

Use of Truncated Infinite Impulse Response (TIIR) Filters in Implementing Efficient Digital Waveguide Models of Flared Horns and Piecewise Conical Bores with Unstable One-Pole Filter Elements

Maarten van Walstijn

Ph.D. Student, Faculty of Music, University of Edinburgh
12 Nicholson Square, Edinburgh EH8 9DF

Julius O. Smith III*

Assoc. Prof., Center for Computer Research in Music and Acoustics (CCRMA)†
Stanford University Stanford University, Stanford, CA 94305 USA

Original version published in the
Proceedings of the International Symposium on Musical Acoustics
(ISMA-98 at Leavenworth, Washington), pp. 309–314, June 28, 1998.

Abstract—We describe computational modeling of flaring horns and piecewise conical bores using “Truncated Infinite Impulse Response” (TIIR) digital filtering techniques. The approach yields highly efficient and accurate computational models and is therefore appropriate for real-time simulations of woodwind and brass musical instruments.

1 Introduction

Efficient time-domain models of wind instruments generally involve at least three components: the mouthpiece and (lip-)reed interface, the main bore plus any valve segments or tone holes, and the bell. Mouthpiece, lip and reed models pose challenging problems which continue to be addressed. A uniform cylindrical or conical bore is easily modeled in the digital waveguide formulation [7] using a single delay line to represent the associated round-trip propagation delay. (Associated losses may be lumped elsewhere, such as at the mouthpiece.) The bell of the instrument, assuming linearity, may be characterized by its reflection and transmission impulse responses, and is therefore efficiently modeled as a “lumped” horn reflectance and transmittance. Other non-uniform tubular segments are potentially well modeled using the well known piecewise conical bore formulation, but its time-domain application has so far been strongly limited due to numerical problems. In this paper we present practical and efficient methods for modeling and implementing piecewise conical bores and horn reflectances.¹

*Work supported in part by Staccato Systems, Inc.

†<http://ccrma.stanford.edu/>

¹Since the transmittance does not participate in the sound generation mechanism, and can be modeled as an external component, we consider explicitly only the reflectance from this point forward.

A straightforward but computationally expensive discrete-time model of the bell is to use its impulse response as a convolution filter. This yields a trivially designed finite-impulse-response (FIR) filter model for the reflectance. Alternatively, an infinite-impulse-response (IIR) digital filter can be designed to approximate the bell reflection response. In general, IIR filters can approximate a given impulse response with much less computation because they are recursive. However, prevalent phase-sensitive IIR filter-design methods perform poorly when applied to a measured bell reflectance. This is due mainly to the long, slowly rising, quasi-exponential portion of the time-domain response, arising from the smoothly flaring bore profile that is characteristic of musical horns. As a result, there is a need for more effective digital filter design techniques in this context.

Inspections of horn reflectances in the time domain suggest that a natural modeling approach might consist of dividing the response into at least two sections: an initial growing exponential, followed by a more oscillatory “tail.” The tail can be faithfully modeled using more conventional filter-design methods. The most efficient way to model a growing exponential is by means of an *unstable* one-pole filter, just as we encounter in piecewise conical acoustic tubes [4]. Thus, the problems of modeling flared horns and piecewise conical bores give rise to the problem of how to utilize unstable digital filters as modeling elements without running into numerical problems.

It turns out that growing exponential impulse-response segments can be efficiently and practically devised using “Truncated Infinite Impulse Response” (TIIR) digital filtering techniques [10]. The basic idea of a TIIR filter is to synthesize an FIR filter as an IIR filter minus a delayed “tail canceling” IIR filter (which has the same poles as the first). That is, the second IIR filter generates a copy of the “tail” of the first so that it can be subtracted off, thus creating an FIR filter. When all IIR poles are stable, TIIR filters are straightforward. In the unstable case, the straightforward implementation fails numerically: While the filter tails always cancel in principle, the exponential growth of the roundoff-error eventually dominates. Thus, in the unstable case, TIIR filters must switch between two alternate instances of the desired TIIR filter (i.e., two pairs of tail-canceling IIR filters). The state of the “off-duty” filter is cleared in order to zero out the accumulating round-off noise. The key observation is that, because the desired TIIR filter functions as an FIR filter, it reaches exact “steady state” after only N samples, where N is the length of the synthesized FIR filter. As a result, a “fresh instance” of the TIIR filter, when “ramped up” from the zero state, is ready to be switched in exactly after only N samples, even though the component IIR filters have not yet reached the same internal state as those of the TIIR filter being switched out.

