CONF-9508153--1

LA-UR- 95-2446

Title: TRANSIONOSPHERIC CHIRP EVENT CLASSIFIER

Author(s): Paul E. Argo
T. Joseph Fitzgerald
Matt J. Freeman

Submitted to: Proceedings on Applications of Time-Frequency and Time-Scale Methods
University of Warwick
Coventry, UK
August 30-31, 1995

DISTRIBUTION OF THIS DOCUMENT IS UNLIMITED

MASTER

Los Alamos
NATIONAL LABORATORY

Form No. 836 R5
ST 2629 10/91

TransIonospheric Chirp Event Classifier

Paul E. Argo, T. Joseph Fitzgerald and Matt J. Freeman

Space and Atmospheric Sciences Group
Nonproliferation and International Security Division
Los Alamos National Laboratory
Los Alamos, New Mexico, USA 87545

Abstract

In this paper we will discuss a project designed to provide computer recognition of the transionospheric chirps/pulses measured by the Blackbeard (BB) satellite, and expected to be measured by the upcoming FORTE satellite. The Blackbeard data has been perused by human means -- this has been satisfactory for the relatively small amount of data taken by Blackbeard. But with the advent of the FORTE system, which by some accounts might "see" thousands of events per day, it is important to provide a software/hardware method of accurately analyzing the data. In fact, we are providing an onboard DSP system for FORTE, which will test the usefulness of our Event Classifier techniques in situ. At present we are constrained to work with data from the Blackbeard satellite, and will discuss the progress made to date.

1. Background

This classifier system must respond to a variety of RF signals that will be detected at the satellite. The background, essentially narrowband signals, will be treated in the section under prewhitening or excision (see figure 1). These are simply the scores of RF carriers for radio and TV stations. A subtler form of RF emissions come from radar systems. Those most commonly seen by the BB system are pulses that look essentially narrowband on the bandwidths dealt with here (>1 MHz). They tend to last for a few microseconds and are relatively bright. The FORTE multiband trigger system is designed to help alleviate problems with them, but if the system is operating near the noise threshold (as it usually must be) then they often will push the trigger threshold into the 'event' category. It is then incumbent upon the classifying system not to spend too much time and energy trying to categorize the pulse, and certainly not to misclassify it. Undispersed, broad band pulses from onboard electronics are also a major part of the noise environment. The multiband trigger system cannot discriminate these from dispersed pulses unless specifically designed with an anti-coincidence detector. These broad band pulses tend to affect the lower portions of the VHF band (30 - 60 MHz) most. Although BB was designed with the recognition that noise from onboard digital systems would be a problem, and major efforts were made to reduce this kind of noise, it is still a limiting feature of the broad band measurements. Finally, in percentage of occurrence, come the dispersed events of interest. As naturally occurring events they are seen as TransIonospheric Pulse Pairs, or TIPPs (see Figure 2). Several other reports have been written about these (Massey and Holden, 1995), so their general features will be discussed only briefly here, as they affect an event classifier system. Most characteristically, as indicated in the name, they are seen as pairs of dispersed pulses. The distribution of power between the pulses varies considerably, from the first pulse being bright with the second nearly invisible, and vice versa. These pulses are usually wider than several microseconds in the Short Term Fourier Transform (STFT) diagram (at a single frequency), indicating perhaps an extended source. In fact, some very long events (hundreds of microseconds) have been observed, and may be the connection of these events to classical lightning discharges. The dispersion can spread the arrival times of the high and low frequencies by as little as a few microseconds to greater than 100 microseconds. The manmade ionospherically dispersed pulses are similar to these TIPPs in some ways, and very different in others. They show the same range of dispersion effects, but are narrower (in the STFT domain) than any of the observed TIPPs. This is primarily because these are coming from localized sources, and very short (several nanosecond) electrical discharges. The spectral power distribution within the received pulse can vary significantly as well (far more than the variation of power with frequency in the source), with at least part of the effect being due to constructive (destructive) interference between various propagation modes. In any case, some show maximum power above 50 MHz, while others have almost no power there at all. The longer lasting TIPP events tend not to show such abrupt frequency characteristics, perhaps because there is more 'smearing" and less chance of modes interfering.

