TBAD Operation Manual

Revision 6

June 23, 2015

Contents

1 Introduction

The TBAD (Transponder-Based Aircraft Detector; or alternatively Tom and Bill's Airplane Detector) is a passive monitoring device of RF transmissions at 1090 MHz, conceived at UCSD by Bill Coles and Tom Murphy, and largely designed by Allen White, George Kassabian, and Mike Rezin in the Physics Electronics Shop at UCSD. Revision 5 units were the last to be built at UCSD, after which licensing transferred to Aircraft Avoidance Systems, LLC—but with the same individuals participating.

The central concept behind the antenna is to monitor the *power* at 1090 MHz in two beam patterns on the sky, comparing the two to decide if the source of transmission is within a "protected" cone. This is done using a "narrow" beam generated by an array of eight patch antennas, and a broad beam formed by a single, shared patch—in our case the central patch in the arrangement of eight. Though neither antenna is technically omni-directional, we have adopted the parlance—for better or for worse—of DIRECitonal and OMNI to refer to these two antennas.

The sidelobes of the narrow beam are always weaker than the broad beam response for a given offset angle, so that only signals within the central lobe of the narrow beam will have a stronger "array" signal than the "omni" signal. This ratio is a feature of the beam patterns alone, and is insensitive to absolute signal strength, polarization of the source, transmitter distance, etc.

1.1 The Physical Device

The physical apparatus separates into three primary subcomponents, shown in Figures 1 and 2 (and Figure 3):

- Array antenna with summers, splitters, attenuators, presenting two SMA connections for the OMNI and DIREC signals. This component is passive (requires no power).
- Discriminator box (weather-proof) to be located near the antenna. This box accepts the antenna inputs, connects to the decoder box via a power/signal cable, and is where thresholds are set, signals can be monitored, and behavior can be set by jumpers. The discriminator box is low power (~ 5 W), so it can be placed in front of a telescope without creating a thermal disturbance.
- Decoder box (rack-mount) that can be far from the discriminator (by as much as 200 ft by cable), containing the microprocessor, power supply, laser shutter control, and serial communications port.

More pictures of the electronics appear in Section 14.

1.2 Introduction to Signal Types

TBAD detects transmissions in a narrow band (± 20 MHz) around 1090 MHz. Legitimate transmissions from aircraft are pulsed, in a fixed-time format. The bulk of traffic at present is Mode A or Mode C (referred to collectively as Mode A/C), which has 450 ns pulses on a 1.45 μs cadence. Mode S uses 500 ns pulses on a 1.0 μs cadence, consisting of a preamble followed by a pulse-position modulated scheme.

Figure 4 shows the formats. For Mode A/C, the X-bit is never present, in practice, and the four octal data values, A, B, C, and D—each 3-bits or 0–7—are present or not in the remaining 12 pulses.

When TBAD senses arrival of a signal—via interrupt—it assigns this as the first framing pulse (F1). It then executes a sequence of seven interrupt tests to look for pulse edges (low to high) and one sampling (relevant only for the C1 bit) in order to differentiate between a Mode S preamble and a Mode A/C transmission. Depending on the result, this is followed by a sequence of pure interrupts for Mode A/C or pure sampling for Mode S.

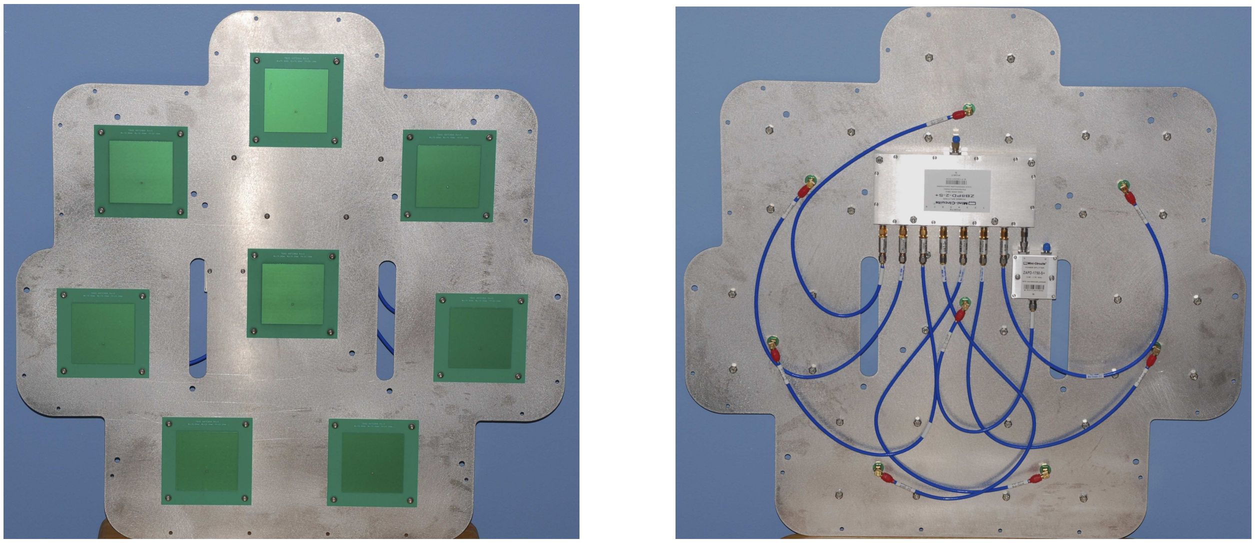

Figure 1: Antenna Front and Back. As oriented here, the antenna is sensitive to polarization in the up-down direction, which would need to be oriented vertically on the sky to match the vertical polarization of transponder signals. The array signal emerges from the large 8:1 summing box on the back, while the broad-beam (OMNI) signal emerges from the small 1:2 splitter just right of center.

