Astron. Astrophys. Suppl. Ser. **124,** 183-188 (1997)

Recovering line profiles from frequency-switched spectra

H. Liszt

National Radio Astronomy Observatory, 520 Edgemont Road, Charlottesville, VA 22903-2475, U.S.A.

Received July 29; accepted September 25, 1996

Abstract. The usual shift and subtract algorithm for recovering line profiles from frequency-switched data succeeds only when the switching interval exceeds the extent of the signal-bearing region. Nonetheless, noiseless spectra can in principle be reconstructed without error for any frequency-switching interval of at least one channel. This work presents a simple recursive algorithm which recovers line profiles from spectra which might otherwise appear hopelessly damaged, with only a slight cost in signal/noise. It is functionally equivalent to the dual-beam restoration algorithm of Emerson et al. (1979, A&A76, 92) which is commonly used to map objects bigger than a telescope's beam throw.

Key words: techniques — image processing — methods: observational — methods: data analysis

1. Introduction

Frequency-switching is a commonly used technique in spectral line observing. As illustrated in Fig. 1, a spectrum is accumulated in two interleaved phases, one of which is inverted and shifted with respect to the other. If the signal is in emission, the upward-pointing phase is commonly described as the "signal" while the other is called the "reference". The signal and reference phases are observed sequentially at moderate switching rates (typically 1 Hz) and combined to form an output spectrum (Fig. 1 at bottom) which has shifted and oppositely- directed representations of the incoming profile.

The advantages of frequency switching are that the 1 Hz rate is rapid with respect to variations in the hardware and environment, which are partly cancelled when the two phases are combined; that there is no need to find a signal-free reference position elsewhere on the sky; and that the source is observed without interruption. Its main defect is an often-poor output passband shape which prevents detection of broad lines. If a spectral shape is impressed on the signal and reference phases, shifting and differencing them will cause something resembling a crude numerical derivative to appear in the output. If the spectral shape is of zero or first order, differencing will remove it. If the spectral shape is discontinuous or complicated, the output frequency-switched spectrum may be hopelessly compromised.

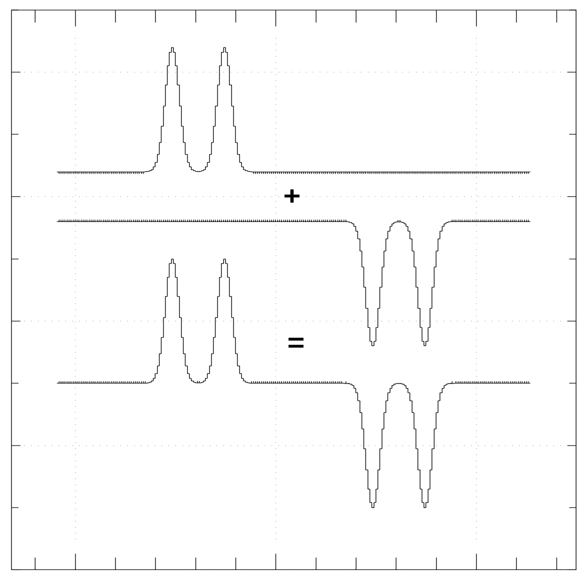

Fig. 1. Construction of a frequency-switched spectrum. A spectrum (top) observed in the signal phase is shifted and inverted (middle) in the reference phase. The signal and reference are added (bottom) to produce the "observed" spectrum

Proper data-handling in extant reduction algorithms requires that broader lines be observed with wider frequency-switching intervals. Such wider switching further degrades the output passband, often rendering the frequency-switching technique inappropriate for all but the narrowest and simplest spectra. This is a significant hindrance when emission-free regions of the sky are either unknown, difficult to find, or very distant.

The standard algorithm for recovering the signal from a frequency-switched spectrum like that at the bottom of Fig. 1 is simply to shift by the known frequency-switching interval and subtract, dividing by a factor 2, as shown schematically in Fig. 2. When both the signal and reference phases contain a good representation of the complete line profile, with adequate signal-free regions between and to either side, the output of the shift and subtract algorithm is an error-free reconstruction of the line profile (Fig. 2 at bottom).