2. Program Operation

The computer program (or event classifier) described here is an IDL version of a system that might fly on a satellite, or that might be used by ground stations to monitor events. The BB data is brought into the classifier in one megabyte packets, which is about 6.7 milliseconds. A series of sliding STFTs are then performed on the data. The spectrograms are 16384 samples long (109 microseconds), consisting of 128 transforms of length 256 and an overlap of one half..

prewhitening/frequency excision: The major problem with broad band VHF data at the satellite is the number of narrow band carriers (FM and TV stations) that contain significant power (see figure 1a). Because the BB trigger system uses the entire broadband signal (including the carriers) to trigger, the trigger level must be set at a very high threshold. Future systems will use narrow band trigger units to avoid this flood of power, and will also contain "prewhitening" circuitry to take out the power in the carriers. For the BB data we can do the prewhitening (or in our case excision) by operating on the raw data. A true prewhitener would take the FT of a long data sequence, and modify the power spectrum so that it contains equal power across the entire bandwidth. A true excision unit would do similarly, but threshold the process so that no frequency had a net increase in power. This is important in the noise issues, because the prewhitener might increase the net noise level by boosting essentially "dead" bands.

In our processing we do the excision in the spectrogram, or STFT domain. A long series of relatively short Fast Fourier Transforms (FFTs) are taken, to provide a time history of the signals spectral content. We average the set of power spectras, which will give us essentially the same character as the long term power spectra (with much lower frequency resolution). Then, in the STFT domain we divide each individual Fourier transform by the normalized summed power spectra (limited to being greater than 1.0 - so it is excision). This has the effect of removing the power from carriers in the STFT domain. Figure 1b shows this effect, as the only difference between figures 1a and b is that figure 1b has had the carriers excised.

time excision -- 20 kHz torque coils on BB: There can be a large amount of noise in the form of undispersed pulses at the satellite. Since these are not of interest to the transionospheric measurements, it is necessary to eliminate their effects on the classifier system. We have chosen to perform an excision along the time axis -- in other words we add up power in the STFT along the ordinate (time) axis and then divide each pixel by the appropriate normalized power. This tends to lower the effects of events that stay mainly in one time bin, relative to all other events. There is a noticeable lowering of power even in the desired chirps, but at an amount far smaller than that to the undispersed chirps. This post processing technique has essentially eliminated false classifications of undispersed pulses as events of interest.

trigger: A second major problem with identifying periods of transient pulse reception on a transionospheric VHF signal is that of triggering on the event. The BB method of triggering on the time domain signal limits the threshold of transients that can be detected, because power in the transient pulse must be comparable to the power in the entire broad band spectrum, which may well contain some powerful CW signals.

A first step solution is to have a hardware signal prewhitener, which will lower the power of the CW carriers. This can potentially gain the system about 10 dB in signal to noise. FORTE will contain not only the hardware prewhitener, but also a set of narrow band triggers. This set of eight triggers will be operated in an "n out of m" coincidence mode (m=8). This does not solve all the problems, however. First, the pulses of interest are created below the ionosphere, and hence are dispersed as they pass through the ionosphere. This dispersion causes the lower frequency portions of the pulse to arrive later than the high frequency portions, possibly up to several tens of microseconds later. So in this case either the trigger box must have a fairly wide time coincidence window (increasing the base noise level and false triggers) or it must have frequency dependent time offsets to allow for the expected dispersion (increasing the system complexity). In the FORTE case there is no mechanism to provide time offsets, and so the frequency selection for the trigger boxes must be matched to the time apertures.

The FORTE trigger is set to have eight channels spread across 20 MHz, and will require n of eight channels to trigger in a given time window (adjustable width) as a condition for keeping a data set. As we will show on various weak pulse sets (coming from the Los Alamos Portable Pulser (or LAPP), the distribution of power across the "low" band (from 28 MHz to nearly 100 MHz) is quite variable, and it will be difficult to establish where to place the 20 MHz window for optimum pulse capture.