Figure 2: Discriminator box (top) and decoder box (bottom). The discriminator box is contained within a weatherproof shell, and has connectors for the two antenna signals and for the umbilical to the decoder box (carries power and communications).

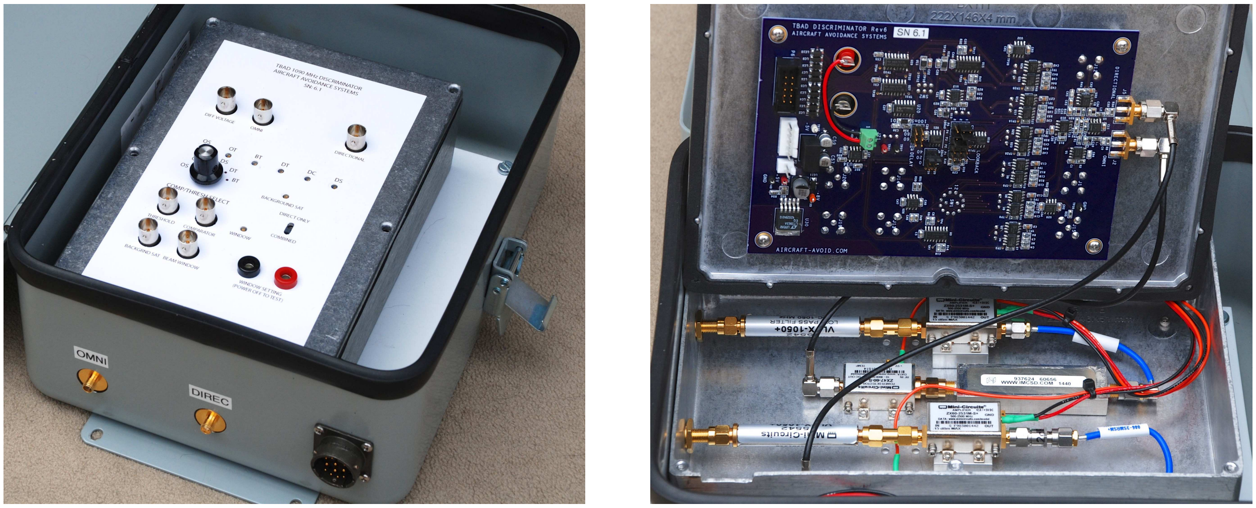

Figure 3: The discriminator box housed within the gray weatherproof enclosure pictured in Figure 2. Note the SMA antenna inputs and the data/power connector on the bottom of the box. Threshold settings and signal monitors are located on the discriminator front panel. Removing the lid of the inner discriminator box reveals the circuit board (where jumpers are located) and the RF input chains.

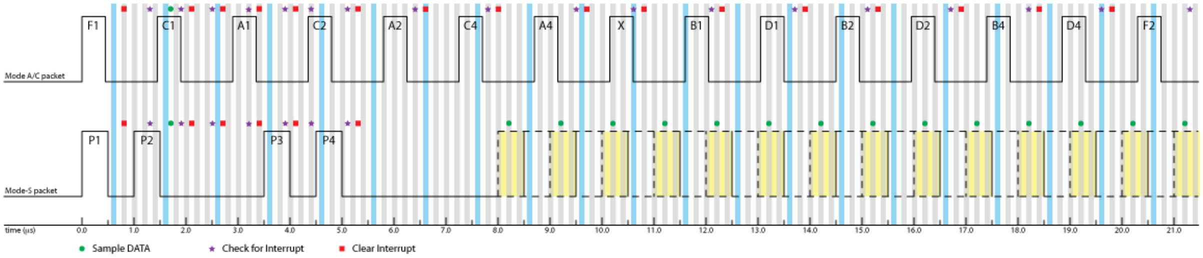

Figure 4: Pulse formats for Mode A/C and Mode S transmissions. For Mode A/C, the first and last framing pulses (F1 and F2) must be present, while the rest may or may not be present depending on the transmistted code. Mode S begins with a static preamble and then fat or skinny pulses with guaranteed transitions at the solid lines and data values represented at the positions of the yellow regions. The vertical bars denote computer instruction moments within the interrupt service routine.

Distance Measuring Equipment (DME) signals can fall in the 1090 MHz band, and appear as two broad pulses 12 μs apart. Since the first pulse does not at all mimic the Mode S, the signal is interpreted in the Mode A/C context, and lights up either bit B1 or D1 via interrupt 12 μs later. Also, the C1 sample is usually satisfied as well, so that typical ABCD codes for DME signals will be 0110 or 0011.

