two first-order allpass filters.

The block diagram presented in Fig. 4.33c is one possible allpass implementation of the FD junction. Other configurations can also be derived but this one has some desirable properties. We realize the "double" fractional delay elements using allpass filters $A_D(z)$ and $A_\Delta(z)$. Consequently, the desired delay D and complementary total delay Δ are defined by

$$D = 2d \tag{4.41}$$

$$\Delta = 2\delta = 2 - D \tag{4.42}$$

This implies that when the noninteger part d of the position of the junction can have values in the interval [0, 1], the delays implemented by the two allpass filters will be in the interval [0, 2].

Let us consider implementation of the allpass FD junction using the Thiran allpass filter that was discussed in Sections 3.4.3 and 3.4.4. The choice of the optimal range for the delay parameter D was studied in Section 3.4.4. For example, the first-order Thiran interpolator yields best results when the delay varies between $D_0 = 0.42$ and $D_0 + 1 = 1.42$. For Thiran allpass filters of different order, the optimal range is also different. However, an easy rule of thumb for the delay parameter is to maintain it in the range $[N - 0.5, N + 1.5]$ where N is the order of the allpass filter.

The implementation of the allpass FD junction described above needs special care, since both the allpass filters used in this structure must be able to approximate fractional delays within two unit delays. In the case of the first-order Thiran allpass filter, the junction can be implemented as shown in Fig. 4.33c only when the junction is located 0.25–0.75 samples away from the nearest unit delay. Only then is the delay of the allpass filters $A_D(z)$ and $A_\Delta(z)$ within the desired range [0.5, 1.5].

When the delay d is smaller than 0.25 or larger than 0.75, the delay D approximated by the filter $A_D(z)$ goes beyond the desired range. In order to change the delay of the reflection function, we may add or subtract unit delays to or from that allpass filter simply by changing its input point one unit delay to either direction. If the input point of the allpass filter $A_D(z)$ is moved one unit delay to the left in Fig. 4.33c (i.e., the input of the allpass filter is taken one sample interval earlier), then one sample of delay must be added to D in order to maintain the overall delay. Similarly, if the input point is switched one unit delay to the right (i.e., the input to the allpass filter is delayed by one unit delay), one sample of delay must be subtracted from the filter's nominal delay D. Another possibility is to change the location of the filter's output point, i.e., the point where the output of the allpass filter $A_D(z)$ is added to the lower delay line in Fig. 4.33c. These are two alternative methods that help to maintain the nominal delay of the allpass filter within the desired range. One of these techniques must be used for the complementary allpass filter $A_\Delta(z)$ as well.

The above discussion concerned the first-order Thiran allpass filter. However, similar procedures—change of the location of input or output point—should be used to maintain the delay of the filter in its best range. In the case of higher order allpass filters, one or more unit delays must usually be added to the filter's nominal delay value because these filters cannot implement an arbitrarily small delay. Remember that the Thiran allpass filter is unstable for delay values $D < N - 1$. For example, the second-order Thiran filter cannot approximate a delay smaller than one sample.

4.6.2 Approximation Errors due to an Allpass Fractional Delay Junction

Let us consider the approximation error due to allpass interpolators. In the structure of Fig. 4.33c, there are no FD elements in the straight path in neither direction. Thus there is *no approximation error in the transmission function* through this allpass FD junction. The only limitation that this solution brings about is that the point where the transmission occurs is not explicitly implemented and the signal value at that point cannot be directly obtained. At the open end of a tube this causes an error in the delay of the output signal (radiated sound). This does not affect the formant structure nor other important properties of the signal and can be easily avoided—if desired—by incorporating an additional FD filter at the output.

The reflection from the FD junction suffers from an approximation error since the "double" FDs in Fig. 4.33b are realized using allpass filter approximations $A_D(z)$ and $A_\Delta(z)$. When the MF (Thiran) allpass FD filters are used, the nature of this error can be predicted from the phase delay responses given in Section 3.4.3.