Occasionally, the frequency-switching interval is not chosen so felicitously. In Fig. 3, the interval accidentally coincides with part of the structure of the line in such a way that the output spectrum (at bottom) is corrupted. In this particular case, use of the shift and subtract algorithm to recover the profile produces a result which has the correct shape but only half the correct amplitude. This example also illustrates the fact that the frequency-switching interval cannot necessarily be deduced from the appearance of a frequency-switched spectrum. There is a common aliasing problem such that a given output frequency-switched spectrum can often be reproduced by various combinations of signal shape and switching interval.

The failure of the standard algorithm to recover the proper spectrum from data like that in Fig. 3 is unfortunate because the output frequency-switched spectrum at bottom clearly contains a faithful copy of every signal-bearing channel; the component on the left appears in the signal phase while that on the right appears in the reference. In fact, noiseless spectra are identically recoverable for *any* frequency- switching interval of at least one channel. No matter how complicated or damaged they may appear, it should not be necessary to abandon even noisy, imperfect data simply for lack of a competent algorithm. The essentials of this matter were recognized some time ago by Emerson et al. (1979: EKH) in considering how to map spatially extended objects with a small telescope beam throw. The solution we propose for spectral line work is functionally equivalent to the EKH algorithm, although this may not be apparent at first glance.

Section 2 discusses several methods of recovering frequency-switched spectra. The effect of each one, operating on real-world data, is shown in the next- to-last figure. Section 3 briefly discusses a generalization to the case when the frequency-switching interval is not a whole number of channels, using the formalism of EKH.

2. Algorithms

2.1. Frequency-switching viewed as a convolution

Ideally, frequency-switching may be described as a process of convolution; a pattern $c(x)$ disposed as at the top of Fig. 4 can be convolved with an array to reproduce the appearance of a frequency-switched spectrum. For some but not all choices of the frequency-switching interval, the original spectrum may be recovered by a naive linear deconvolution in the ideal case of noise-free data. That is, the convolution of $c(x)$ and a spectrum $f(x)$ is the observed spectrum $g(x) = c(x) * f(x)$. If the Fourier transform of

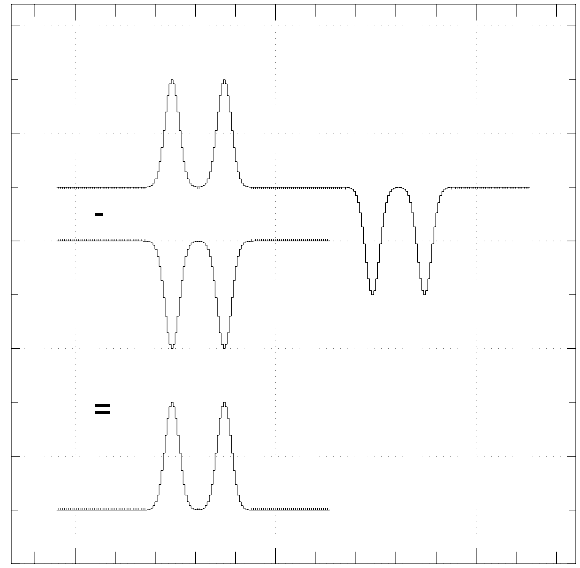

Fig. 2. Reconstruction of a frequency-switched spectrum by the standard algorithm. The observed spectrum (top) is shifted (middle) and differenced with itself, and divided by 2 to recreate the original signal

a function is denoted by $F()$, the deconvolution theorem states that $F(g(x)) = F(c(x)) \times F(f(x))$ and $F(f(x)) = F(g(x))/F(c(x))$. As long as no Fourier components of $c(x)$ have 0-modulus, $f(x)$ can be recovered exactly from noiseless data simply by doing an inverse Fourier transform. As shown by EKH $F(c(x)) \propto \sin(\pi\delta x/x)$ and many choices for the frequency-switching interval δx yield zeros in the components of $F(c(x))$ when a simple FFT is performed, including but by no means limited to channel shifts of 2^n $(n > 1)$ channels.