TEC correlation: Once the data has made it through the preliminary filters, and is believed to be an event of interest, we apply a sliding correlation of the STFT to a predefined set of frequency chirps. We have selected 32 Total Electron Content (TEC) levels (which define the chirp slopes), varying logarithmically from the lowest to highest expected TEC levels (10^{16} to 5×10^{18}). In a straight forward manner we step along each chirp, adding the power at each pixel lying on the chirp. In addition, we keep a counter of how many pixels contributed to the chirp correlation, where only pixels with powers within 6 dB of the brightest pixel in the STFT are counted. The chirp defined at TEC(0) is the undispersed chirp, to ensure we "trap" them at this point if they have gotten through the other filters.

neural classifier: At this point nothing has been said about whether the event is a TIPP or a manmade chirp. In order to ascertain this, we have 'dechirped' the data, using the best fit of the TEC to do so. We look at the nine time bins centered on the chirp, and smoothed in frequency from 128 values to fifty. We then sum up the power in each column (time) and frequency (row), yielding two vectors (9 and 50 long) rather than a 9x50 array. We then pass all the information to a neural net, training for TIPP and pulser events.

4. Blackbeard Data and the Event Classifier Performance

We have a set of 115 files, containing data taken when the BB RF system triggered due to the time domain trigger unit. These events are classically found to be the TIPPs seen in Figure 2, and are likely to be very "bright". pronounced events. When the Event Classifier (EC) system scans through these sets, we identify a TIPP chirp in every data set. The character of these chirps can vary widely, from well defined pulse pairs, to pulses with a very dim mate, and even one or two very long poorly defined events. In fact, the EC found three data sets that contained not one, but two or more separate TIPPs. This is interesting, because it indicates that these pulses may be much more common than expected, and with more sophisticated triggering systems we may be saturated with data. In one of the multiple chirp sets the 'extra' chirps are so weak that they are much more like the Los Alamos Portable Pulser (LAPP) in amplitude, and would not have triggered the BB data acquisition system.

The LAPP data set, meant to simulate more closely events of interest, has been much harder to accurately classify. The chirp events are typically very weak, and some data sets appear not to contain chirps at all. In addition, even when chirps are readily obvious to the human eye, the constructive interference between propagating modes can cause a "braided" appearance, with regular dropouts along the frequency axis. If we allow the EC program to select its own "quiet" bands for trigger selection, there are a couple events for which we cannot satisfy the trigger criteria. If by human intervention we force the trigger selection to choose a higher frequency band, where these two events have more consistent power, then the trigger criteria are also easily met, and the EC finds and classifies the events. This indicates that before FORTE flies we need to provide a good estimate of the location of the 20 MHz band in which the narrow band triggers will be placed.

As discussed in the Classification section, we use various measures taken from the STFT to determine whether an event has happened. These include the excess power in the spectrogram (above some predefined mean), the peak pixel brightness, the best fit TEC value, the correlation of the TEC chirp to the STFT, and the number of pixels meaningfully involved in the correlation value. In the IDL version of the EC program we routinely plot these parameters for each megabyte of data analyzed. Figure 3 shows the graphical outputs for a LAPP event. In the LAPP events the spectrogram power might show no increase during the chirp event. In this case other methods are necessary to detect them than simply looking in the time domain.

5. The Event Classifier and Signal to Noise Issues

The event classifier software described above has been operated extensively on data taken from the Blackbeard satellite. As has been discussed, the trigger system for capturing data on BB is a wide band system, and so very weak chirped signals cannot trigger the system. The LAPP data set, which constitutes data with much weaker signal strengths, was obtained by pulsing an RF signal aimed at the satellite, and capturing a longer data set that would be expected to contain the chirped pulse. This data will contain chirped signals that have a much lower signal to noise value than the more common TIPPs events. We have chosen to use one of these data sets, and to add progressively stronger random Gaussian noise to the digitized data, in order to simulate even weaker signals. In this way we can study the effectiveness of the present EC software.