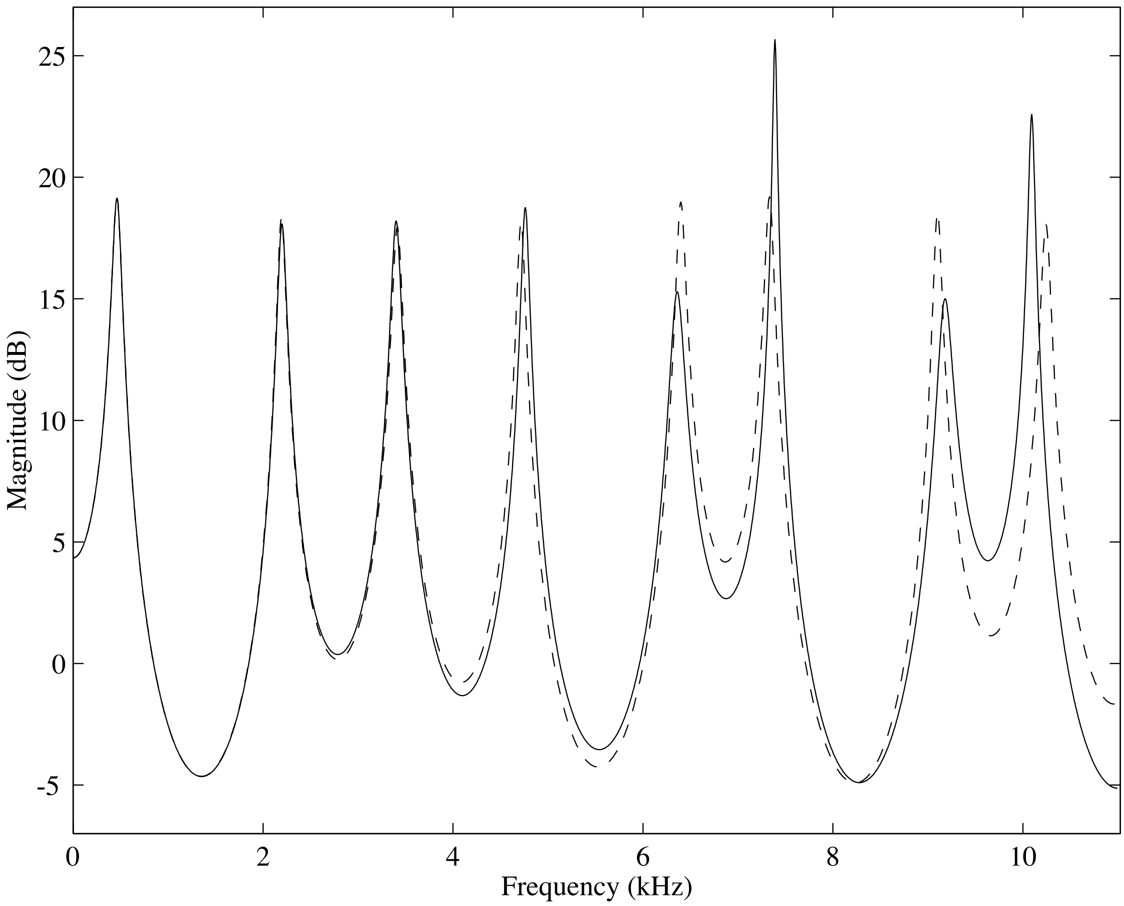

Fig. 4.34 Transfer function of a two-tube system realized using first-order Thiran allpass filters (solid line). For comparison, the simulation result obtained with ideal interpolation (dashed line) is also given. The location of the junction is $d = 0.25$.

The approximation error is zero when the nominal delay of the filter is 1.0. This implies that the approximation error disappears when the junction is located half way between unit delay, i.e., $d = 0.5$. This is in sharp contrast with the FIR FD junctions, where the approximation error is largest when $d = 0.5$.

For other values of fractional delay, the approximation error is nonzero, and its effect can be characterized as follows. At very low frequencies the junction (impedance discontinuity) is located very nearly in the correct position, but with increasing frequency the junction tends to move towards some erroneous location. This is because the Thiran allpass filter approximates the delay accurately at low frequencies but poorly at high frequencies. We may assume that at low frequencies the formant structure of a tube model implemented with these allpass filters is similar to the ideal one, whereas at higher frequencies the center frequencies and amplitudes of formants are incorrect. The following example confirms that this is actually the case.