In practice the naive linear deconvolution is nearly useless with real data. Because the signal and reference phases are accumulated independently, neither their noise nor their baseline shapes, etc. are exact shifted copies of one another: the data only approximate a convolution. Even more importantly, a one-sided function like the $c(x)$ shown at the top Fig. 4 is grossly insensitive to some Fourier components even when it has no exact zeros. This is manifested by the many small moduli in the periodogram power spectrum at the bottom of that figure and is not rectified by available alternatives such as switching symmetrically to either side. Unless noise in the frequency-switched spectrum $g(x)$ is coincidentally missing or heavily filtered at these frequencies, a simple linear deconvolution will amplify it greatly.

Some degree of selective attenuation of the Fourier components of the incoming signal is inevitable with frequency-switching schemes. One insight of EKH was to show that components of $F(c(x))$ may be selected such that a region of appreciable extent can be recovered without unduly increasing the noise through deconvolution.

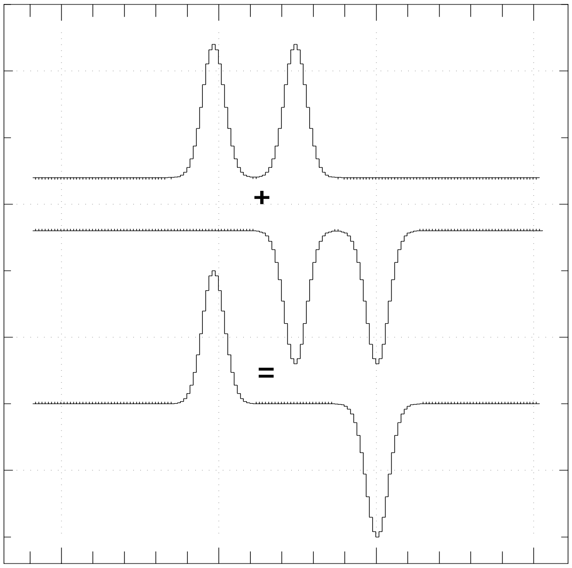

Fig. 3. As in Fig. 1, but for a pathological case. The resulting spectrum (bottom) does not reproduce the original signal when presented to the algorithm described in Fig. 2

Note that the aliasing referred to at the end of Sect. 1 (see Fig. 3) is equivalent to linear deconvolution with improper switching intervals $\delta x' = \delta x/n, n = 2, 3...$ etc.

2.2. The conventional shift and subtract technique

Symbolically we represent a spectrum by $f(x)$. The process of creating a frequency-switched spectrum involves forming the difference $g(x) = f(x) - f(x - \delta x)$ where δx is the frequency-switching interval (most telescopes actually specify offsets for the signal and reference phases individually; δx is their separation). The standard algorithm for creating an approximation $f'(x)$ to the unknown $f(x)$ from the data is (see Fig. 2)

$$g(x) = f(x) - f(x - \delta x) \tag{1a}$$

$$f'(x) = \frac{g(x) - g(x + \delta x)}{2}. \tag{1b}$$

Substituting 1a into 1b, one finds

$$f'(x) = f(x) - \frac{f(x - \delta x) + f(x + \delta x)}{2} \tag{2}$$

which shows that an arbitrary spectrum is recovered at position x only when δx is large enough that $f(x \pm \delta x) = 0$. When this the case, the error in $f'(x)$ is smaller than the noise in $g(x)$ by a factor $1/\sqrt{2}$. Switching within the signal region, as in Fig. 3, creates insurmountable problems for this algorithm.

