DOPPLER BANDWIDTH CHARACTERIZATION OF ARTM CHANNEL SOUNDING DATA

David Landon
Department of Electrical and Computer Engineering
Brigham Young University

Faculty Advisor:
Michael Rice

ABSTRACT

Data from ARTM channel sounding test flights is examined to characterize the dynamic channel behavior of aeronautical telemetry channels. The dynamic behavior is characterized using the Doppler power spectrum. The width of the Doppler power spectrum is the Doppler bandwidth of the channel which indicates the required bandwidth of adaptive detection techniques such as adaptive equalization, adaptive modulation, adaptive channel selection and adaptive error control coding. Data collected from ARTM Flight 11 suggest a Doppler bandwidth exceeding 6.7 Hz for the channel, but greater accuracy and resolution will only be possible with more data.

KEY WORDS

Doppler Bandwidth, Doppler Spread, Channel Characterization, Multipath Fading, ARTM.

INTRODUCTION

Multipath fading is a common telemetry problem, which may vary substantially with time for any given channel. As the transmitter or receiver move relative to each other, the types of terrain or individual scatterers which separate them often change. As a result, frequency fades may be stronger or more common at certain times than at others.

A common solution to the problem of channel distortion is to estimate the worst-case channel distortion and design a channel equalization filter to undo the effects of the distortion. Clearly, if the channel changes over time, we would be undoing a distortion that is not even present. That may distort the channel and make it worse than doing nothing at all. The typical approach, then, for a time-varying channel is to use a filter whose coefficients are updated at some appropriate interval.

Adaptive algorithms form a whole class of techniques which can be applied to telemetering problems. Adaptive channel equalization is just one example. As a further example, adaptive modulation schemes can vary the number and position of constellation points, which forms the basis of modem operation as they negotiate the optimal baud rate at a given instant. The key principle of adaptive algorithms is that they are not designed to work by brute force during one set of worst-case conditions. By updating channel model parameters at an appropriate interval they react to the actual channel conditions rather than initial conditions, which have long since changed.

But that all just raises one more very important question: what is an "appropriate interval?" How often should the channel be sensed and the models updated? That is the essence of the Doppler bandwidth parameter B_d and its reciprocal, $(\Delta t)_C$, the coherence time of the channel. Basically, the coherence time is a measure of how long the channel can be viewed as though it did not change over time. A long coherence time would mean, that the appropriate update interval would be rather long. Likewise, a short coherence time would mean that the appropriate update interval would be quite short. Algorithms designed to adaptively compensate for channel conditions, should be set to update at least as often as B_d times per second.

This paper examines the standard approach of quantifying the channel dynamics and applies this technique to the ARTM channel sounding data.

DETERMINATION OF THE DOPPLER BANDWIDTH FOR A RAPIDLY CHANGING CHANNEL

The Doppler bandwidth is defined as the range of values over which the Doppler power spectrum is nonzero [1]. As an example of the procedure for computing the Doppler power spectrum, consider the time sequence of 10 consecutive power spectra illustrated in Figure 1a. The ten spectra are plotted one in front of another and each represents a snapshot of the channel response at 150 msec time intervals during ARTM Flight 11. Each one is quite different from at least one of its neighbors suggesting that our snapshot sample rate is too slow to capture all of the channel behavior. The "snapshot sample rate" will be used in this paper to describe the rate at which snapshots are taken. Thus for the plots in this paper, the snapshot sample rate is measured in the time variable t, not the delay variable t, which is tied to the sample rate used to collect the data (100 Msamples/second for Flights 10 and 11). It is clear that there is a null in the front frequency response near +2 MHz relative to the carrier frequency. Though somewhat concealed, in the fourth, seventh and tenth plots, there are also nulls which sweep from right to left as you move back in the cascade of plots. The other plots in the ten-plot sequence have negligible fading. This indicates that during the first, fourth, seventh and tenth snapshots, there was a strong (-30 dB) frequency-selective fade near the carrier frequency. The frequency-selective fading lessens and almost disappears altogether

within another 150 msec. Thus, as the plane progressed through its flight the nature of multipath distortion changes drastically with time.

Normally, we assume that the channel has a specific frequency response or transfer function $H(f)$ and the power spectrum corresponds to $|H(f)|^2$, but that would be incomplete here. Figure 1a clearly shows that the frequency response of the channel changes with time, so it makes more sense to consider the plot in Figure 1a as a plot of $|H(f,t)|^2$, where $H(f,t)$ is a transfer function reflecting the channel's time dependence. It is exactly this time dependence of the channel that we will explore.