The discussed simulation tests only part of the EC system, because for simplicity we passed the EC a single 16K data set that contained a LAPP chirp, and "told" the classifier that the triggerbox returned a value of '8'. We let the EC do a TEC correlation search, and used the standard criteria for determining if an *event* was found. At each noise level we reset the event background by adding the same level of Gaussian noise to the sample of data that determined background conditions. We proceeded to apply one hundred realizations of the Gaussian noise to the chirped data set, each time searching for the chirp. We recorded the number of times the Event Classifier correctly detected the chirp at that noise level, increased the noise level and proceeded.

When the EC software is tested against such a data set, it performs similar to, but perhaps slightly worse, than the human eye. At 15 dB of added noise (minus 11 dB signal to noise in the time domain) the EC is giving 90% accuracy (see Figure 4). At 18 dB added noise the classifier has dropped to about 25%. Investigation of this data set showed that the requirement that the brightest pixel in the spectrogram be more than 8 times brighter than the average was no longer being satisfied most of the time. If we lower the threshold level, we find a corresponding increase in detections. If we go so far as to eliminate the threshold then the classifier eventually gets deceived by noise, and "finds" events where they do not exist. We cannot arbitrarily move this threshold, it was set at a level that protected us from false alarms due to the random signals that are actually in the data, and we will probably have to settle for a classifier effectiveness that is slightly below the human potential.

6. Discussion

The hardware/software Event Classifier project has been developed to cope with the expected increase in the data load, as the more sophisticated FORTE system starts collecting data. The present method of having a data analyst examine the data sets, searching for transionospheric pulse chirps, will not work for such increased data set. Therefore, we have been developing an Event Classifier system that will scan the data, pick out times of "interest", and determine the character of the detected event. This might include categories such as: 1)transionospheric pulse pairs, 2) single short pulses (manmade in origin), 3) radar systems, and 4)anomalous signals (often called "zoo" animals). The present paper discusses a part of the EC system that singles out the chirped transionospheric pulses to provide to later stages of the entire Event Classifier. It has been designed to be insensitive to undispersed pulses and narrow band radar systems. No real "zoo" events have been seen by either human observation or this system, so we cannot say much about that situation here. The Blackbeard RF system has imposed its own limitations on the data set, in particular that detected pulses will be very bright, especially when viewed in the time-frequency domain. We expect that there will be a plethora of TIPP-like events once the triggering thresholds are lowered -- in this case we may discover a jungle of new event characteristics. For the present we are constrained to work within the BB data set, and have shown that the EC system is either as "good", or nearly as good as the human observer, and clearly can examine the upcoming increase in data flow far more easily.

Acknowledgement: This work was funded by the United States Department of Energy

REFERENCES

Massey, R.S. and D. H. Holden, *Phenomenology of Transionospheric Pulse Pairs*, accepted by Radio Sci., 1995