2.3. A recursive, bootstrap technique

To ameliorate the difficulties caused by switching within the line, imagine applying a different algorithm to the spectrum at the bottom of Fig. 3, i.e. step to the right in the spectrum calculating

$$f_{\rm s}(x) = g(x) + f_{\rm s}(x - \delta x). \tag{3a}$$

This recursive approach, which simply adds back to each channel a close-at-hand estimate of what should have been subtracted from it, identically recovers the signal phase of the spectrum from the $g(x)$ at the bottom in Fig. 3. Similarly, a leftward progression calculating

$$f_{\rm r}(x) = g(x) - f_{\rm r}(x + \delta x) \tag{3b}$$

recovers the reference phase, whence

$$f'(x) = \frac{f_{\rm s}(x) - f_{\rm r}(x + \delta x)}{2}. \tag{3c}$$

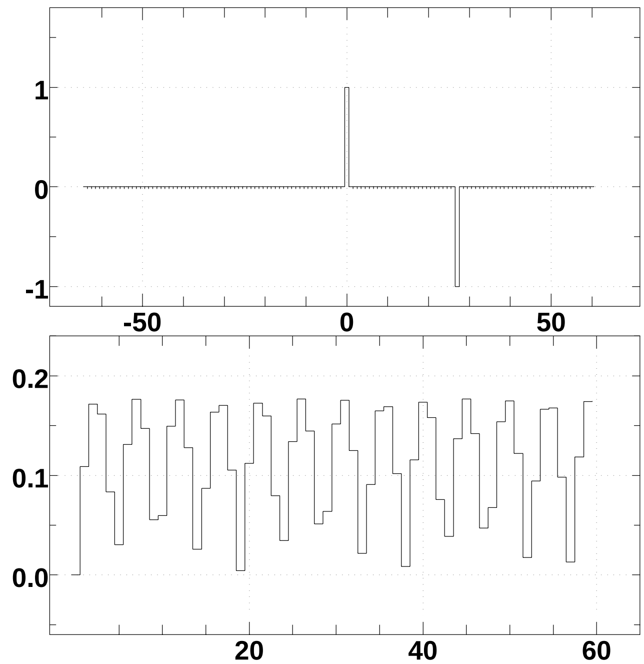

Fig. 4. Frequency switching viewed as a convolution. At top, a typical convolution function $c(x)$ for frequency-switching (the interval is 27 channels). A periodogram power spectrum of this function appears at bottom showing the amplitudes (not their squares) of its Fourier components. Many of these are quite small, making it hard to do deconvolution on noisy data. Many choices for $c(x)$, such as intervals of $2^n (n > 1)$ channels, have actual zeroes in their power spectra, making linear deconvolution impossible. Some Fourier components of the incoming signal are heavily attenuated by the switching process

In fact this procedure is equivalent to the dual beam restoration algorithm of EKH, although such may not be entirely obvious from a casual reading of the earlier work. EKH show that it is possible to do a worthwhile selective recovery of the signal using only those Fourier components which are not attenuated by more than a factor of two in

