RECEIVED

JUL 01 1996

OSTI

*Operation and Modeling of the FORTE
Trigger Box*

Los Alamos
NATIONAL LABORATORY

MASTER

*Los Alamos National Laboratory is operated by the University of California
for the United States Department of Energy under contract W-7405-ENG-36.*

DISTRIBUTION OF THIS DOCUMENT IS UNLIMITED

one

An Affirmative Action/Equal Opportunity Employer

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither The Regents of the University of California, the United States Government nor any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by The Regents of the University of California, the United States Government, or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of The Regents of the University of California, the United States Government, or any agency thereof. The Los Alamos National Laboratory strongly supports academic freedom and a researcher's right to publish; therefore, the Laboratory as an institution does not endorse the viewpoint of a publication or guarantee its technical correctness.

*Operation and Modeling of the FORTE
Trigger Box*

Tim Murphy

Operation and Modeling of the FORTE Trigger Box

by

Tim Murphy

Abstract

The fast on-orbit recording of transient events satellite (FORTE) will carry a multiple-narrow-band trigger designed to detect impulsive VHF signals embedded in a high-noise background. The FORTE trigger boxes consist of eight VHF channels spaced across twenty MHz of bandwidth. A trigger is generated when a sufficiently bright signal is seen in a user-defined number of these channels within a specified coincidence window. In addition, the trigger circuitry incorporates a feature to reject events caused by the actuation of narrow-band carriers. This report describes the trigger's operating principles and their implementation in the satellite hardware. We then discuss a computer model which can be used to simulate the performance of the trigger circuit.

Introduction

One of the outstanding issues in broadband VHF signal detection is how to detect transient-impulse signals in an environment that includes continuously operating transmitters such as television or FM-radio stations. Although the impulse may be quite bright in many parts of the spectrum, the presence of the carriers can make the overall broadband signal-to-noise ratio (SNR) small; this makes traditional threshold trigger techniques ineffective. A promising approach to this problem is the use of multi-channel coincidence trigger techniques. This technique has demonstrated detection of broadband impulses in complex noise environments where the SNR is 0 dB or less.

Los Alamos National Laboratory is building a research satellite designed to detect and characterize terrestrial VHF impulses, known as the fast on-orbit recording of transient events satellite (FORTE). The FORTE satellite payload includes two broadband VHF receivers which incorporate multiple-narrow-band triggering hardware of the type described above. In order to understand the expected performance of this system we have developed a software model of the FORTE trigger box. The software is designed to faithfully imitate the actual hardware while also allowing some additional flexibility. In this report we describe the operation of the actual hardware trigger in some detail and then the software model of it.

Principle of Operation

The problem we are faced with is the detection of a very brief, transient VHF signal embedded in a background consisting of both "white" noise and deliberately transmitted signals such as TV and FM broadcasts. The white noise is characterized by having a continuous spectrum with an approximately uniform, low amplitude power spectral density, while the broadcast signals are typically much higher amplitude but are confined to specific frequency bands. Because the transient signals of interest are inherently broad-band it is desirable to record them with a broad-band receiver. The amplitude of the transient signals is such that they are clearly discernible against the terrestrial white noise background using simple level triggering. However, the presence of powerful transmitters within the frequency band of interest makes it impossible to detect the transient signals using just a level trigger. A narrow-band receiver tuned to the appropriate band can be used to attempt to avoid the transmitters, thus allowing the use of level triggering, but one then loses most of the information in the transient signal and there is still a problem with false triggers generated by stochastic variations in the background.

The idea behind the multiple-narrow-band trigger is to combine the best features of both narrow- and broad-band schemes. The broad band of interest is subdivided into a number of narrow-band channels, each of which is independently monitored. A true transient event will cause simultaneous or sequential signals in many or all of the channels, while stochastic background fluctuations will be uncorrelated. Thus requiring multiple-channel coincidence will eliminate most of the random false event problem. It is possible that one or two of the channels may be occupied by a transmitter, making detection of the transient in that channel impossible. We therefore do not require coincidence between all channels but only a subset of them, typically 60-70 per cent.