An empirically derived trumpet bell reflectance was found to have an impulse response duration on the order of 10 ms (which is on the order of 400 samples at a 44.1 kHz sampling rate). While a length 400 FIR filter can faithfully model the trumpet-bell reflectance, use of TIIR methods reduces the complexity by well over an order of magnitude with good matching of the principal time-domain and frequency-domain features (accurately preserving the horn resonances in particular).

In the remainder of this paper, we will discuss (1) the needed class of TIIR filters, (2) TIIR modeling of flared horns, using a theoretically derived Bessel horn reflectance as the desired response, (3) construction and calibration of a TIIR-based digital waveguide trumpet model, based on experimental data, and finally (4) TIIR modeling of piecewise conical acoustic bores.

2 TIIR Filters

An FIR filter can be constructed in general as the difference of two IIR filters [10]. The output of the second IIR filter is delayed, scaled, and subtracted, so as to cancel the “tail” of the first IIR

filter. The overall output is that of an FIR filter, but with great computational savings when the delay is large compared with the IIR order.

Figure 1 shows the case of a one-pole based TIIR filter which is sufficiently general for purposes of this paper; it can be configured for either a truncated growing exponential or a truncated constant impulse response. There are various choices of filter structure even in this simple case. Figure 1 shows the “shared delay” form. By “pushing” all four one-pole filters forward through the subtraction block, one obtains additionally the “shared dynamics” form suggested in [10]. For simplicity, however, we will describe the version in Fig. 1.

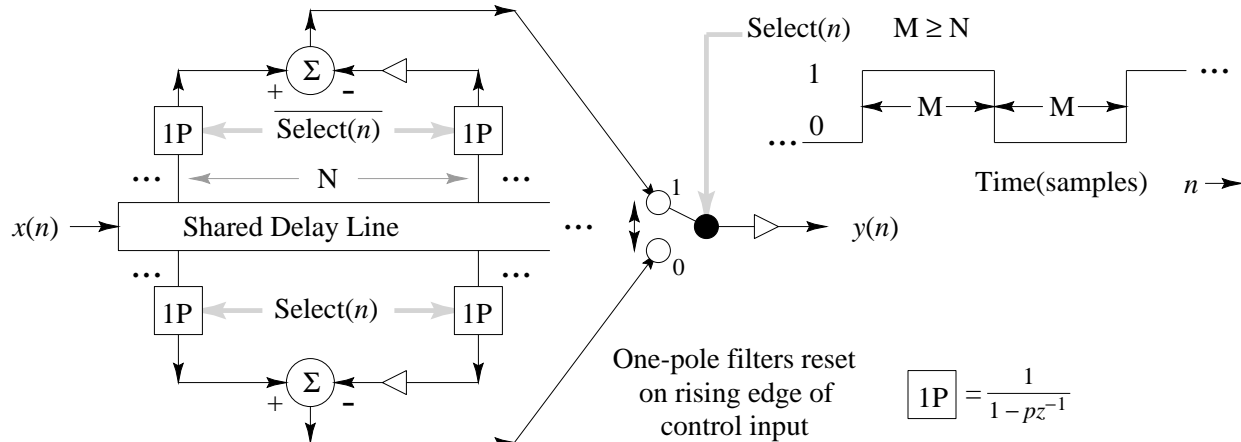


Figure 1: Example of a TIIR filter for generating a growing exponential or constant segment.