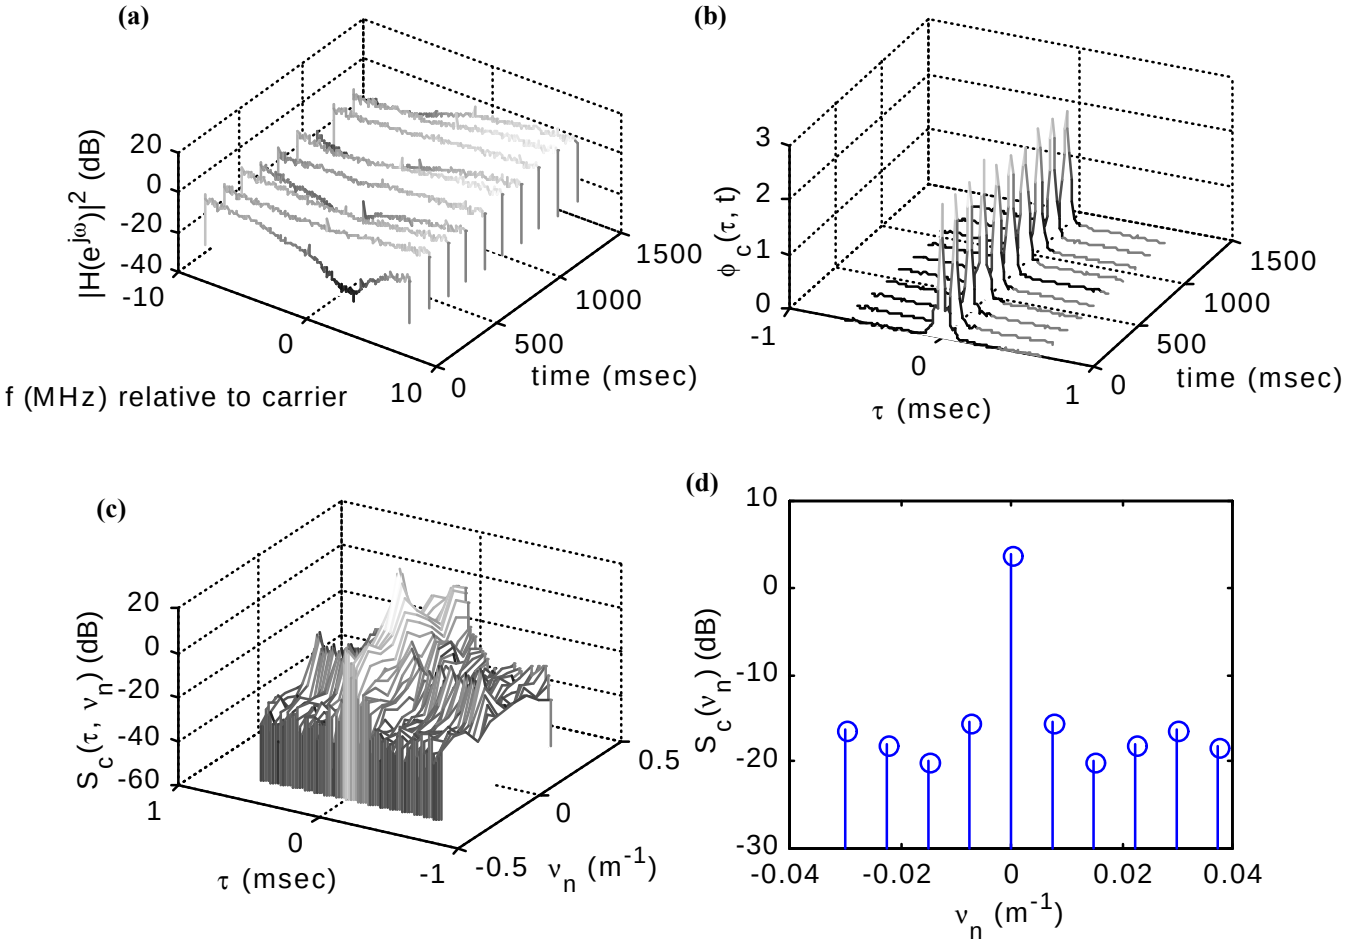

Figure 1 Data segment from ARTM Flight 11 starting at 18:53:38 showing its (a) magnitude squared channel transfer function, $|H(f,t)|^2$, (b) delay-correlation function, $R(t,t)$, (c) scattering function $S_C(t,n_n)$ and (d) Doppler power spectrum, $S_C(n_n)$.