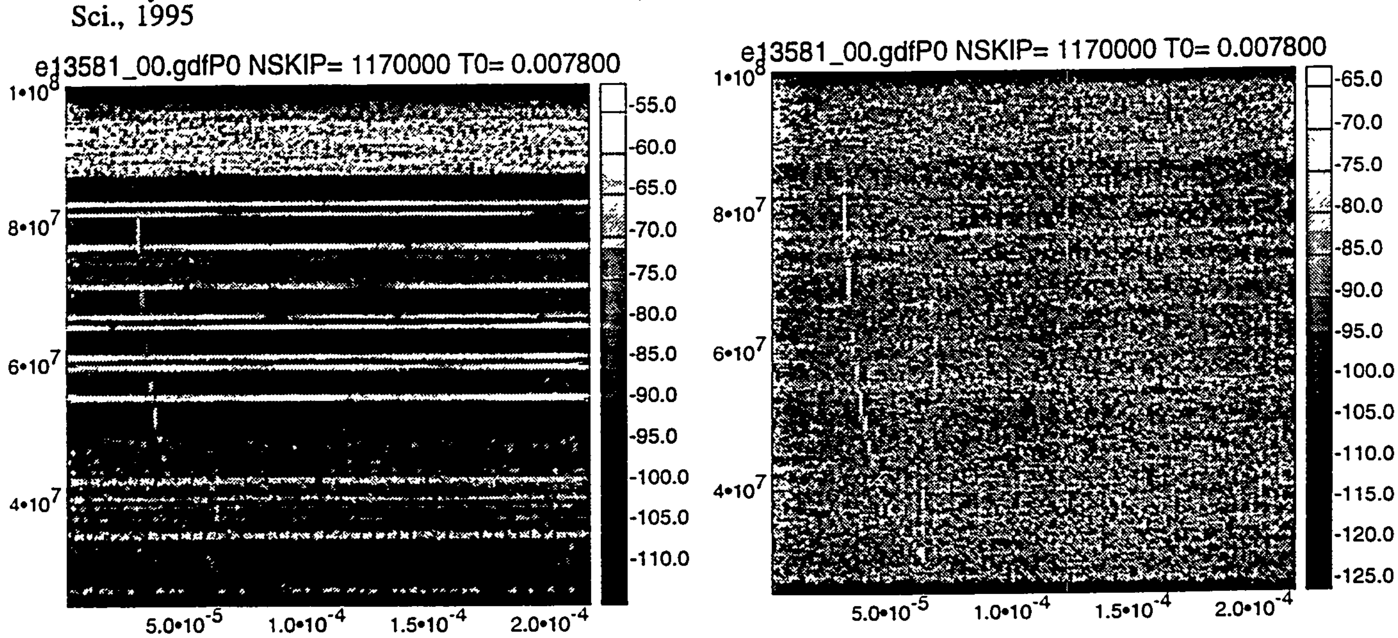

Figure 1 a)This figure shows the 'raw' spectrogram of a LAPP pulse. The ordinate axis is time in seconds, the abscissa is frequency in Hz. The signal levels are in dB. Note that the predominant signals are the CW carriers (horizontal lines), and that the chirp (towards the left) could be masked by the carriers. In narrow band frequency ranges the pulse will be the brightest signal, b) similar to figure 1a, but after frequency excision has lowered the power in the CW carriers. The LAPP is now the brightest feature on the spectrogram, and is far more easily recognized by feature recognition techniques. Note the characteristic curvature of the pulses - this is due to the dispersive nature of the ionosphere, through which lower frequencies travel slower.

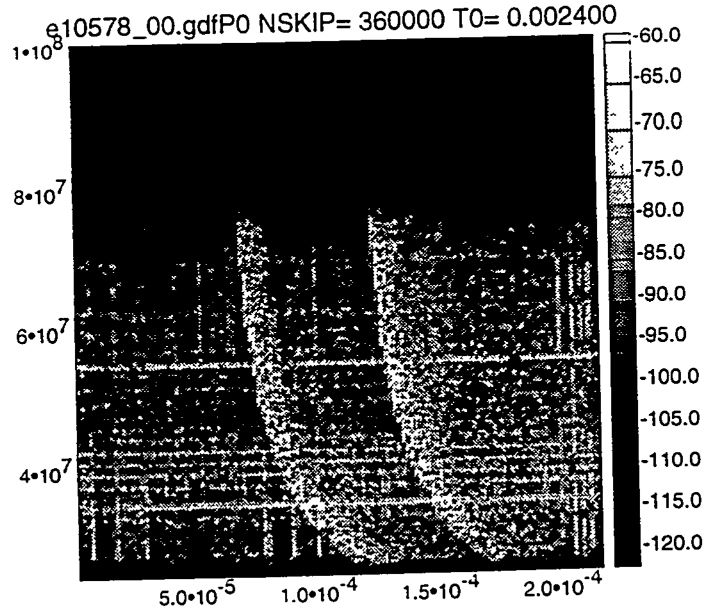

Figure 2 A 'classical' TransIonospheric Pulse Pair (or TIPP). The reason that these pulses are found in pairs has not been explained.