Interference will typically have a single transition within the 21 μ s interval (interpreted as F1), and could also satisfy the C1 sample, producing codes 0000 or 0010, respectively, with no X or F2.

2 Setup

The setup is very simple. Most of the cables have unique connectors/genders so that connection mistakes are unlikely. A quick procedure:

1. Situate components so that the antenna has a proper sky view, the discriminator box is very near the antenna (sheltered if possible), and the decoder box is in a sheltered environment with access to AC power and a computer or terminal server.

2. The antenna should be oriented so that the feed point (solder bump) on each patch is centered horizontally and offset vertically (up or down does not matter). This ensures vertical polarization—which is irrelevant for zenith pointings. If in a sun-exposed location, it may be wise to put a plastic sheet/bag over the antenna to take the UV hit, rather than the patches themselves.

3. Connect the antenna to the discriminator using the shortest practical RG-58 coax cables with SMA connectors ($<$ 2 m recommended). This is the one step that could result in crossed connections. The DIREC signal emerges from the 8:1 summing box on the back of the antenna. The OMNI signal emerges from the 1:2 splitter connected to the central patch. (See Figure 1.)

4. Run the multi-conductor gray cable between the decoder and the discriminator. The connectors are of opposite gender on each end, so direction matters.

5. Connect an AC power cord to the decoder box and turn on the power switch. The LEDs on the decoder convey the status. The POWER and NO FAULT LEDs (both green) should be ON. The SHUTTER CLOSED LED (red) will be ON at turn-on, but should go OFF after 10 seconds if there are no offending conditions. The other three LEDs (Fig. 14) flash when the DIREC signal is stronger than some "saturation" setting, when an airplane is perceived to be in the protected beam, and when the OMNI signal is stronger than its "saturation" threshold.

6. Until you know you want to operate differently, switch the discriminator toggle to COMBINED data and the decoder toggle up to OMNI ENABLE. Set the knob on the decoder to the number of "in-beam" triggers needed in the last 10 s to close the laser shutter. For locations with sparse skies, such as over Mauna Kea, a low number like 3 is recommended. For busier skies, a higher number like 20 may be more suitable. For very busy skies, a dipswitch in the decoder multiplies the front-panel number by ten for increased range (Section 6).

7. Connect a computer to the serial output of the decoder box or connect this to a terminal server and establish one-way (listening only) communication at 115200 8N1 (Section 10 below).

3 Threshold Settings

TBAD operation is largely based on tunable thresholds. The behavior is therefore *highly* configurable to suit differing environments.

In what follows, the DIREC (array; narrow beam) signal strength will be denoted D, and the OMNI (single patch; broad beam) signal strength will be denoted O. The ratio, $R = D/O$ will usually be expressed in logarithmic form (dB), so that $R = D - O$ when both D and O are expressed in dBm or equivalent scale. Also, we refer to the various thresholds by names assigned on the circuit schematic and discriminator front panel, which will be summarized after the threshold concepts are covered.

3.1 Threshold Concepts

When a source is on the boresight of the antenna, D is 11 dB stronger than O. In practice, we require $R > 5$ dB. This gives us some margin over simple equality, which is otherwise prone to spurious, noise-induced false alarms. The result is about a 12° half-angle on the sky of "protected" zone. This spatial zone can be tuned somewhat by changing the ratio threshold. The corresponding threshold is labeled **BT**, standing for Beam Threshold. Increasing this threshold narrows the effective protected zone (disappearing if BT exceeds ~ 1 V), while lowering produces a broader zone (approaching 25° as BT$\rightarrow$0 V).

Because the ratio is insensitive to absolute power of the signal, we need to restrict our attention to signals strong enough to be at a viable distance. For instance, an airplane at 40,000 ft at an elevation angle of 15° has a slant-range of about 50 km. We can estimate the received power for a transponder transmission at that distance, and require that a signal reach this level before we react to it. This threshold is called **DT**, for Directional Threshold.

We also want to protect against airplanes that saturate the array power detector, because this could result in an artificially small ratio for a plane. When D approaches the saturation of the detector (around 0.55 V), we should therefore trigger a closure condition. We call this the **DS**, or Directional Saturation threshold. Likewise, if O is near saturation, an airplane must be *very* close, and we should close the laser shutter as a precaution, because a nearby airplane will have a high angular rate and may slip into the protected zone faster than TBAD can respond—given that TBAD relies on the interrogation rate of an aircraft by ground stations and other aircraft. We call this the **OS**, or Omni Saturation threshold. Both *DS* and *OS* can be used not only to protect against saturation, but also adjusted deliberately shy of saturation to create exclusion zones out to some tunable distance, augmenting the angular protection. This is especially helpful to mitigate high-angular-rate aircraft in areas where interrogation is sparse.