Another feature of a true transient signal is that it is of short duration. We can eliminate false events due to the sudden actuation of a transmitter by requiring the signal duration in each channel be limited. By combining these requirements of signal coincidence across multiple channels and short duration of the signal in a channel we can effectively trigger on transients in a challenging noise environment.

Hardware Trigger Description

The FORTE satellite is equipped with two 20-MHz-bandwidth, twelve-bit-resolution receivers. These are superheterodyne VHF receivers that are tunable across the range from 25-300 MHz. The trigger circuit consists of eight 1-MHz-wide channels spaced evenly across the 20 MHz receiver bandwidth. The trigger box sits behind the mixer stages of the receiver, so although the radio can be tuned over the entire VHF band the input to the trigger box is always mixed down to the band from 0 to 20 MHz (nominally). The trigger box contains eight separate 1-MHz-bandwidth receivers located in the 0-20 MHz band at 1.25, 3.75, 6.25, 8.75, 11.25, 13.75, 16.25, and 18.75 MHz respectively.

The average power is monitored in each trigger channel. The average is calculated in a logarithmic sense, that is the power is monitored in dB and the average power is taken to be the mean of the dB values. The average is updated approximately ten times a second. Data is read out from the trigger channels at two million samples/s, the rate appropriate to the information bandwidth, and then goes to logic circuitry to be processed.

The trigger box has several adjustable parameters that govern its operation and allow its performance to be tailored to changes in the noise environment. These are the threshold power level (S), the carrier reject time (C), the coincidence window (W), and the number of coincident channels (N). The effect of these parameters will be described below.

A signal must meet three criteria to generate a system trigger: a threshold test, a duration test, and a coincidence requirement. The first criterion is a threshold test applied separately to each channel; the signal in that channel must exceed the average power level by a user-specified amount. Because the noise environment changes significantly in the course of an orbit, the threshold is specified in relative rather than absolute terms. The threshold is defined by first specifying a single overall relative threshold in the form of S dB above average power levels. The exact threshold level in each channel is then calculated by taking the average power, as discussed above, and using the specified offset to derive a threshold level for that channel. Thus the absolute detection threshold will be different for each channel; it depends on the noise level in the channel. This procedure is used so that one noisy channel does not drive the whole system to needlessly high thresholds. Typical values for the threshold are 16 or 17 dB.

The second trigger criterion is also a single-channel test. The criterion is that once there is a signal above threshold it must drop below threshold again within a specified time interval C. This test is designed to reject events caused by the sudden actuation of a transmitter whose frequency happens to lie in the channel. Choices for this parameter are 1, 2, 5, and 10 μ s.

Each channel is connected to a pulse-forming circuit which generates a 1-bit binary data stream whose default value is zero. There is a user-specified parameter W that determines the coincidence window for the trigger. If the first two criteria are met, this circuit generates a string of ones W μ s long that begins at the time at which the analog signal first crossed the threshold. This procedure is shown graphically in Figure 1 where we show the power envelope for an hypothetical signal. The envelope briefly crosses the detection threshold early in the record, producing an digital output pulse. Later the record crosses the threshold and remains above it due to the actuation of a transmitter. Because this second event does not drop back below threshold quickly enough it does not produce a digital output pulse. Allowed values for the duration of the coincidence window are 5, 10, 20, 40, 80, 100, 324, and 640 μ s.

Figure 1 : Input and Output Data Streams

The third trigger criterion is a comparison that involves all of the channels. The binary data streams from each channel are added together, and whenever the sum equals or exceeds a user-specified number (N), typically five, a system trigger is generated. This process is shown graphically in Figure 2. Here we see uncorrelated pulses occurring at various times in the different trigger channels, with channel 6 being particularly noisy. A trigger is generated when a broadband signal produces pulses in five or more channels simultaneously, as happens about two-thirds of the way through the record shown in the figure. The intermediate step of creating binary data streams is used to allow us to weight each channel equally (a simple linear addition of the unscaled raw data would be dominated by the loudest channel) and also makes it simpler to incorporate a variable length coincidence window.