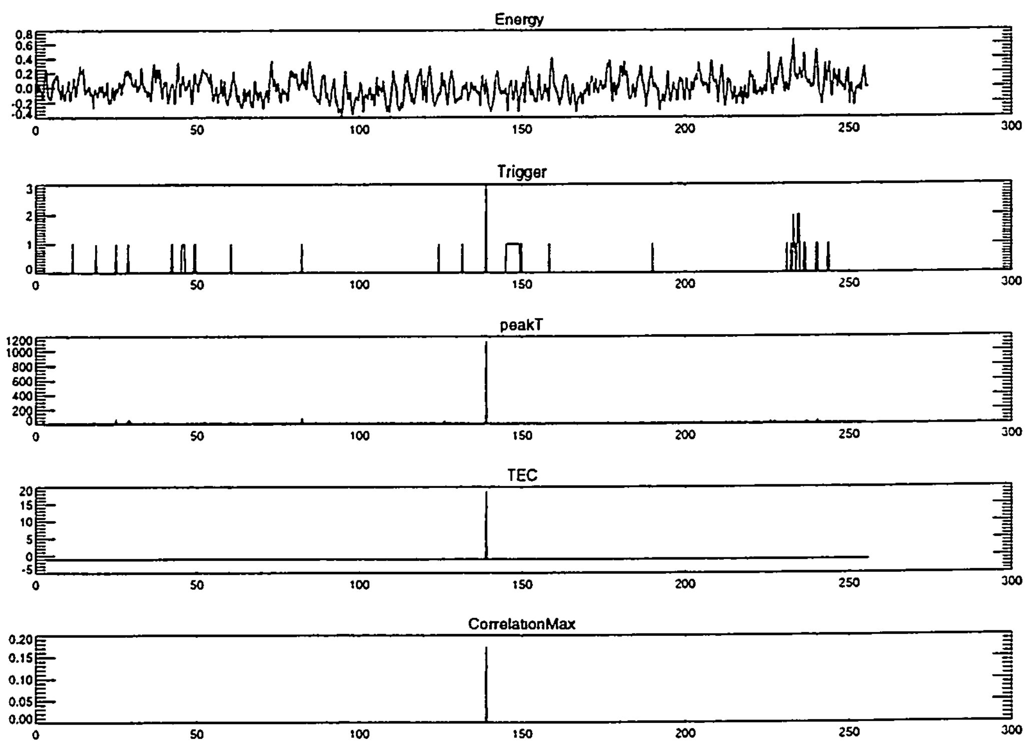

Figure 3 This set of curves is output from the Event Classifier routine for a LAPP data set, and shows the summary of a specific megabyte of data (approximately 1/150th of a second). The ordinate for all five curves is the time, indicated by the count number of the individual spectrogram (in 1 MB there are 256 total 16K spectrograms with overlaps of 1/2). The top curve shows the relative excess power in each spectrogram -- notice that there is no increase at the time of the pulse time 136. The next curve down is the output of the simulated trigger box, which shows much more activity than the normal TIPP data, probably because the satellite is over the continental US,

and the RF noise is much higher. The third curve, showing the peak brightness in each normalized spectrogram, does show a large increase during the LAPP event. This is indicative of the value of moving from the purely time domain into the time-frequency domain. The fourth plot down is the TEC value corresponding to the maximum correlation of the individual spectrograms to the theoretical chirps defined by TEC levels. The bottom curve is the maximum correlation in each spectrogram, normalized to the average spectrogram power, and indicates how many bright points were lying along the theoretical chirp. It is in the bottom two curves, when the data is correlated against expected chirp shapes that we see a significant enhancement in the system resolution. In this case the LAPP stands out nicely against the background. Given the set of available LAPP data we have been able to "tune" the system to detect LAPP events, and to have very few false detections. We expect once the much wider dynamic range of the FORTE data is fed into the classifier we will need to rework the algorithms extensively. In this data set a "non-event" can be seen at time 35 -- the values of peak brightness and correlation were small enough to discard this as noise.