Referring to Fig. 1, suppose the upper pair of one-pole filters is switched in (as the figure indicates). When the **Select** signal transitions, the alternate one-pole pair below is selected, and the upper one-poles can be cleared and halted (or simply not computed in a software implementation). If the TIIR impulse-response length is N samples, then the first upper filter on the left is restarted N samples before it is to be switched back in, while the second upper filter is restarted on the *same* sample as when it is switched back in. This works because, even though the upper pair will not be in the same state as the lower pair after N times steps, its tail-canceling difference, which synthesizes an FIR filter, *is* identical (ignoring round-off errors). Therefore, the switching resets can be as often as every N samples. It is desirable, however, to switch much less often than every N samples in order to minimize computations. The minimum switching rate, at the other extreme, is determined by the exponential growth rate and available dynamic range [10].

Note that the multiply-add which forms the tail-canceling subtraction can be shared since only the output of the actively selected branch is needed.

Finally, we note that when the structure of Fig. 1 is used to implement a truncated constant impulse response, the one-poles become digital integrators (no multiplies), and the tail-canceling multiply-subtract becomes only a subtraction. The resets for digital integrators can be considerably less often than for growing exponentials, because the round-off error grows more slowly in an integrator [10].

In summary, a TIIR filter for making a truncated constant or rising exponential impulse response segment can be computed at a cost close to that of a one-pole filter and a multiply-add, plus some

associated switching and control logic.

3 Horn Reflectance Filter

A general characteristic of musically useful horns is that their internal bore profile is well approximated with a Bessel horn [2]. Although any real instrument bell will show significant deviations from this approximation in its bore shape and acoustic reflectance, a theoretically derived Bessel horn reflection function may serve as a suitable generalized target-response for developing effective digital filter design techniques. In order to obtain such a target-response, the pressure reflectance of a Bessel horn that approximates the shape of a trumpet bell was computed as in [9].

As shown in Fig. 2, the Bessel horn reflection impulse response has a slow, quasi-exponentially growing portion at the beginning, corresponding to the smoothly increasing taper angle of the horn. A one-pole TIIR filter gives a truncated exponential impulse response $y(n) = ae^{cn}$, for $n = 0, 1, 2, \dots, N - 1$, and zero afterwards. We can use this truncated exponential to efficiently implement the initial growing trend in the horn response ($c > 0$). We found empirically that improved accuracy is obtained by using the *sum* of an exponential and a *constant*, i.e.,

$$y(n) = \begin{cases} ae^{cn} + b, & n = 0, 1, 2, \dots, N - 1 \\ 0, & \text{otherwise} \end{cases}$$

The truncated constant b can also be generated using a one-pole TIIR filter, with its pole set to $z = 1$. In this case, no multiplies are needed, except for the single scale factor b . The transfer function of the TIIR filter for modeling a single segment of the horn impulse response as an offset exponential can be written as

$$H(z) = h_0 \frac{1 - p^{N+1} z^{-(N+1)}}{1 - pz^{-1}} + b \frac{1 - z^{-(N+1)}}{1 - z^{-1}}. \quad (1)$$

The remaining reflection impulse response has a decaying trend, and can therefore be modeled accurately with diverse conventional filter design techniques. Here, the Steiglitz-McBride IIR filter design algorithm was applied [3].

In Fig. 2, the TIIR horn filter structure (using a 3rd-order IIR tail filter approximation) is compared with the theoretical response. The *phase delay* (directly proportional to the “effective length” of the bell for standing waves), has a particularly good fit, which is important for accurate musical resonance frequencies of a brass instrument.

4 Application to the Trumpet Using Empirically Derived Data

Acoustic pulse reflectometry techniques [6] were applied to obtain the impulse response of a trumpet (without mouthpiece). A piecewise cylindrical section model of the bore profile was reconstructed using an inverse-scattering method [1], taking into account the viscothermal losses (see Fig. 4). The piecewise cylindrical model corresponds well to the physical bore profile for non-flaring tube-segments, thus giving a good physical model up to the bell. The remaining cylindrical sections do not provide valid geometrical information, but they retain all relevant acoustical information of the bell reflectance, including the complex effects of higher transversal modes and radiation impedance.