4.6.3 Simulation of a Two-Tube Model with Allpass Filters

In order to compare the properties of allpass and FIR implementations of the interpolated waveguide model, a two-tube model with one FD junction was simulated. The example system is the same as the one presented in Section 4.3.7. The total length of the

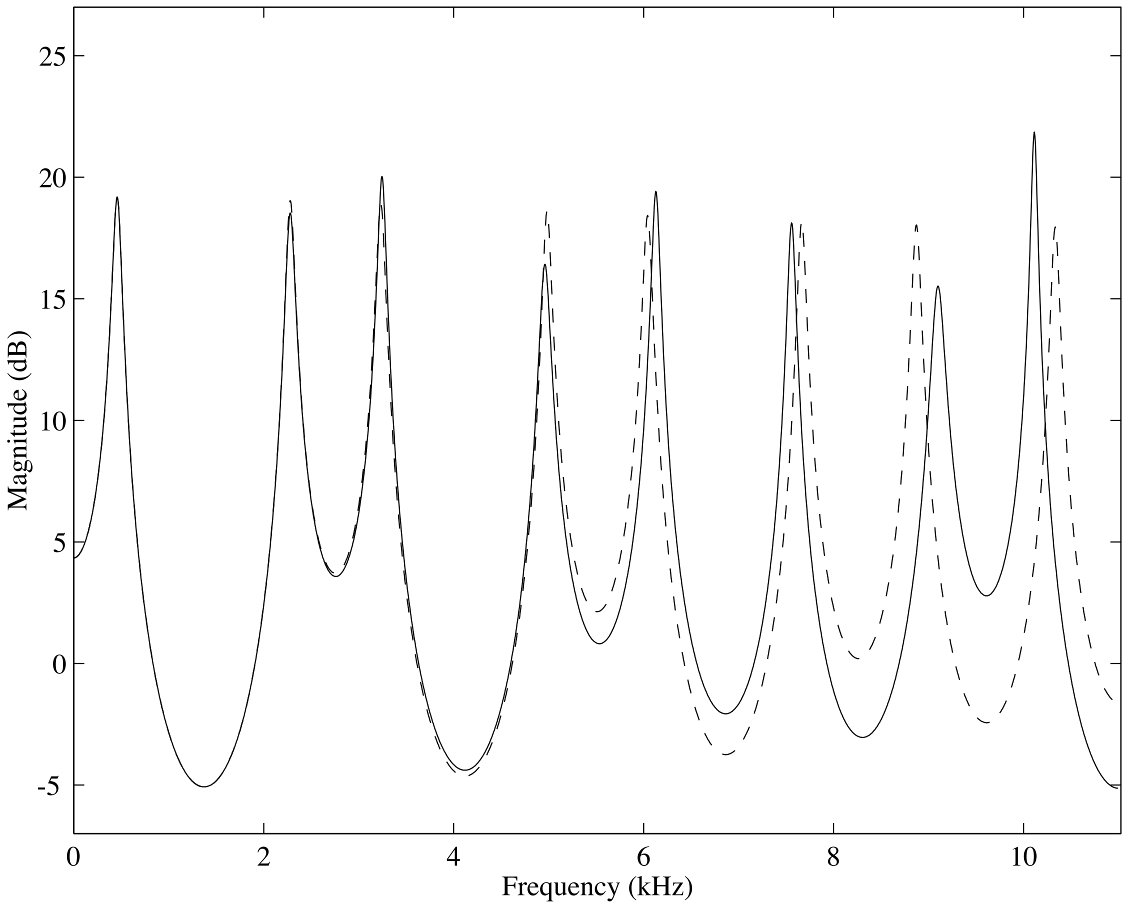

Fig. 4.35 Transfer function of a two-tube system realized using first-order Thiran allpass filters (solid line). For comparison, the simulation result obtained with ideal interpolation (dashed line) is also given. The location of the junction is $d = 0.75$.

tube system is eight unit delays. The junction between the tubes is located between the third and fourth unit delays, and the reflection coefficient is –0.5. One end of the tube is assumed to be closed and the other one open. To simplify the example, the terminations have been approximated with purely resistive loads, $r_1 = 0.9$ and $r_2 = -0.9$.