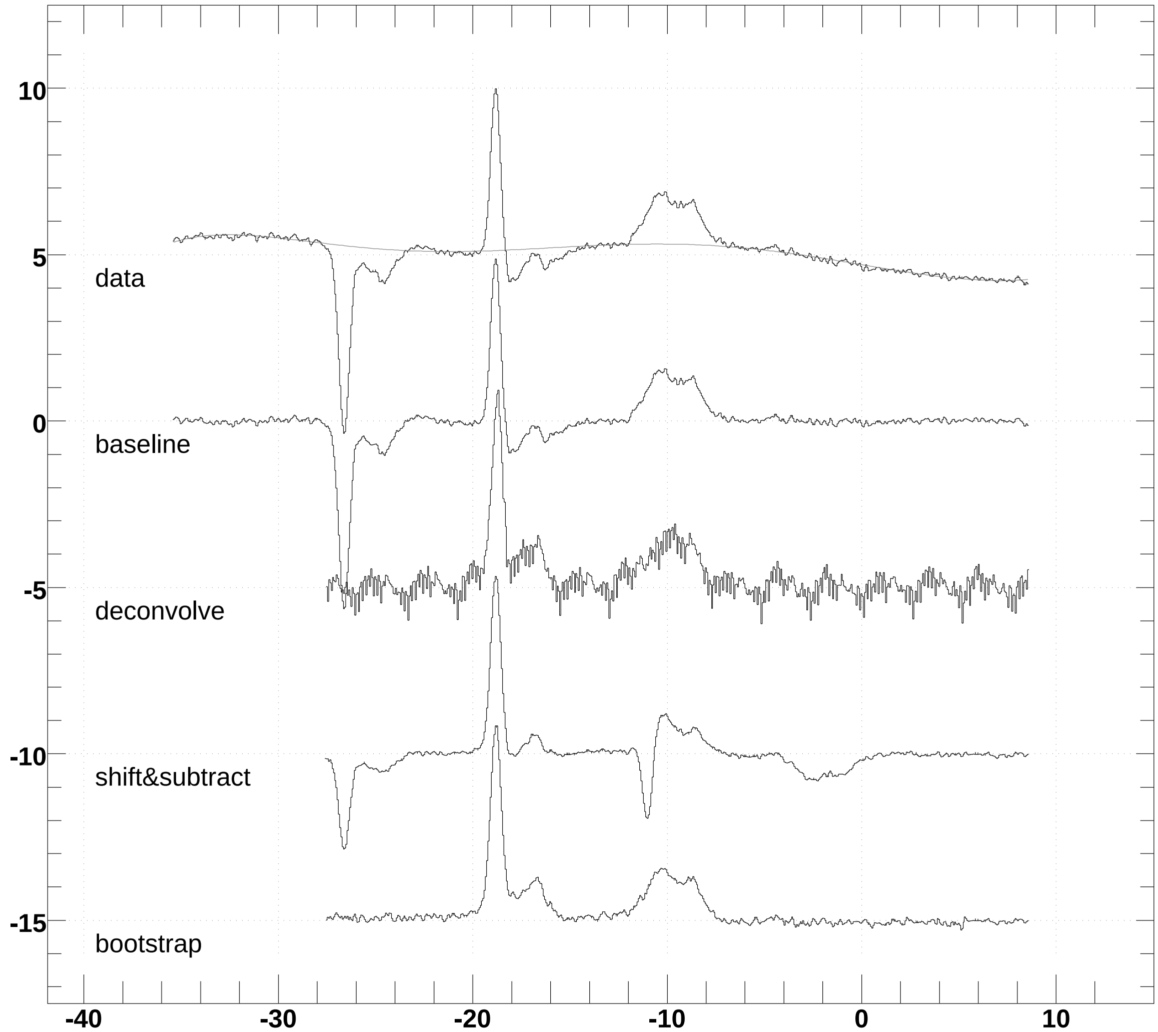

Fig. 5. Treatment of actual frequency-switched data. At top is a CO $J = 2-1$ spectrum observed with a 122.88 channel frequency-switching interval (the strongest feature is telluric). A linear baseline and standing-wave have been removed to create the "baseline" spectrum and this baselined spectrum is shown as processed by naive linear deconvolution, the convential shift and subtract, and by the bootstrap algorithm described here

switching. They then show a) that this selection is equivalent to requiring six samples per cycle of $F(c(x))$; b) that a region of extent $3\delta x$ can be recovered; and c) that the deconvolution may be implemented without Fourier transforms simply by convolving $g(x)$ with a simple comb (their Fig. 6iii; see our Fig. 6 and Sect. 3.4 here as well). In the caption for their Fig. 6, it is noted that using only half the comb is eqivalent to our Eqs. (3a) or (3b), from which it follows that use of all of it is equivalent to our Eq. (3c).

The non-recursive approach in Eq. (1) should of course be used whereever possible. However the recursive bootstrap is worthwhile even when it is not absolutely needed because it is an efficient means of testing whether the incoming signal was inadvertantly corrupted by improper frequency- switching. It produces not only a faithful representation of the more obvious signal-bearing channels, but a cleaner view of the passband as well (i.e. the presumed signal-free channels on either side of the line), which could in principle harbor signals which otherwise would be missed in the process of shifting and subtracting.