Figure 2: Bessel horn response (solid) compared with digital filter approximation (dashed) in terms of impulse response (a), magnitude (b) (up to Nyquist) and phase delay (c) (up to bell cut-off). The vertical line in (a) indicates the segmentation into an growing exponential and a decaying tail.

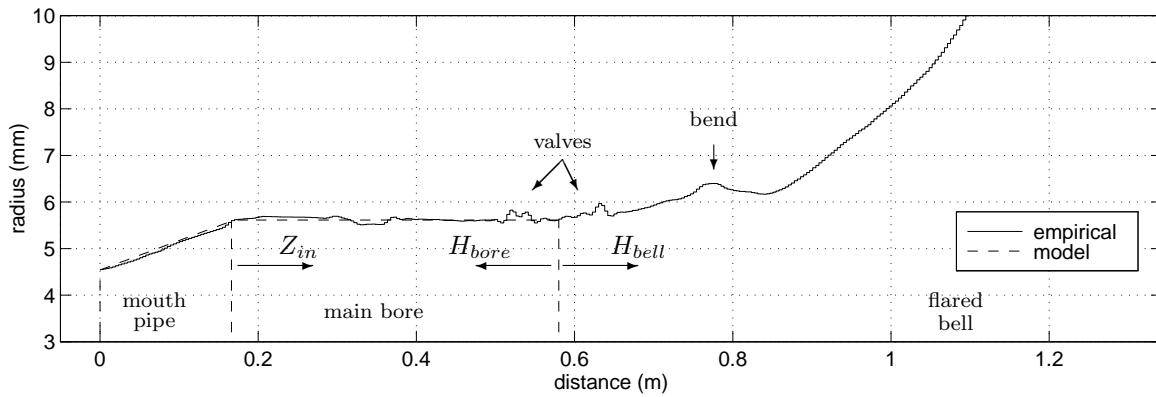


Figure 3: Trumpet bore profile reconstruction. The valves and the final tubular bend show as 'dents' in the profile. The main bore plus mouthpipe can be modeled with a cylindrical section preceded by a truncated cone (dashed lines).

The main bore of a trumpet is essentially cylindrical, with an initial taper widening (mouthpipe) (see Fig. 3). Thus, an accurate digital waveguide model of the trumpet can be derived by approximating the bore profile data with a cylindrical bore, plus a conical section to model the mouthpipe, and modeling the remaining part of the reconstruction as the isolated bell reflectance $H_{bell}(\omega)$. The complexity of the model can be further reduced by lumping the viscothermal losses of the main bore with the bell reflectance filter, yielding the “round-trip filter” $H_{rt}(\omega)$:

$$H_{rt}(\omega) = \frac{H_{bore}(\omega)}{H'_{bore}(\omega)} \cdot H_{bell}(\omega), \quad (2)$$

where $H_{bore}(\omega)$ represents the response “seen” from the bell (see Fig. 3) while assuming an ideal closed end at the junction between the mouthpipe and the main bore, and $H'_{bore}(\omega)$ is the theoretical value of $H_{bore}(\omega)$ assuming no losses. The inverse Fourier transform $h_{rt}(t)$ differs from the theoretical Bessel horn response primarily in its two-stage build-up towards the primary reflection peak (see Fig. 4). This characteristic was observed for a variety of brass instruments. By adding another offset-exponential TIIR section (Eq. (1)) to the basic horn filter structure, the filter design methodology is sufficiently flexible to cover the two-stage build-up. The resulting impulse response and corresponding input impedance curve $Z_{in}(\omega)$ (“seen” from the start of the main bore) are depicted in Fig. 4. The small amplitude deviations are mainly due to the fact that the TIIR approximation of the initial slow rise is insensitive to reflections caused by bore profile dents. Note that the resonance frequencies, controlled by the phase delay of $H_{rt}(\omega)$ are accurately modeled.




Figure 4: Round-trip filter (a) and “main-bore” input impedance (b) according to empirical data (dashed) compared to TIIR horn filter (solid). The vertical (dash-dot) lines in (a) indicate the response segmentation into 2 growing exponentials and a tail. The tail is modeled with a 4th-order IIR filter.