We proceed by taking an inverse Fourier transform of the data with respect to our frequency variable and plot the result in Figure 1b. The inverse Fourier transform of the time-varying power spectrum $|H(f,t)|^2$ is the time-varying correlation function R(t,t), i.e.:

$$|H(f,t)|^2 \xleftarrow{F^{-1}} R(t,t) \tag{1}$$

where the correlation is in t, the delay variable [2]. The ten plots of Figure 1b appear much more similar to each other than the plots of Figure 1a do. Each looks like a narrow peak with some small unevenness at the base and then is virtually zero over the rest of the range. The unevenness at the bases of the peaks are actually the result of multipath with delays too small to be resolved using this technique [3]. (The delays of a couple of other data sets are quite resolvable as can be seen in Figure 3b. For a discussion of the multipath characterization ARTM Flights 10 and 11 see [3].) The plots in Figure 1b show that the data transmitted through the channel is uncorrelated with delayed versions of itself except at very small delays. Ideally it should be uncorrelated with itself except at a delay of zero, because the transmitted signal is a pseudo noise (PN) sequence. Because of the multipath distortion, however, one or more copies of the signal arrive at some small delay after the signal first arrives by way of the line of sight (LOS) path and widens the bases of the peaks.

Though less variation is visible in Figure 1b as compared to Figure 1b, the same trends can be seen. A close look at Figure 1b reveals that the peaks in the first and fourth plots are somewhat wider than the peaks of the other plots. This indicates that more signal energy is arriving with a delay relative to the LOS path in the first and fourth plots than in the others. Thus, strong frequency selective fades tend to occur during times of more intense multipath distortion.

We would like to determine how quickly the channel changes its nature over time so that adaptive techniques can keep up with the changing channel. That reduces to a question of the frequency content of the time variations of the channel. Intuition suggests that we should take a Fourier transform with respect to our time variable t, as strange as that may seem. Rather than dealing with transforms of the data in the $t \leftrightarrow f$ dimension, which corresponds to each of the ten individual plots of Figures 1a and 1b, we take the Fourier transform *across* the ten plots for each value of t. Then we compute the magnitude squared of the data. The result is called the Scattering function, $S_c(t, n)$, of the data and is plotted in Figure 1c based on the following relationship:

$$R(t,t) \xleftarrow{F} S_C(t,n). \tag{2}$$

$S_c(t, n)$ is called the scattering function of the data by analogy to various basic channel models as discussed in [4] in which multipath distortion is the result of waves arriving from point scatterers. It is the same function used by radar engineers, though the formula

may not be in the same form they are used to seeing. In the plots in this paper, the scattering function is a function of the delay t, and the *normalized* Doppler shift n_n. The normalization is performed by dividing out the dependence of the Doppler shift, n, on the velocity of the source and the angle, a, between the direction of travel and the line connecting the source and receiver as defined below:

$$n_n = \frac{n}{(velocity)\cos a} \tag{3}$$

The Doppler shift indicates the apparent change in frequency of the signal arriving at the receiver relative to the frequency it would have, had it arrived by the LOS path [5]. Thus, the scattering function shows the output power of the signal as a function of delay and Doppler shift.

First note in Figure 1c that there is a large peak in the scattering function at zero delay and zero Doppler shift. That corresponds to the LOS path, which contains the most signal power. Next one notices that the scattering function is a symmetric in the delay variable. This comes from the fact that $R(t,t)$ is a correlation function and is symmetric in t. Although the scattering function is higher in the center and generally drops off as $|n_n|$ increases, the scattering function seems to come back up. That indicates that there is still a fair amount of signal power at the edge of our Doppler shift range. This will be explored a little later in this section when the Doppler power spectrum is treated.

To simplify the information presented by the scattering function, $S_c(t, n_n)$ is reduced to a function of just one variable relating the Doppler shift to the time variations of the channel. Following the development in [5], we look at the impact of the channel on the transmission of a single tone. Setting the frequency variation to zero in our derivation is the equivalent of integrating out the delay variable in the scattering function [5] and obtain the Doppler power spectrum:

$$S_C(\nu_n) = \int_{-\infty}^{\infty} S(t;\nu_n)d\tau \tag{4}$$

The Doppler power spectrum $S_C(\nu_n)$ for our data is given in Figure 1d.