Figures 4.34 and 4.35 show the magnitude of the transfer function of this system for two junction positions. The solid line has been computed using first-order MF allpass filters and the dashed line was obtained with ideal interpolators. The sampling rate is 22 kHz. In Fig. 4.34, the location of the FD junction is $d = 0.25$ corresponding to delay parameters $D = 0.5$ and $\Delta = 1.5$. In Fig. 4.35, the junction is situated at $d = 0.75$ between the unit delays, which means that $D = 1.5$ and $\Delta = 0.5$. These can be considered the 'worst-case' situations since the delay values are on the edge of the range.

The transfer functions of the ideal and approximative system are nearly identical at low frequencies but at high frequencies the deviation is quite large. However, the magnitude error of the fourth formant (at about 4.7 kHz) is only 1 dB in Fig. 4.34 and 2 dB in Fig. 4.35. These figures should be compared with Figs. 4.24–4.27 that have been computed with FIR interpolators. It can be concluded that the result obtained with the first-order allpass filters is clearly better (at low and middle frequencies) than with the linear interpolators and about as good as with the third-order Lagrange interpolators.

The computational complexity of the allpass filter implementation is less than that of

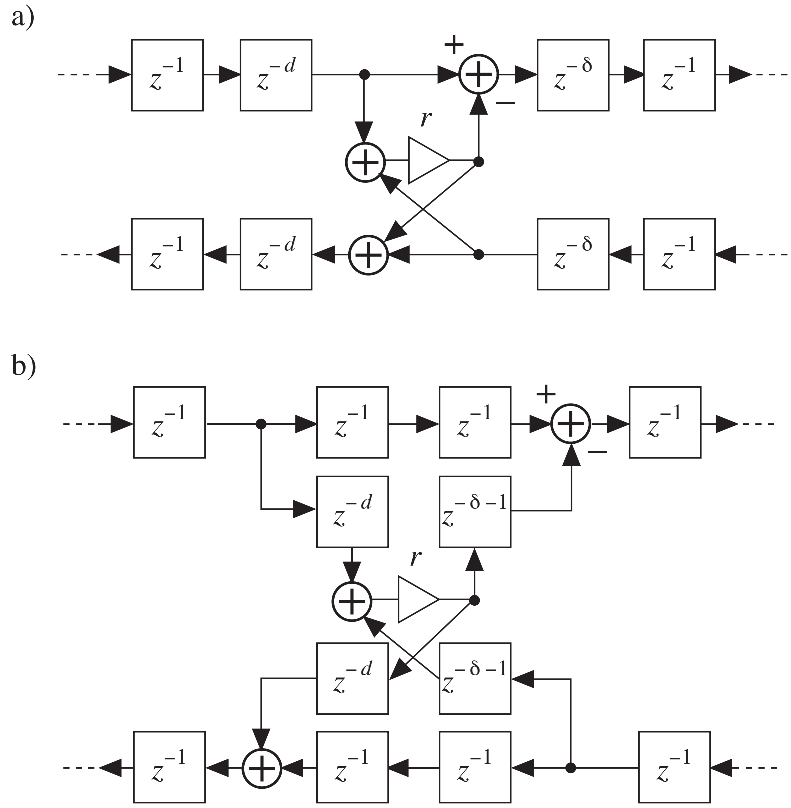

Fig. 4.36 An example of an alternative FD junction structure that at first seems to be correct. Deeper analysis reveals that an allpass implementation of this structure will not yield a stable waveguide model.

the FIR filter implementation: a first-order allpass filters is implemented with one multiplication and two additions and a second-order allpass filter can be realized with two multiplications and three additions, whereas a third-order FIR filter requires four multiplications and three additions. Two filters per FD junction are needed both in the case of allpass and FIR filters. Thus the allpass filter appears to be an attractive choice for the implementation of fractional delay waveguide models.