5 Piecewise Conical Bore Modeling

It is well known that a growing exponential appears when waves traveling within one conical taper angle reflect from a section with a smaller (or more negative) taper angle [4]. This phenomenon has precluded the use of a straightforward recursive filter model [5, 8] since such a filter would have to be unstable. However, using TIIR principles, it is possible to use unstable digital filters in this way while resolving practical difficulties.

The main difference in the piecewise conical modeling case is that conical segments are not strictly FIR. However, in practical musical acoustics, they have quite short decay times. Therefore, we may apply TIIR principles with t_{60} replacing the FIR filter length in determining the maximum switching rate, where t_{60} is the time for the *external* impulse response of the model component to decay by 60 dB. Note that the dc response must also decay to insignificance by t_{60} .

6 Conclusions

We have presented a computationally efficient modeling framework applicable to flaring horns and piecewise conical bores. The horn models use tail-canceling IIR filters to implement finite exponential and constant impulse response segments, with periodic replacement of unstable filter components used to avoid indefinite build-up of round-off errors. The piecewise conical models follow analogous principles with t_{60} replacing the FIR filter length. Compared with previous practical approaches to modeling these musical acoustic elements computationally, the TIIR approach offers compelling advantages.

References

- [1] N. Amir, G. Rosenhouse, and U. Shimony, “Discrete model for tubular acoustic systems with varying cross section—the direct and inverse problems. Part 1: Theory,” *Acta Acustica*, vol. 81, pp. 450–462, 1995.
- [2] A. H. Benade and E. V. Jansson, “On plane and spherical waves in horns with nonuniform flare. I. theory of radiation, resonance frequencies, and mode conversion,” *Acustica*, vol. 31, no. 2, pp. 80–98, 1974.
- [3] L. Ljung and T. L. Soderstrom, “The Steiglitz-McBride algorithm revisited—convergence analysis and accuracy aspects,” *IEEE Transactions on Automatic Control*, vol. 26, pp. 712–717, June 1981, See also the function `stmcb()` in the Matlab Signal Processing Toolbox.
- [4] J. Martinez and J. Agullo, “Conical bores. Part I: Reflection functions associated with discontinuities,” *Journal of the Acoustical Society of America*, vol. 84, pp. 1613–1619, Nov. 1988.
- [5] G. P. Scavone, *An Acoustic Analysis of Single-Reed Woodwind Instruments with an Emphasis on Design and Performance Issues and Digital Waveguide Modeling Techniques*, PhD thesis, CCRMA, Music Dept., Stanford University, March 1997, <http://ccrma.stanford.edu/~gary/>.
- [6] D. B. Sharp, *Acoustic Pulse Reflectometry for the Measurement of Musical Wind Instruments*, PhD thesis, Dept. of Physics and Astronomy, University of Edinburgh, 1996, <http://acoustics.open.ac.uk/pdf/thesis.pdf>.
- [7] J. O. Smith, “Physical modeling using digital waveguides,” *Computer Music Journal*, vol. 16, pp. 74–91, Winter 1992, special issue: Physical Modeling of Musical Instruments, Part I. <http://ccrma.stanford.edu/~jos/pmudw/>.

- [8] V. Välimäki and M. Karjalainen, “Digital waveguide modeling of wind instrument bores constructed of truncated cones,” in *Proceedings of the 1994 International Computer Music Conference, Århus*, pp. 423–430, Computer Music Association, 1994.
- [9] M. van Walstijn and V. Välimäki, “Digital waveguide modeling of flared acoustical tubes,” in *Proceedings of the 1997 International Computer Music Conference, Greece*, (Thessaloniki, Greece), pp. 196–199, Computer Music Association, 1997.
- [10] A. Wang and J. O. Smith, “On fast FIR filters implemented as tail-canceling IIR filters,” *IEEE Transactions on Signal Processing*, vol. 45, pp. 1415–1427, June 1997.