A particular disappointment in the data is the strikingly low resolution of the plots in Figure 1c in the n_n direction. That is an important problem, which should be easy to solve at a later date. Looking back at Figure 1b, one sees that there are only ten points in the time dimension at each delay, so all we can do is perform a 10-point DFT. The result is a large collection of 10-point plots in the Doppler shift dimension—one for each delay value. When we have more consecutive snapshots of the channel in time, we will be able to see much finer resolution and capture much more of the channel behavior.

To calculate what resolution is theoretically possible with our data, one looks first at the snapshot sample rate of the data. A snapshot of the channel is taken every 150 msec, so the sample rate is

$$F_s = \frac{1 sample}{150\, m\sec} = \frac{20}{3} Hz\ . \quad (5)$$

Now a 10-point DFT gives us ten frequency bins, so each bin corresponds to Fs/10 = 2/3 Hz per bin or data point. Our data will be able to tell us about Doppler shifts as large as ± Fs/2 = ± 3.33 Hz. But Doppler shifts are actually a function of relative velocities and angles. To allow us to recompute the Doppler shifts for the channel under different velocities and angles, we normalizing this range by dividing by the velocity (see Equation 3). Our points occur with a per-data-point resolution of 7.46 cycles/mm, which corresponds to a range of normalized Doppler shifts of ± 37.3 cycles/mm. To increase the range of measurable values, one must increase the snapshot sample rate. In summary, Figures 1c and 1d show us the Doppler shift behavior of the channel at a resolution of 0.67 Hz over the range ± 3.33 Hz, which corresponds to a normalized resolution and range of 7.46 cycles/mm and ± 37.3 cycles/mm, respectively. Increasing the number of consecutive samples would increase the resolution of our results and increasing the snapshot sampling frequency would increase the range of the Doppler shifts which can be observed.

As with the scattering function, we notice that the Doppler power spectrum is choppy and lacks much detail, which comes from the fact that we only had ten points of resolution in each of the DFT's which are averaged together. That problem will be reduced as soon as longer strings of data can be collected and used. We notice as with the scattering function that the Doppler power spectrum is largest for $\nu_n = 0$ and generally drops off as $|n_n|$ increases. This indicates that most of the power arrives by the line of sight path.

On the other hand, at + 37 cycles/mm, there is a secondary peak. Though little power (-20 dB) is coming through there, it is well above the noise floor, which falls at roughly –30 dB. It suggests that there might be a very specific Doppler shift in the data arriving at the receiver due to specific wave scatterer. One also notices that the spectrum never dips below –25 dB over the entire range. Compare this with the idealized time-invariant case of Figure 2d in which most of the data points are well below the noise floor (–300 dB). The output power of this segment of data is significant at all Doppler shifts in this range of ± 37.3 cycles/mm. Thus the Doppler bandwidth $B_d \geq 74.6$ cycles/mm or 6.7 Hz. (Here, "significant" is taken to mean that the receiver would easily notice the signal arriving at these strengths.)

The coherence time of the channel is the reciprocal of the Doppler bandwidth and quantifies how long the channel can be viewed as though it did not change over time. The reason for the reciprocal relationship for the two values is that they quantify functions

related through a Fourier transform [1]. For our section of Flight 11, then, the coherence time is:

$$(\Delta t)_C = 1/B_d < 150 \text{ msec.} \quad (6)$$

This is not at all surprising, since we are dealing with a rapidly changing channel. From Figure 1a it was clear that the channel changed from showing severe multipath distortion to showing very little within 150 msec. Clearly the channel cannot be viewed as time invariant for as long as 150 msec. A higher snapshot sample rate would be necessary to determine how wide large a value of B_d might suffice to capture most channel variations.

Another problem with our data is that it is heavily skewed toward the channel behavior during periods of noticeable multipath distortion. A sequence of ten channel snapshots is triggered when the bit error rate (BER) exceeds a certain threshold. That is most likely when the channel is particularly distorted through interference such as multipath. Periods of low BER are simply not represented. If rapid changes occur in the channel during periods of low BER, these will not show up in our data. Intuitively, there seems to be little danger of this. Regions of low BER generally coincide with a high signal to noise ratio (SNR) and with low channel distortion. Major changes in the channel's frequency response are unlikely and one might not worry much about updating adaptive algorithms much. Still, that hypothesis can only be confirmed when more data taken during periods of low BER is available.