Figure 2 : System Trigger Derived from Multiple Channels

Software Model

In order to explore the effects of varying the parameters that control the trigger system we have written an IDL computer program, called TRIGBOX, that is designed to duplicate the functions of the FORTE trigger box in software. The model can be set up to duplicate the FORTE trigger box exactly but also has the flexibility to allow some experimentation with the trigger design. The features that are fixed in the hardware but variable in the model are the number of channels, the spacing between the center frequencies of the channels, their bandwidths, and the detection thresholds, which can be set individually in the model but which are fixed by a single parameter in the hardware.

The first step in the computer model is to read in the data that specify the trigger parameters. These include the number of channels (N_c), the sampling rate of the data (150 million samples/s for Blackbeard), and the parameters N , C , and W discussed previously. For each channel the user must specify the center frequency, bandwidth, and detection threshold (S_n), where $1 \leq n \leq N_c$. Although the hardware trigger always sees input in the range 0–20 MHz, the channel frequencies in the

model are given in terms of the frequencies of the raw input signal. This is done to simplify analysis using spectrograms and other tools which show the input data in its true frequency range. The model allows the S_n to be different, while in the real hardware they are specified by the single parameter S .

Once the trigger configuration has been read in, we find the ratio of half the sampling rate (i.e. the Nyquist frequency) to the largest trigger bandwidth and then calculate the largest power of two that is less than or equal to this ratio. We use this as a reduction factor when determining the approximate information-bandwidth limited data rate for the trigger channels. For example, if we are using a sampling rate of 150 million samples/s the Nyquist frequency is 75 MHz, and therefore for 1-MHz-wide trigger channels the ratio of Nyquist to trigger bandwidth is 75 and the reduction factor is 64.

After we configure the trigger, the next step is to read in the Blackbeard (or other) data files. The program is currently designed to accept data in the form of Blackbeard general data format (GDF) files, but it can easily be modified to accept other time-series data. The list of files to be processed is read in from a source file, and the data files are handled in the order in which they appear in the source file. Blackbeard files are too large to be conveniently handled as a single unit so we process them in 256 kbyte segments. A small amount of additional data preparation is required because, although Blackbeard takes data in the bands from 25–100 MHz or 100–175 MHz, the data received on the ground has been mixed with a 100 MHz oscillator and is always in the band from 0–75 MHz. The simplest way to account for this in the model is simply to subtract 100 MHz from the trigger channel frequencies when analyzing high-band data or to reverse the order of the channels and subtract the channel frequencies from 100 MHz when analyzing low-band data. These steps would not be necessary for unmixed input data.

Once the trigger frequencies have been suitably adjusted, the input data is sent to a filtering routine called TRIGFLTR, which is designed to imitate the action of a single trigger channel. The input to this routine is the time-domain waveform, the center frequency of the channel, the bandwidth of the channel, and the bandwidth-limited data rate reduction factor described above. The first step in the processing for this routine is to convert the time-domain data into the frequency domain by using a fast Fourier transform (FFT). The FFT indices that represent the frequencies within the trigger channel are calculated, and a square-edged filter is applied by setting all elements outside the range of in-band indices to zero. The in-band data is then mapped into the lowest-index elements of the FFT array. The resulting array is truncated to a length equal to the original length divided by the reduction factor defined above. For our example of 1-MHz-wide trigger channels and 150 million sample/s input data the reduction factor is 64. The final step is to convert the shortened FFT back into the time domain to give the in-band signal in the channel, sampled at the rate appropriate to the information bandwidth of the broadest channel. If the channel bandwidths are different, we could in principle use a different reduction factor for each one, but using the same factor for all channels makes the length of the output data stream the same for all channels and allows a simple comparison algorithm to be used later in the program.