We have found that multi-path interference sometimes creates false "in-beam" alarms by destructively interfering a signal against its reflection. A similar effect can be produced when signals from two aircraft overlap in time. If the destructive interference/summing happens in the O channel, the D channel looks comparatively higher than it should, triggering the ratio threshold. We get around this by only looking at the first ~ 75 ns of the pulse, and also—optionally—requiring that we had no distracting signals on the leading edge of the pulse. We accomplish the latter via the **DC** (Directional Clear; lookback) threshold, so that the directional signal must initially be weaker than a certain level to activate a comparison.

Finally, it can be very useful to log transponder activity that does not pose a threat to closure, so that we may learn more about the wider world of the sky and gather useful characterization data. The threshold, **OT** for Omni Threshold, dictates when we record these extra signals, as seen by the broad-beam antenna. In high traffic-density areas, it is recommended that this feature not be used outside of characterization tests, lest TBAD be overwhelmed by irrelevant traffic. Bypass switches on the discriminator box and the decoder box govern the behavior of the omni-log feature (see Section 8).

3.2 Threshold Implementation

Thresholds are easy to change and monitor. The thresholds have trim-pots accessible from the outside of the discriminator box (but within the weatherproof enclosure), and a BNC connector for monitoring the threshold setting (labeled THRESHOLD). A knob selects which threshold is being monitored. Figure 2

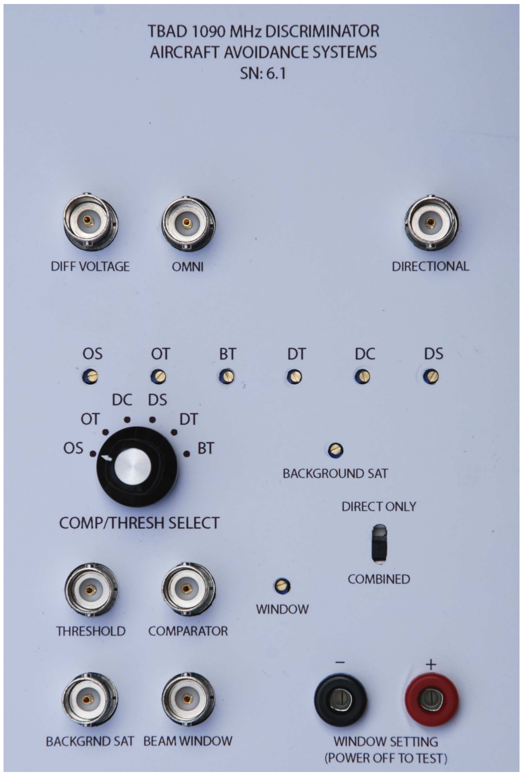

Figure 5: Discriminator box front, where threshold adjustments are made. The knob selects which threshold is made available to the THRESHOLD BNC, with trimpots labeled accordingly. Signals from the two RF power detectors are available, as is the differential amplifier (gain = 4) output. The various comparators associated with the six thresholds may also be monitored at the COMPARATOR connector, likewise knob-selectable. One may also adjust and monitor window width and background saturation duty cycle from the front panel. Finally, one may select whether transmitted DATA is based on directional signal (DT) alone, or combines omni (OT) as well.

shows the six front-panel threshold adjustment pots, and the selector switch that determines which signal is presented to the THRESHOLD output.

Except for the ratio threshold, BT, the threshold values all refer to levels on the RF power detector. The logarithmic detector delivers a voltage between about 0.5–2.0 V, and is perhaps counter-intuitively *inverted*, so that *low power corresponds to* 2 V, and *high power* (saturation) *corresponds to* 0.5 V. The response is shown in Figure 6—note particularly the green curve most appropriate for 1090 MHz. Over most of the range, the response is rather linear (on a logarithmic plot), so that each factor of 10 in power (10 dB) is 0.25 V in output.

From the figure, we can fit the 1000 MHz line reasonably with:

$$P(\mathrm{dBm}) = 25.82 - 40.98V, \tag{1}$$

or:

$$V = 0.63 - 0.0244P(\mathrm{dBm}). \tag{2}$$

The signal strength from a transmitter of known power at known distance can be turned into a received, amplified power in dBm, which can then be expressed as a voltage produced by the power detector.

The RF power delivered to the power detector is

$$P_\mathrm{D} = G_\mathrm{A} + G_R + G_\mathrm{T} + P_\mathrm{T} - 20\log_{10}(4\pi) - 20\log_{10}(R/\lambda), \tag{3}$$

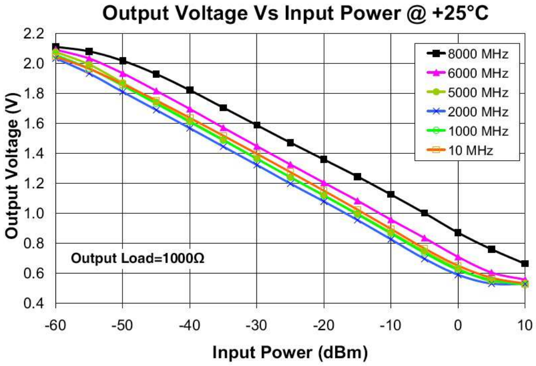

Figure 6: Power detector response to input signal strength. For the purposes of the end-user, it is only important to realize that for high power inputs, the detector saturates at about 0.55 V (green 1000 MHz curve applies). When the detector sees no power, it will deliver about 2.1 V. In practice, background RF noise may result in a floor between 1.9–2.0 V. It is useful to characterize the background level in the setup location.