THE TIME-INVARIANT CHANNEL

To gain an intuitive understanding of what this Doppler bandwidth data means, it is instructive to compare our results with those for a time-invariant channel. In a time-invariant channel, each snapshot of the time-varying transfer function $|H(f,t)|^2$ would be identical. Thus, the delay-time correlation function $R_C(t,t)$ would remain constant over our time variable t. But the scattering function would look very different from what we have seen in Figure 1c. Because there is no variation between snapshots of $R_C(t,t)$, each of the Fourier transforms which generate $S_c(t, n_n)$ would be the Fourier transform of a constant for any given delay t. That means that we should see impulse functions of n_n, and these impulse functions would all integrate to a single impulse function of n_n in the Doppler power spectrum $S_c(n_n)$.

Figure 2 gives a concrete example of the theory discussed above. We form a time-invariant system by taking the first snapshot of the channel in Figure 1a and repeating it to create ten identical snapshots of an idealized channel (see Figure 2a). Next, we compute the delay-time correlation function for the channel as described in the last section and plot it in Figure 2b. Comparing these plots with those of Figure 1a and 1b, one notices that they look very similar with one important difference: they do not vary

from one snapshot to the next. But the time-invariant scattering function $S_c(t, n_n)$ shown in Figure 2c is very different from the time-varying case. Rather than an interesting group of peaks, we have a single ridge which varies only slightly in height. The incredible drop off on either side of the ridge (-350 dB) and the discrete nature of our snapshots suggests that the scattering function is nothing but a row of impulse functions in t. These integrate to a single impulse function, $S_c(n_n)$, as shown in Figure 2d, which we plot relative to a reasonable noise floor of –30 dB, rather than showing an overly idealized dynamic range.