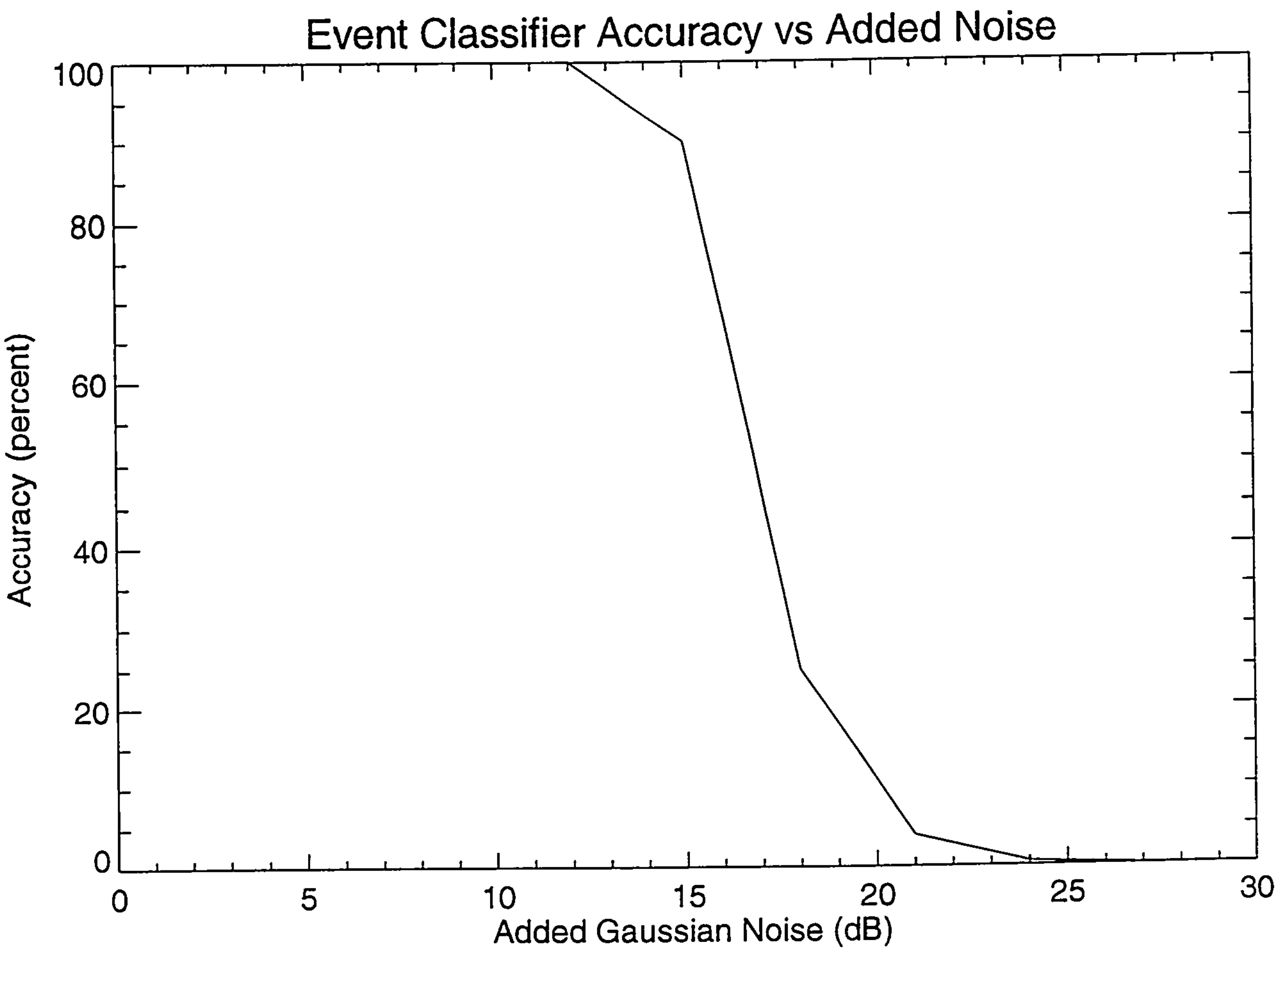

Figure 4. A plot showing the capability of the EC to identify the transionospheric chirp in the presence of increasing Gaussian noise. The ordinate is the percentage of tries that the classifier found the event, the abscissa is the added noise measured in dB relative to the background noise already present in the data. Visually we determined that the addition of 3 to 4 dB of additional noise lowered the signal to noise ratio to 0 dB in the broadband data. At about 13 dB of added noise (or -10 dB signal to noise) the EC began to start missing detections; by 18 dB added noise (-15 dB signal to noise) it was severely impaired.