where $G_{\mathrm{A}} \approx 33$ dB is the net amplifier gain (including 3 dB attenuation employed in antenna synthesis); G_{R} is the receiver gain (5 dB for omni, on axis; 11 dB for the DIREC antenna at the "edge" of the protected zone; 16 dB for DIREC beam center); G_{T} is the transmitter (aircraft antenna) gain, which should be about 5 for a half-space dipole transmission; P_{T} is the transmitter power (48.5 dBm for 70 W; 51 dBm for 125 W; 57 dBm for 500 W—see the next paragraph for why these levels are relevant); R is the line-of-sight distance to the transmitter (aircraft), in meters, and $\lambda = 0.275$ m is the wavelength. For example, a 125 W transmission at 50 km would produce a signal at the power detector of -27.2 dBm on the "edge" of the directional beam, registering 1.29 V according to Eq. 2.

Figure 7 provides a graphical means to determine the voltage that will be delivered by the power detector for various transmitters in various locations within the beam for the omni and directional antennas. First, it is useful to know that the FAA requirements on transponders demand at least 70 W of peak (pulse) power for *any* transponder, and 125 W for any commercial aircraft. Meanwhile, no transponder should have peak power above 500 W. For example, an airplane at 100 km and centered in the beam will produce pulse signals reaching down to 1.15 V if operating at the maximum allowable 500 W, 1.32 V if operating at the minimum commercial aircraft power of 125 W, and 1.38 V if operating at the minimum power for any aircraft of 70 W. But sensitivity at the edge of the beam is more important, and these translate to 1.3 V, 1.45 V, and 1.51 V for the three decreasing power levels, respectively. Meanwhile, the omni antenna, with lower gain, will produce signals at 1.45 V, 1.59 V, and 1.65 V if in the center of the (broad) beam.

The "in-beam" threshold criterion, BT, is referenced to the output of a difference amplifier with a gain of 4.0. Thus a 1 dB signal difference results in an amplifier input difference of 0.0244 V, which when multiplied by the gain, becomes 0.0976 V, or 0.1 V in practical terms. We then have a convenient relationship that the BT threshold, in volts, is just the ratio threshold in dB divided by 10. So if we require the difference/ratio to be 5 dB, this means we will set the BT threshold to 0.5 V.

A reasonable set of default thresholds is:

BT	0.5 V→5 dB difference: ratio of narrow to broad signal strength
DT	1.30 V: sets "in-beam" sensitivity, or range
DC	1.80 V: sets condition of low background before pulse for valid "in beam"
DS	0.85 V: sets saturation level for directional antenna

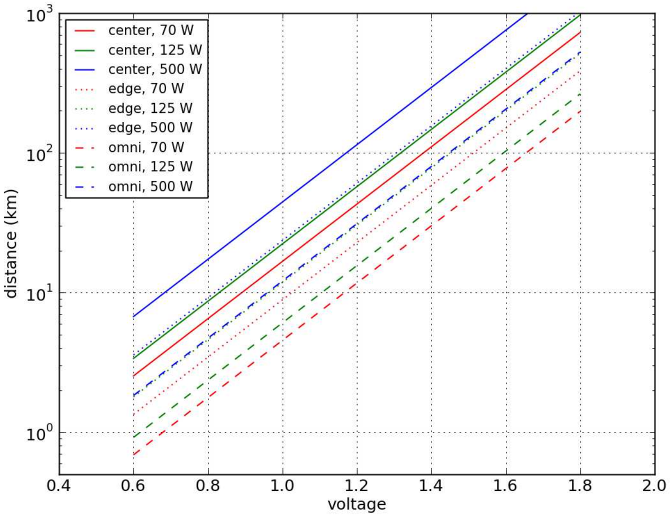

Figure 7: Voltage-Distance relationship for various transponder powers and locations within the beam pattern, assuming a good polarization match.

OS	0.85 V: sets saturation level for omni antenna (plane very close)
OT	1.40 V: sets range in which to record all aircraft, when enabled. 1.4 V corresponds to 30–80 km range, approximately.

One way to characterize performance given a set of thresholds (as in the above list) is by showing the distances to which various types of planes are protected. Figure 8 gives such an example for the above threshold settings.

One cautionary note: these calculations all assume that the transmitter's polarization is vertical (plane not in steep bank), that the antenna is likewise aligned vertically, and that we are in plane of the the transmitter beam pattern, which looks like a donut. This is appropriate for planes at large distance, where we push sensitivity limits. But a plane directly overhead may transmit little power straight down (hole of toroidal beam pattern). Luckily, this is largely compensated by the closer range, so we don't worry too much about missing these signals.

There is also some empirical evidence that our end-to-end RF signal chain suffers 3 dB (factor of 2) attenuation via accumulated cabling and connector insertion loss. So the distances computed above might be more accurate if a $\sqrt{2}$ factor is applied.