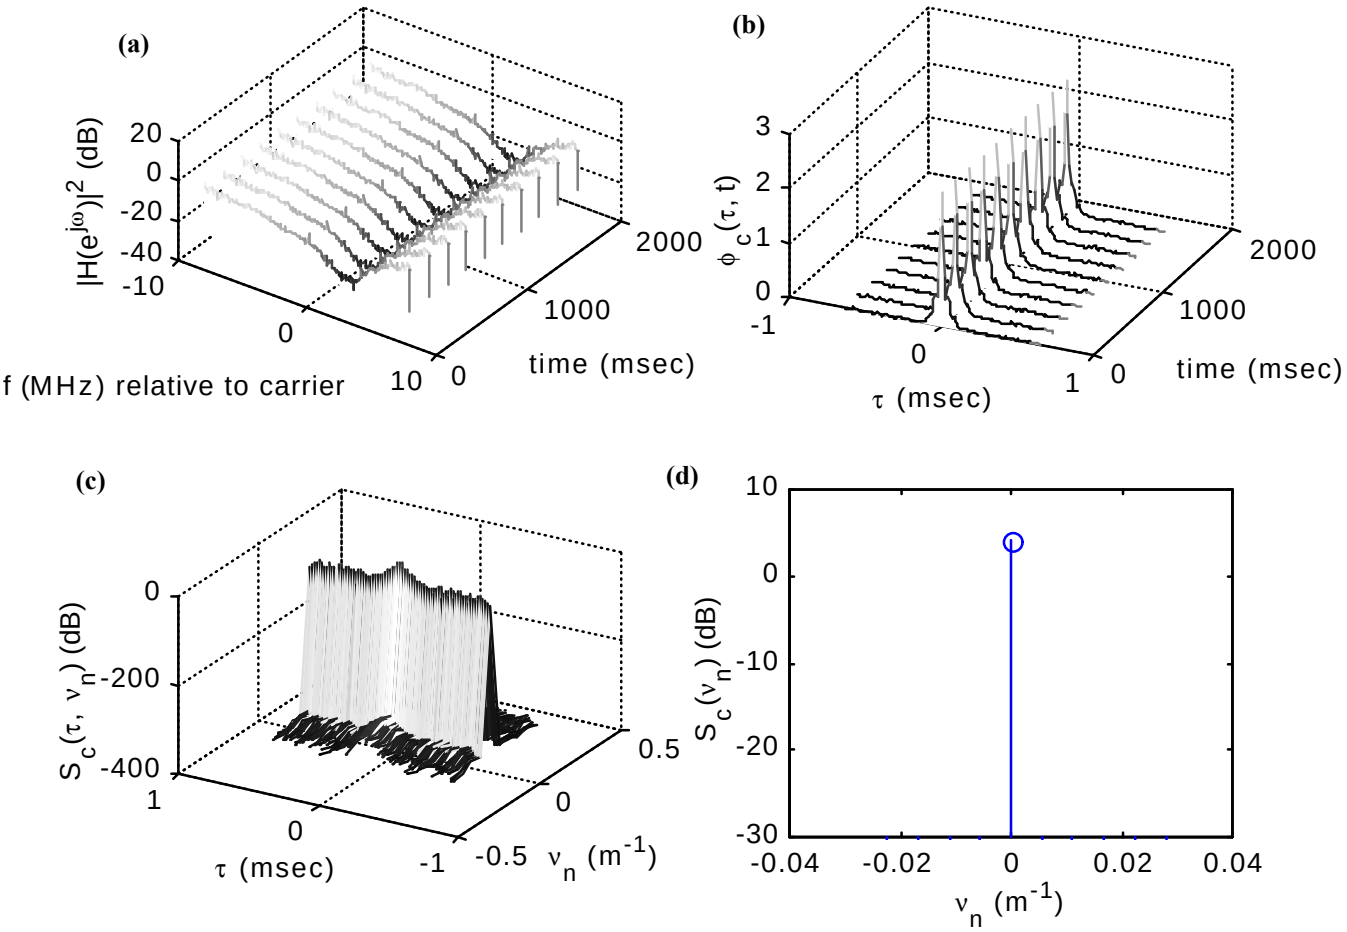

Figure 2 Data segment created from ARTM Flight 11 by duplicating the channel frequency response at 18:53:38, showing its (a) magnitude squared channel transfer function, $|H(f,t)|^2$, (b) delay-correlation function, $R(t,t)$, (c) scattering function $S_C(t,n_n)$ and (d) Doppler power spectrum, $S_C(n_n)$.

Because all of the energy is concentrated at $n_n = 0$ in the Doppler power spectrum of a time-invariant channel, as seen in the impulse function of Figure 2d, the Doppler bandwidth is zero. That makes good physical sense. There is no variation in the channel, so the snapshot frequency content of channel with respect to time is zero as characterized by a Doppler bandwidth $B_d = 0$. Channels that change slowly would be expected to have

small but nonzero Doppler bandwidths and rapidly changing channels should have larger Doppler Bandwidths. Because the coherence time is the reciprocal of the Doppler bandwidth, we find that the coherence time, $(\Delta t)_C = \infty$ for the time-invariant channel. That indicates that knowing the nature of the channel at some point in time tells us the nature of the channel at every other time. If the channel were to change slowly, the coherence time would still be large, but finite. Rapidly changing channels would have a much smaller coherence time.

DETERMINATION OF THE DOPPLER BANDWIDTH FOR A LESS RAPIDLY CHANGING CHANNEL

We found in dealing with rapidly changing channels, that the snapshot sample rate of our data may limit our ability to calculate the Doppler bandwidth. Here we look briefly at a data set from Flight 10 in which the snapshot sampling rate may actually be adequate. Looking at Figure 3a, we notice that the changes in the frequency response of the channel appear to be rather gradual. There is a deep null of –30 dB at 4 MHz above the carrier frequency in the front plot which seems to shift gradually toward the carrier frequency. As the channel progresses in time, the null eventually splits into two nulls by the tenth plot.

The plots of the correlation function $R(t,t)$ in Figure 3b suggests an explanation for the splitting of the null. The front plots all look much like those of Figure 1b which show a channel with only one strong multipath component. By the tenth plot, however, they look quite different. The main peak has split into three peaks and a couple of much shorter peaks at a delay of 370 nsec. This plot corresponds to Figure 9 in [3], with the plot reflected and made symmetric through the point $t = 0$. The difference comes from the fact that only one strong multipath component is present in the front plots but by the tenth plot at least two strong reflections arrive at the receiver in addition to the LOS signal [3]. The key idea to keep in mind for our discussion, though, is simply that these changes appear to happen slowly enough that our snapshot sampling rate (5 Hz for Flight 10) captures the channel behavior.