4.6.4 Alternative Allpass Filter Structure

Figure 4.33 presents one FD junction structure that can be used for implementing an interpolated waveguide model. Now we briefly discuss another structure that is perhaps more obvious than that of Fig. 4.33, and at first seems to be quite correct. Nevertheless, this structure is nonrealizable.

In Fig. 4.36a we present a one-multiplier scattering junction that is located between the sampling points of a waveguide. The FD and CFD elements are explicitly shown. It is advantageous to try to remove these FD elements from the waveguide. Then the FD junction can be connected to an arbitrary point on the waveguide and it is not needed to ‘cut’ the delay lines by adding a fractional delay filter in the middle of the unit delay chains. One way to achieve this is to commute the FD elements with the adder and branch nodes of the scattering junction. In Fig. 4.36b, all four FD elements have been pushed inside the scattering junction. Those copies of the FD and CFD elements that were left in the delay line chains, have been combined into unit delay elements. In addition, one extra unit delay element from the right-hand side has been pushed into the junction and has been combined with the CFD element. This helps to realize an allpass junction with better accuracy. The structure of 4.36 is aimed at an implementation using a first-order allpass filter. Then the fractional delays z^{-d} would be implemented with two allpass filters whose delay parameters would vary within [0, 1]. The delay parameters of the other two allpass filters (approximating transfer functions $z^{-\delta-1}$) would vary in the range [1, 2].

Unfortunately this design proves to be useless. Let us consider the transmission function through the junction from the left to the right. We denote the allpass filter approximations for the FD and CFD elements by $A_D(z)$ and $A_\Delta(z)$, respectively. The transmission function is then $z^{-2} - rA_D(z)A_\Delta(z)$. We would like this to be an allpass function or at least a passive function. However, the approximation errors in the two allpass filters assure that the phase function of $rA_D(z)A_\Delta(z)$ is nonlinear, and although this transfer function and also that of the other signal path, z^{-2}, both are allpass functions (the product of two allpass filters is also an allpass filter), the resulting transmission function does not have this property. In fact, a sum of two allpass filters is in general not an allpass function: the differences in the phase functions of the summed functions lead to cancellation or emphasis of different frequencies in its magnitude spectrum. This approach can be and has been used for designing lowpass and highpass filters, for example.

The above explanation implies that the transmission functions through the FD junction of Fig. 4.36b are not guaranteed to be either allpass or passive. This may endanger the stability of the overall waveguide system, which consists of several feedback loops. For this reason, the structure is useless—although in theory it has been derived in a correct way.

4.7 Allpass Fractional Delay Finger-Hole Model

Section 4.4 was concerned with implementation of a finger-hole model using FIR FD filters. Now we present a corresponding model realized using allpass filters (Välimäki, 1995b). It is worth noting that this is merely one possible configuration. Also in this case there are more but the one discussed here turns out be an efficient one in terms of computational cost. For an example of another approach, see Cook (1995).

The derivation follows the same guidelines as above in the case of the FD two-port scattering junction. We begin with the finger-hole junction where there are four reflection filters as illustrated in Fig. 4.37a: the two filters $R(z)$ simulate the reflection of the volume velocity waves and the two filters $1 + R(z)$ simulate the transmission of volume velocity through the finger hole. It is assumed that this finger hole is located at an arbitrary position along the waveguide bore model. Thus one of the unit delays of the digital waveguide has been split into two parts and the hole has been positioned between them. The fractional delays d and δ are explicitly shown in Fig. 4.37a.

We proceed by pushing the four FD elements inside the junction. When an FD element is pushed through a branch or summation node, this element has to be appended into each input path of that branch. This has been done in Fig. 4.37b. It is seen that now there are altogether eight FD elements.