The recursive approach does not necessarily achieve a noise level corresponding to the full integration time because the signal and reference phases may not be fully present in the spectra upon which it works. There are more subtractions, some of which increase the noise (Eqs. (3a) and (3b)) because they are not compensated by averaging which only occurs in Eq. (3c). Unlike most conventional algorithms, however, this one produces a noise level which may be lower over its representation of the signal-bearing region than over the baseline regions. If the support of the signal and reference phases is apparent, the procedures implied by Eqs. (3a) and (3b) can be limited to that region.

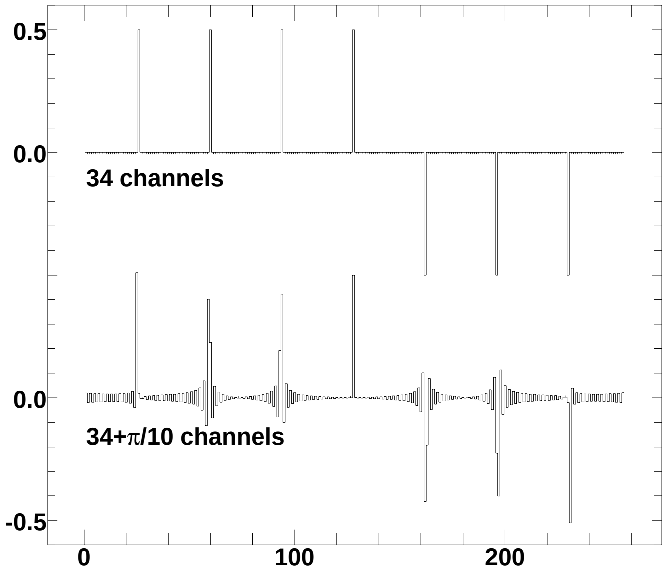

Fig. 6. Unfolding functions. The convolution of these functions with frequency-switched spectra implements the bootstrap in the mannner discussed by Emerson et al. (1979) for shifts of 34 (top) and $34 + \pi/10$ channels. There are 256 channels and the central spike is the 128th channel from the left in either case. The phase of the comb has been chosen so that the recovered spectrum appears in channels occupied by the so-called signal phase of the unfolded data

If the switching interval is very narrow compared to the signal-bearing region, use of recursion in traversing a region of extent x_0 in steps of δx will degrade the noise by a factor $\approx (x_0/\delta x)^{1/2}$ (also see EKH). When $x_0/\delta x \approx 1$, noise in the output spectrum of the recursive algorithm varies in broad contiguous patches across the band, which we explored by doing Monte Carlo simulations of frequency-switching a broad, flat-topped profile. For a profile with noise ΔT which would have been well-recovered by the shift and subtract technique (resulting in a uniform noise level $\Delta T/\sqrt{2}$) the bootstrap had differing noise levels in four regimes; $\Delta T/\sqrt{2}$ over the region of the line; between ΔT and $\Delta T/\sqrt{2}$ in the adjacent baseline region; $\sqrt{2}\Delta T$ in that portion of the output passband where the reference phase was originally present; and ΔT elsewhere. For a profile switched within itself by half the width, the resultant noise level was ΔT over the region of the line and the adjacent passband, and $\Delta T/\sqrt{2}$ elsewhere.

It is very important to recognize that the recursion can introduce serious artifacts from imperfections such as bad baselines. On the other hand, the use of narrower frequency switching intervals, made possible by the use of the recursive algorithm, may do much to eliminate the root cause of such problems.