As a result, the scattering function has a tighter concentration of energy at lower Doppler shifts as can be seen in Figure 3c. That same concentration of energy at lower Doppler shifts is especially evident in the Doppler power spectrum shown in Figure 3d. Here, the power of the spectrum drops off consistently with increasing $|n_n|$ and drops clear down to the noise floor of –30 dB at 25 cycles/mm. This indicates that the changes in the channel all happen at a rate less than about 22 cycles/mm or at 2.2 Hz. That simply means that changes in the channel occur at a rate of 2.2 Hz or less and can be observed at 5 Hz since that just satisfies the Nyquist sampling criterion. The data suggest a Doppler bandwidth $B_d = 4.4$ Hz and the coherence time $(\Delta t)_C = 230$ msec—just larger than the separation between our samples.

It would be wrong to assume that either the Doppler bandwidth or coherence time given above is accurate for all of ARTM Flight 10 data sets. The data comes from a very brief window of the data and is not at all representative of the rest of the flight. Other segments of the data are much more like that of Figure 1 and exceed the limitations of the data as well. The problem with analyzing all of the data we currently have is that it is not a continuous run. It comes in widely separated segments which were triggered by a high BER. Large intervals are simply missing.

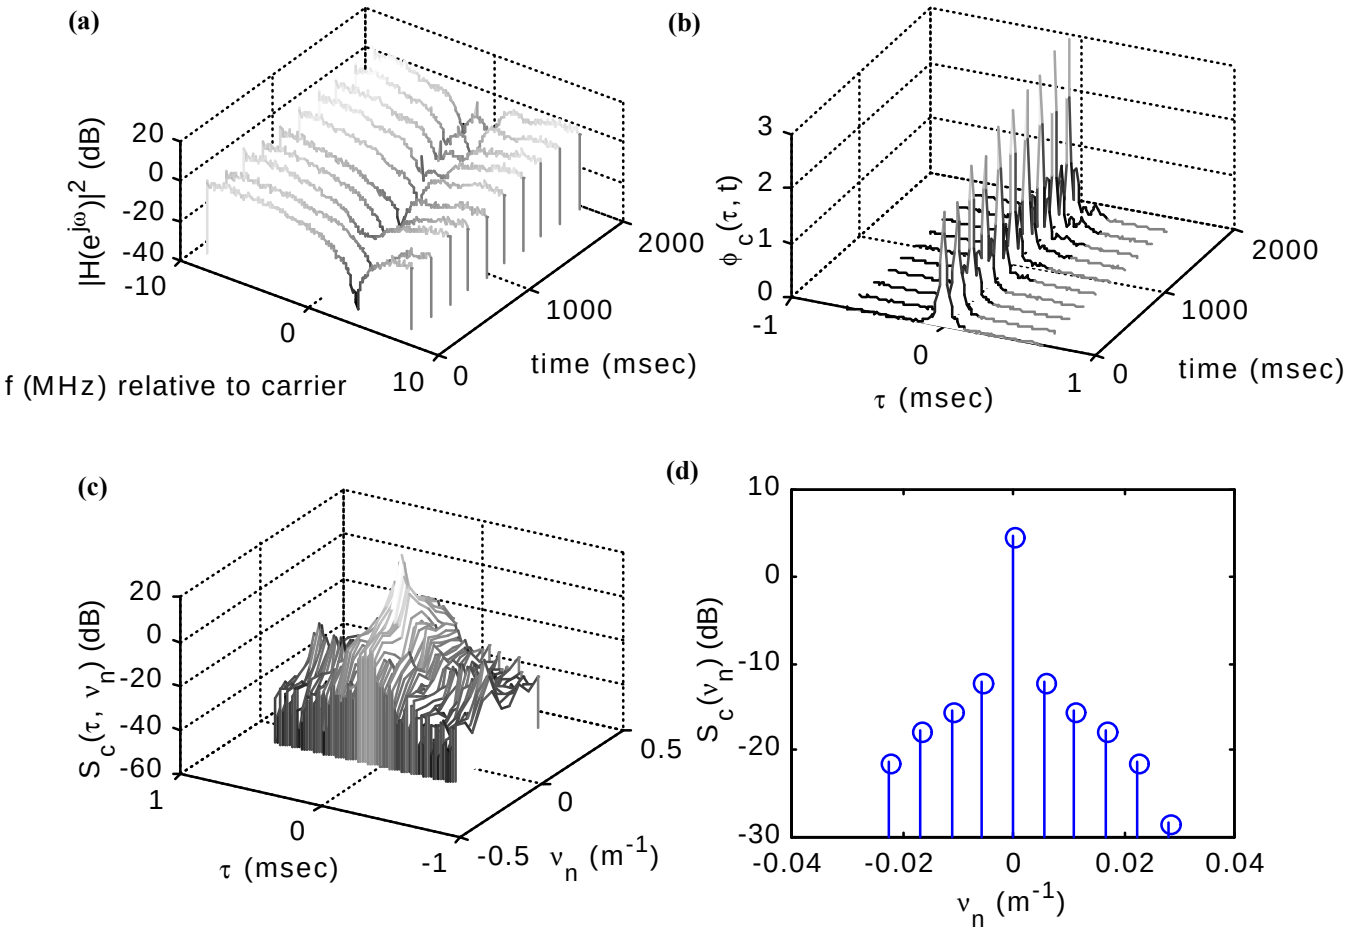

Figure 3 Data segment from ARTM Flight 10 starting at 17:30:11 showing its (a) magnitude squared channel transfer function, $|H(f,t)|^2$, (b) delay-correlation function, $R(t,t)$, (c) scattering function $S_C(t,n_n)$ and (d) Doppler power spectrum, $S_C(n_n)$.

CONCLUSIONS

The Doppler bandwidth and coherence time of a channel are valuable parameters to measure for a given channel. They indicate how quickly channel parameters are changing or equivalently how long a given channel characterization will be valid. This is extremely important in any adaptive channel techniques, such as adaptive channel equalization, adaptive modulation, adaptive channel selection and adaptive error control coding. Knowing the Doppler bandwidth helps specify an update rate for such algorithms and optimize the computational intensity of any adaptive technique.

ARTM Flight 11 data indicates that when angle of travel is ignored, at a velocity of 90 m/s or 200 mph, the Doppler power spectrum is significantly larger than zero over a range of ± 3.3 Hz. The Doppler bandwidth B_d is greater than 6.7 Hz and the coherence time of the channel is less than 150 msec. Thus any adaptive algorithm applied to the channel should update parameters at a rate greater than 6.7 times per second in order to account for the channel variations.