4 Window Settings

We have learned that most of our false "in-beam" events come from crowded airspace where signal pulses can overlap (collide) and interfere with each other. We also suspect a significant contribution from multi-path interference—especially when structures occupy the intermediate foreground. The transponder signals are so powerful that we have enough margin in signal-to-noise ratio to ignore most of the pulse, only looking at the leading edge. This capability partly addresses both problems because we can demand a clean leading edge to ignore many pulse collisions, and obviate multi-path interference which is always delayed from the main line-of-sight pulse. We therefore make a "window" during which we assess whether the directional/omni ratio qualifies it to be "in beam." Figure 9 shows the general pulse timing layout.

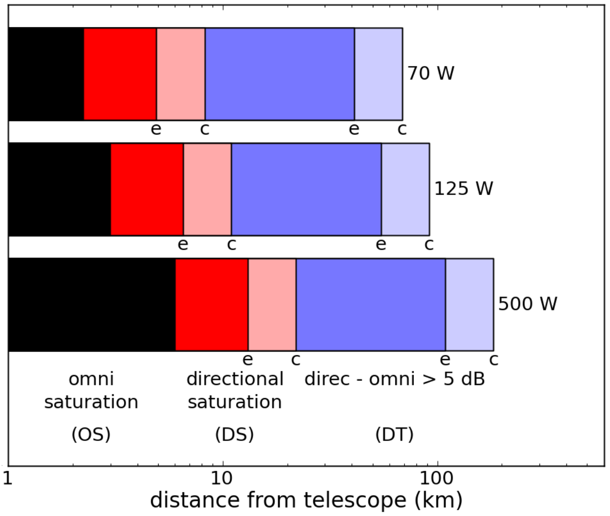

Figure 8: Protected zones for three transponder powers (covering the FAA max/min requirements), for the threshold values listed (DT=1.3; DS=0.85; OS=0.85). The "e" and "c" labels denote edge and center of the protected zone (about 15° half-angle). For example, a 70 W transponder at the edge of the protected zone will trigger out to 40 km range, and will saturate at 5 km.

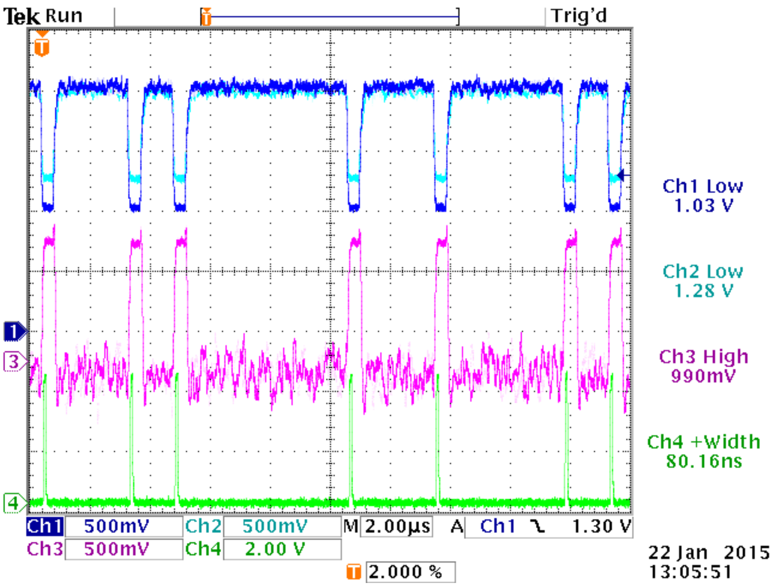

Figure 9: Typical Mode-A/C pattern type, as viewed on an oscilloscope. The pattern is produced by TSIM with mode knob at 12 and parameter knob at 5, delivered through 6 dB 1:2 splitter followed by an additional 10 dB attenuation on the OMNI channel. Dark blue is the DIREC signal; cyan is OMNI; magenta is the difference amplifier, and green is the narrow beam-check window. Each (negative-going) pulse produces a strong difference signal, which is only assessed during the window period.