2.4. An example using real data

Figure 5 illustrates the processes discussed here. At top is a 30-minute frequency-switched integration on the CO($J = 2-1$) line toward the continuum source B0355+508; a profile of the $J = 1-0$ line of CO appears in Liszt & Wilson (1993). The channel interval of the switch was 122.88 channels or 7.8 km s^{-1} in units along the x-axis. The profile was baselined by removing a first-order binomial and a long-period standing-wave (shown superposed at the top) and the baselined profile was processed in three ways: by linear deconvolution using FFT's, by the usual shift and subtraction, and by the bootstrap recursive solution discussed above. Clearly only the bootstrap produces a reasonable profile without substantial further intervention although its noise is noticeably larger than that in the (rather useless) spectrum directly above it.

3. Dealing with fractional channel switching

The bootstrap implementation of the EKH algorithm is a simple and intuitive way of understanding what actually occurs but with the obvious limitation that it is exact only when the amount of the frequency switch is a whole number of channels. The generalization of the bootstrap to non-whole numbers of channels follows trivially after expressing its actions in terms of a convolution in the manner of the original discussion of EKH. As a function of channel number $i (i = 0 \ldots i_>)$ the convolving function for a switch of s channels is

$$E(i) = 0.5 \sum_{m=-N}^{N+1} \mathrm{sign}(2m-1)\ \mathrm{Sinc}(i-(i_>\ \mathrm{div}\ 2)-ms) \quad (4)$$

where $N = 1+\mathrm{int}(i_>/s)$, $\mathrm{Sinc}(x) = \sin(\pi x)/(\pi x)$, and the div operation signifies integer division discarding any remainder. If the convolution is implemented with an FFT, the spectrum should be well padded with zeroes at either end.

Figure 6 shows convolution functions for the cases of 34 and $34 + \pi/10$ channels. The delta functions of the whole number case at top have a clearer ringing behaviour at bottom; the sinc function is used to interpolate band-limited, critically-sampled functions and the action of convolution with a single off-center sinc is simply to translate by a non-integer number of channels. The phase of the comb in Fig. 6 has been chosen so that the recovered signal will appear in the same channels as are occupied by the signal phase of the switching cycle. If both the signal and reference phases are symmetrically offset and the recovered signal should appear between them, the term $-ms$ in Eq. (4) can be replaced by $-(m + 1/2)s$.

Note that the simple existence of a fractional channel switching interval does not by itself mandate the use of the bootstrap. The ensuing sacrifice in noise level (a factor of at least $\sqrt{2}$ if the simple shift and subtract technique can be used) causes larger statistical uncertainties in the estimated line properties, quite possibly exceeding the error introduced by use of a slightly inaccurate switching interval. When features are well-resolved, whole-channel shifts introduce at most only small errors. The convolution implementation of the bootstrap runs much more more slowly than the whole-channel algorithm expressed in Eq. (3).

4. Summary

We considered several algorithms for recovering data from frequency-switched spectra. Naive linear deconvolution, a possibility for some but not all switching intervals, is rendered useless by the noise and baseline imperfections of real-world data (Sect. 2.1). The usual shift and subtract technique (2.2) treats baseline imperfections most benignly (though not necessarily well) and results in the lowest possible noise levels; when the data permit, this is the algorithm of choice. The bootstrap or EKH techniques discussed here (Sects. 2.3 and 2.4) can (*i*) recover line profiles for which the switching interval is too small for the shift and subtract algorithm to function; (*ii*) account as precisely as possible for fractional channel switching intervals; and (*iii*) check that broader-lined signals have not been lost. Though sensitive to baseline imperfections, they are capable of repairing spectra which might otherwise appear badly damaged or useless.

Acknowledgements. The National Radio Astronomy Observatory is a facility of the National Science Foundation, operated by AUI, Inc. under a cooperative agreement. I thank Darrel Emerson for pointing out the equivalence of the bootstrap and dual-beam restoration algorithms. I also thank the referee, Glyn Haslam, for his bemused, approving remarks.

References

Emerson D.T., Klein U., Haslam C.G.T., 1979, A&A 76, 92 (EKH)

Liszt H.S., Wilson R.W., 1993, ApJ 403, 663