These values do not tell the whole story, however, because the data of ARTM Flight 11 is rather limited. Because the number of snapshots of the channel in each segment is small (10), the resolution is low. Little more can be said about the channel except that it is know at a few discrete points and could take on any value in between these points. And because the snapshot sample rate is low (6.67 Hz), we cannot say how much larger the Doppler bandwidth is than 6.7 Hz. Both of these problems reduce to one solution: we need more data. The other question to explore is to determine how skewed the data is toward segments containing channel distortion. It would be valuable to determine for sure that the channel really does appear less time variant during periods of low BER than during periods of high BER. Should that prove true, the current sounding technique will tell us all we need to know about the Doppler bandwidth. In that case, data not represented by the data would have been drowned out anyway by the results of the data which does get recorded.

Similar data can be calculated for ARTM Flight 10, but also paints an incomplete picture of the channel. For one segment of Flight 10, the Doppler bandwidth is 4.4 Hz and the coherence time is 230 msec, but the data is rather out of context. The current channel sounding technique only offers disjoint segments triggered by a high BER and is not representative of the full channel.

Because an unlimited data set is not a possibility, the most important question to address might be what range of the Doppler spectrum is worth examining for our channel. To determine that would require a higher snapshot sample rate to capture a wider Doppler shift range of the Doppler spectrum. The next issue to address might be to double-check that a BER-independent collection of data reveals no more about the channel than the

current data collection method does. Finally, it would be important to collect a large number of consecutive channel snapshots and examine the channel behavior at a finer Doppler shift resolution.

REFERENCES

[1] Proakis, John G., "Digital Signaling over Fading Multipath Channels," Digital Communications, McGraw-Hill Book Company, New York, 1989, pp. 702-716.

[2] Bello, Philip A., "Characterization of Randomly Time-Variant Linear Channels," IEEE Transactions on Communication Systems, vol. CS-11, pp. 360-392, 1963.

[3] Rice, Michael, "ARTM Channel Sounding Results—An Investigation of Frequency Selective Fading on Aeronautical Telemetry Channels," Proceedings of the International Telemetering Conference. vol. 35, October 1999.

[4] Parsons, David, "Wideband Channel Characterization," The Mobile Radio Propagation Channel, John Wiley & Sons, New York, 1994, pp. 161-188.

[5] Haykin, Simon, "Advanced Communication Systems," Communication Systems, John Wiley & Sons Inc., New York, 1994, pp. 730-764.