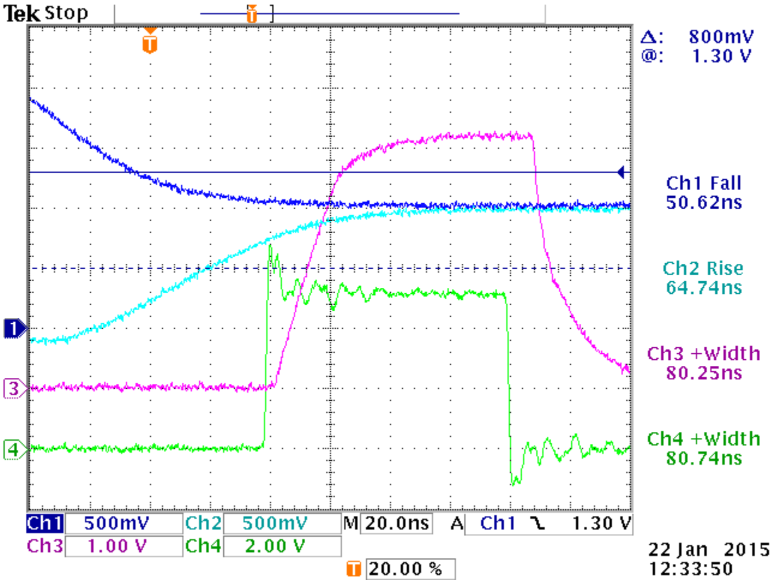

Figure 10: Extreme close-up of the leading edge of one pulse from Figure 9 (different colors), showing pulse and window timing. Dark blue is the directional signal; cyan is the DIFF amplifier output; magenta is the BEAM WINDOW output; and green is the internal logic signal (at R57) upon which the BEAM WINDOW output is based. All but the last are available on the front panel of the discriminator box. The blue line is the DT trigger point (here at 1.3 V), and the window is formed 46 ns after the directional signal crosses the DT mark (jumper at 40 ns, plus 6 ns static delay) with an 80 ns width. The DIFF (ratio) condition is evaluated only within green window. We can see that the DIFF signal has enough time to stabilize before the end of the window.

Mode A and Mode C pulses are 450 ns in duration, and Mode S pulses are 500 ns in duration. DME pulses are microseconds long, and have slow rise times compared to the other types. We therefore generally elect to look at only the first $\sim$ 80 ns of each pulse, where we find the most reliable ratio information. We can set a delay for the start of the window in 20 ns increments (internal jumper), as well as the width of the window (external potentiometer). An additional jumper setting impacts the look-back time for demanding a "clean" leading edge in tandem with the DC threshold.

Before proceeding further, examine Figure 10 to get a sense of the timing. The incoming pulse (having been smoothed by an input filter) has a rise time of about 50 ns. The difference signal is a bit more sluggish, stabilizing about 100 ns after the directional signal crosses the DT threshold (blue horizontal line). The window (green) is placed to just overlap with the initial part of the stable region. Even a 5 ns overlap is sufficient to generate the beam event. Fig. 10 shows a 80 ns window width at a delay jumper setting of 40 ns.

The delay jumper can be placed in any one of the positions J6–J11, detailed in Table 1. It **must** be placed in one of the positions or no window will be generated, and thus no "in-beam" detections will be generated. A static 6 ns delay adds to the value chosen, and is timed from the crossing of the DT threshold. The 40 ns position is a decent nominal delay (46 ns total). Figure 11 helps locate the jumper positions.

Next is the window width, which is set by a potentiometer accessible from the front panel. The setting can be monitored either by measuring the resistance between two test-point jacks on the front panel, or by watching the output signal on the BEAM WINDOW BNC output on the discriminator front panel. If using the resistance, **power to the unit must be off** to make this measurement. The setting is sensitive to the polarity of the DVM leads. For consistent results, put the red (+) lead in the red jack (TP21 on the PCB) and

Table 1: Delay Jumper Settings

Jumper Position	delay (ns)
J6	0
J7	20
J8	40
J9	60
J10	80
J11	100

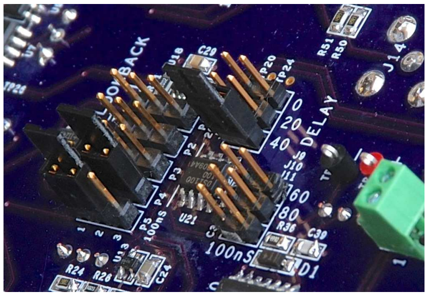

Figure 11: Discriminator board jumper locations. The rightmost two banks contain the (split) sequence from J6–J11. The pictured configuration has the jumper block in J8, corresponding to a 40 ns delay. The leftmost bank is for the lookback time, P1–P5. As shown, one jumper is engaged in the P5 jumper (100 ns lookback), while a second is engaged in P4 for 80 ns lookback (see Fig 12 for valid configurations).

the black $(-)$ lead in the black jack (TP22 on the PCB). Each unit ships with a characterization document, which includes a table for window width as a function of the potentiometer (R35) resistance. Table 2 gives a crude (generic) example of the mapping.

A reasonable fit to the initial range (50–100 ns) is $R = (w-30)/18.8$, or $w = 30 + 18.8R$, where w is the width in nanoseconds, and R is the resistance in kilo-ohms.

Finally, the judgment of whether the edge is "clean" is accomplished by a pair of controls, one being the "DC" threshold (not to be confused with direct current!) and the other being a jumper bank to control timing. The DC threshold is set on the front panel and is easily monitored from the front panel monitor connection (labeled THRESHOLD). The associated jumper (called the **lookback jumper**) is accessed inside the box, and can be configured as displayed in Figure 12.

When no lookback jumpers are present (i.e., not spanning adjacent pins—danglers are okay), this feature is effectively disabled: a window will be generated any time the directional signal exceeds the DT threshold.