The next step in the model is to run the output from TRIGFLTR through a pulse-forming routine, PULSFRMR, which implements the first two trigger criteria. This routine takes the time-domain channel data and the parameters S_n , C , and W as input and returns a one-bit binary data stream indicating at which points the data passed the first two trigger criteria.

The first step in the pulse-forming algorithm is to convert the field values to instantaneous power. This is done by taking the logarithm (base 10) of the square of the field value at each point and multiplying by ten to convert to dB. These values are then added together and the total is divided by the number of points in the record to get an average power. The appropriate S_n is added to the average to get a threshold power level. The threshold power is then subtracted from the time series of power values. The resulting time series has negative values when the instantaneous power is below threshold and positive values when it is above threshold. The second step in the pulse-forming procedure is to go through the time series and find those points that are positive. Whenever a set of positive points is found, we determine the starting point and the ending point for that set. We then check to see if the length of the set exceeds C . If the length of the set is short enough the starting point is recorded for future use, otherwise it is discarded and we move on to the next set. Once the entire time series has been examined we generate a one-bit data stream of the same length as the original time series. The values in this data stream are initially set to zero. We then take the set of points that indicates the start of a legitimate signal, as described above, and for each one generate a string of ones starting at that point in the time series and whose length is W . This binary data stream is the output of this routine. The pulse-forming procedure is shown graphically in Figure 1.

The routines TRIGFLTR and PULSFRMR are applied separately for each trigger channel. Once this has been done for all channels we have a set of N_c binary data streams of equal lengths. These streams are sent to another routine, COMPARATOR, which generates the system triggers. The input to this routine is the set of binary data streams and the parameter N . The output is an array of system trigger times and the number of channels that generated each trigger. The routine begins by setting a flag to zero. It then begins working its way through the data, adding the values from the N_c binary streams point by point. Whenever the sum equals or exceeds N and the flag is zero the flag is set to one and the time and number of channels are recorded. The flag remains set at one until the sum drops below the threshold value, at which time the flag returns to zero. The use of the flag eliminates multiple triggers from the same event.

In principle this is all there is to the trigger model, except for some standard I/O to write the results to an output file. We have added another routine, however, to deal with the problem of spacecraft noise on Blackbeard. There are several noise sources on the satellite which operate intermittently. When they are in operation they produce periodic noise bursts that generate a large number of false triggers. These false triggers can be easily eliminated by inspection by human operators both because of their periodic nature and because they do not show any ionospheric dispersion. The trigger logic is not sensitive to dispersion (and would require

significant modification to take this into account), but we can easily deal with periodicity. This is handled by a routine called DITCHBBNOISE.

DITCHBBNOISE is invoked when the number of triggers in a record exceeds ten. The routine simply looks at the time between triggers. For the first two triggers the time difference between them is calculated, and if that time equals one of the known satellite noise-source periods then both triggers are discarded. For the third and subsequent triggers the same procedure is followed, except that when the time difference between triggers is one of the known noise periods only the later of the two triggers is discarded. This procedure relies on the assumption that a true event is unlikely to occur exactly on top of one of the noise triggers, although this could in principle happen. It appears to be quite effective in practice. If this routine is invoked and actually discards triggers, that fact is noted in the output file.

The last step in the model is to write the results into output files, which is done by the main routine TRIGBOX. For each input data file a separate output file is generated that contains the trigger settings, the number of triggers generated, a declaration if onboard noise reduction was used, and the location and number of channels of each system trigger. In addition, the code writes a summary file which simply contains the filenames and the location of triggers generated for all of the input files processed in that run. Figure 3 summarizes the logic flow of the model.

Figure 3 : Logic Flow in TRIGBOX

100-1000

This report has been reproduced directly from the best available copy.

It is available to DOE and DOE contractors from the Office of Scientific and Technical Information, P.O. Box 62, Oak Ridge, TN 37831. Prices are available from (615) 576-8401.

It is available to the public from the National Technical Information Service, US Department of Commerce, 5285 Port Royal Rd. Springfield, VA 22616.

Los Alamos
NATIONAL LABORATORY

Los Alamos, New Mexico 87545