The FD elements that are on the same path can be combined: the fractional delays z^{-d} and complementary FDs $z^{-\delta}$ can be combined into single unit delay elements (since $d + \delta = 1$) in the upper and in the lower lines (see Fig. 4.37c). The two FD elements z^{-d} on

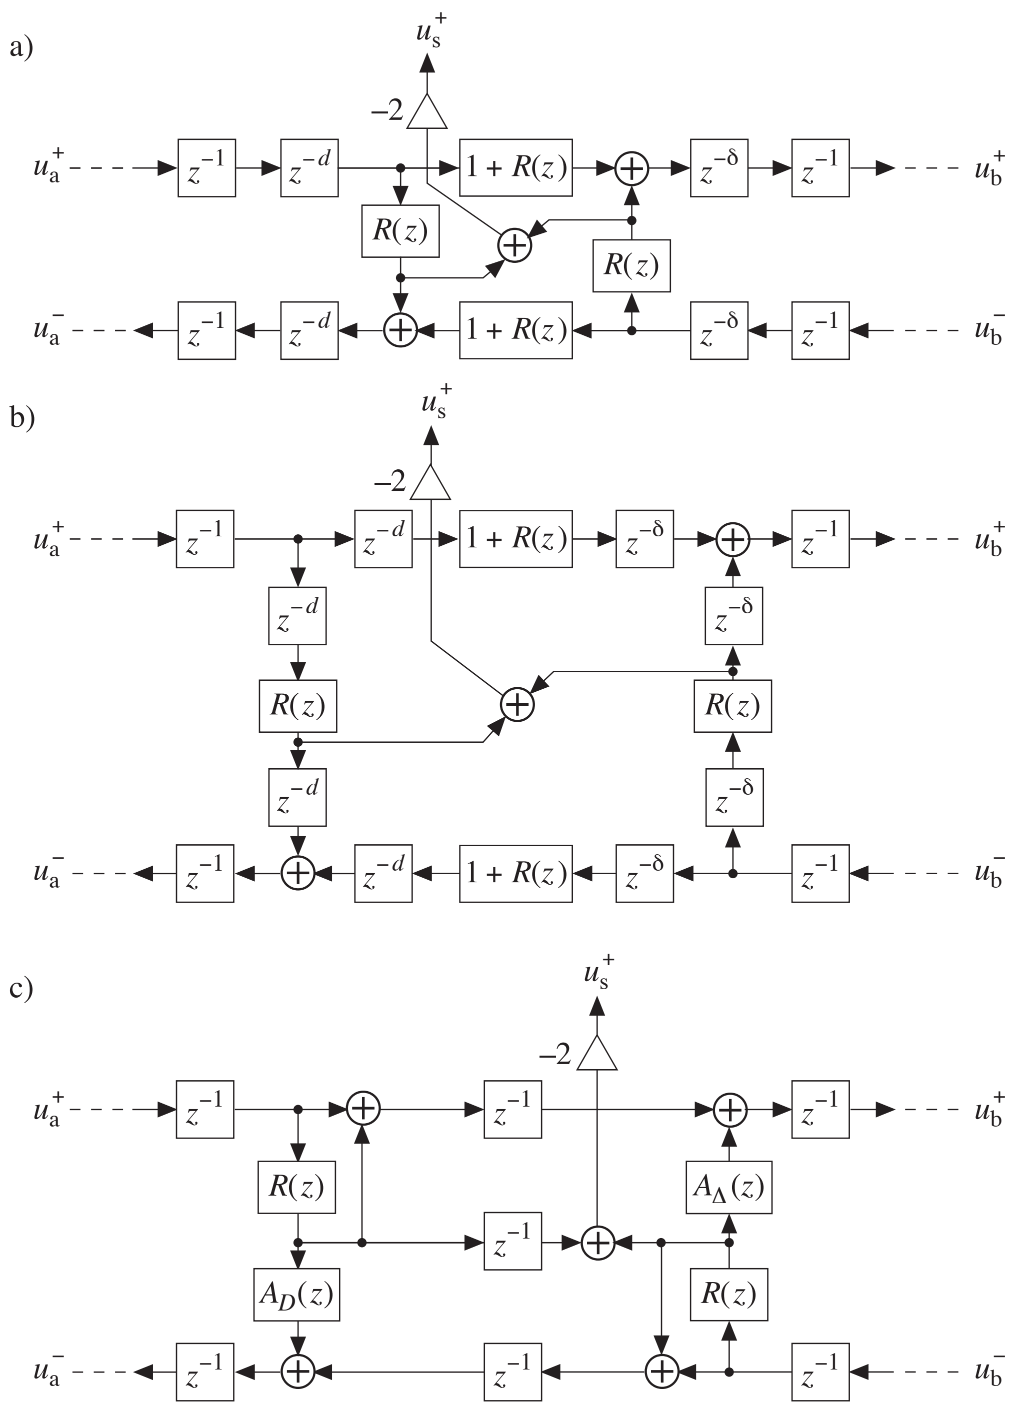

Fig. 4.37 Derivation of the finger-hole model structure where the fractional delays are approximated using allpass filters. a) The finger-hole junction located at an arbitrary point (between unit delays) of a digital waveguide. b) The FD elements have been pushed inside the junction. c) The final structure that can be implemented using two identical reflection filters $R(z)$ and two allpass filters.

the left of Fig. 4.37b can be combined into a single block z^{-2d} while the two CFD elements $z^{-\delta}$ on the right into $z^{-2\delta}$.

These double FD elements are implemented using allpass filters $A_D(z)$ and $A_\Delta(z)$. Consequently, for these allpass filters the desired delay D and complementary total delay Δ are defined by $D = 2d$ and $\Delta = 2\delta$. This implies that when the noninteger part d

of the position of the junction is varied in the interval [0, 1], the delays approximated by these two allpass filters have values in the interval [0, 2]. Also shown in Fig. 4.37c is an efficient implementation of transmission functions $1 + R(z)$ as a sum of the direct signal and the output of the filter $R(z)$.

The unit delay element in the middle of Fig. 4.37c adjusts the delay between the signals from the upper and lower delay line before their addition. Here we have simplified the structure considerably by neglecting the FD elements z^{-d} and $z^{-\delta}$ that should precede the adder in the two signal paths. It is believed that this approximation does not affect the properties of the finger-hole model. In effect we have only neglected a small delay at the output of the system. This trick does, however, reduce the number of additions and multiplications, since it saves the implementation of one more fractional delay filter.

The finger-hole model of Fig. 4.37c can implemented using first-order allpass filters $A_D(z)$ and $A_\Delta(z)$. The total computational load of the finger-hole model is thus not very large since only two one-pole reflection filters and two first-order allpass filters need to be implemented.

4.8 Conclusions

In this chapter, implementation techniques for fractional delay waveguide filters have been discussed. Interpolation is a standard method for estimating the value of a signal between its known sample values. A novel technique, deinterpolation, was introduced. It can be interpreted as an inverse operation of interpolation in the sense that interpolation is used for computing the output at an arbitrary point of a discrete-time delay line, whereas deinterpolation is used for locating a corresponding input point. Both operations can be implemented using recursive or nonrecursive digital filters, but deinterpolation is more intuitively implemented using an FIR filter structure. The use of allpass filters for interpolation and deinterpolation was also discussed.

Interpolation and deinterpolation provide means for creating spatially continuous digital waveguide models. The simplest examples of such systems are a variable-length delay line and a variable-length waveguide. These and the more elaborate waveguide structures are called fractional delay waveguide filters. They are constructed of arbitrarily many digital waveguides that are connected to each other at fractional points. This approach was applied to the time-domain modeling of acoustic tubes and realization of two-port and three-port fractional delay junctions. Both FIR and allpass fractional delay filters were employed for this purpose.

The properties of the FIR and allpass techniques were compared. The results can be summarized in the following way: FIR structures are conceptually simpler since the locations of filter taps have a direct relation to the impulse response of the filter and thus also to the locations of the interpolated input and output points along a digital waveguide. The use of maximally-flat FIR FD filters or maximally-flat group-delay allpass filters leads to a low-frequency approximation of an ideal tube structure. Oversampling by a small integral factor is required to obtain a reasonable quality of approximation. The benefits of allpass filter structures are an identically flat magnitude response in the FD approximation and a smaller approximation error with a lower-order filter than that obtainable with an FIR filter structure with a comparable number of arithmetic operations.