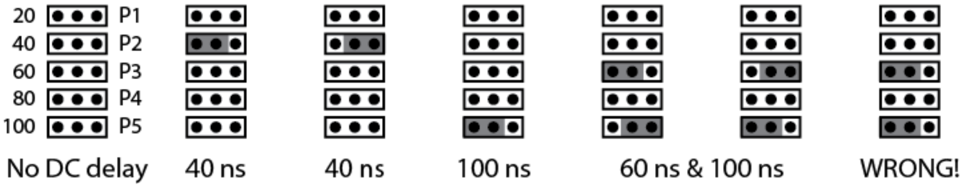

Figure 12: Example lookback jumper configurations for the lookback (DC) delay. If no jumpers are present (or dangling to the side), then this action is effectively disabled. A single jumper can occupy any delay from 20–100 ns in either the left or right columns. Up to two jumpers can be used for a combined effect, but only if their columns differ. The last example to the right is invalid.

Table 2: Window Width Settings

Window Width (ns)	TP22 to TP21 Resistance (kΩ)
50	1.05
60	1.60
70	2.15
80	2.65
90	3.15
100	3.7
150	6.3
200	9.0
250	11.7
300	14.4

If a single jumper is present—in either column—the circuit will effectively produce a window only if the signal was "clean" (weaker than the DC threshold) X ns prior, where X is the jumper position. Because of the finite rise time of normal pulses (50 ns in Rev4/5/6 builds), putting the lookback jumper in P1 or sometimes P2 (20 or 40 ns) prevents window formation (thus beam detection), because the signal will not be in a "low" state so shortly before the DT threshold was crossed. It is possible to add a second jumper, provided that *it is in a different column* from the first. The function in this case is a demand that the signal was clear X ns ago AND Y ns ago. In effect, this helps establish that the signal was indeed low and *flat* before the pulse arrived. In practice, a single jumper adequately eliminates leading-edge crud, but the second jumper adds stringency if needed for extremely busy skies.

It is important to note that due to the slow rising edge of DME pulses, DME transmissions from within the protected zone **will not be regarded as such** when any lookback jumper is present. Thus we sacrifice the ability to use DME as in-beam indicators for the sake of greater suppression of spurious "in beam" triggers. For remote areas like Mauna Kea, this slightly lowers the total rate of triggerable signals one might expect to see, and since false triggers from heavy traffic is much less a problem, it may be best to operate without this condition (allow DME in-beam triggers). In high-traffic areas, signals are so abundant that DME signals are superfluous, in which case the extra demand on pulse shape can lead to a significant reduction of false triggers.

5 Signal Noise Safety Feature

The DC (directional clear) threshold performs a second function, by ensuring that the RF background detected by the system does not creep above a desirable level. For this, we have an integrator on the circuit board with a 10 second time constant so that if the signal spends a significant (tunable) fraction of time above (lower voltage than) the DC threshold level, the shutter will be forced closed with a special code (9998), updating the condition once per second. Simultaneously, any detected airplane will be tagged with the DIREC SAT condition, for additional diagnostic information and further assurance that the shutter will remain closed in this condition.

The duty cycle at which the system triggers is set by the "BACKGROUND SAT" potentiometer on the front panel. The value can be monitored on the front-panel BNC by the same name. If it reads 1 V, for instance, then if the signal spends more than 20% (1 out of 5 V) of its time seeing more power than the DC threshold, the background saturation flag gets set, and the shutter closes. On a typical unit, the background levels measure about 1.95 V, with approximately 25 mV RMS. Thus if DC is set to 1.8 V, we require a 6-σ excursion to cross DC. And if the SAT level is 1 V, it would have to spend 20% of its time with more

Figure 13: Decoder front panel controls, including the beam sensitivity adjustment and the OMNI enable switch. Here, the system is set for 3 events per 10 s, and the omni logging is enabled. The knob at left controls which signals/conditions the speaker reports.

power than this, which should not happen unless something is wrong. Even in the presence of tons of signal traffic, we do not exceed this criterion when DC is set to 1.8 V. The function may be tested by moving DC up toward 2.0 V. At some point, the background saturation condition will trip and close the shutter.

6 Beam Events per Interval

On the front panel of the decoder box is a knob for adjusting the sensitivity of the shutter to beam events. The rationale is that TBAD occasionally will report false triggers of in-beam activity—especially in crowded airspace—due to signal overlap and multipath reflections. The window function and DC (directional clear) lookback function significantly reduce the frequency of such events, but they still happen. We therefore have the option to require multiple occurrences per count interval (10 seconds) in order to shutter the laser. We call this knob "NB" for number of beam events, or more phonetically, in-beam (NB) events.

The count interval is set to be the same as the shutter closure interval of 10 s. The adjustment knob on the front of the decoder box (Figure 13) specifies how many beam events must be accumulated within this sliding interval in order to cause the shutter to close. At its minimum setting (zero), the shutter will *always* be closed, regardless of signal activity. This provides a way to manually test shutter activation. At "each event," (one), the shutter will close any time it classifies an event as being "in-beam." From there, higher numbers correspond to more stringent requirements for shutter closure. For instance, at a setting of 20, one would need to see 20 "in-beam" events in the last 10 seconds in order to shutter the laser.

For the busiest skies, a dipswitch option on the decoder circuit board "boosts" the front panel knob setting by a factor of ten (until saturating at 255). Thus one may even require, for instance, 150 events per ten seconds to trigger the beam closure by setting the front panel knob to 15 and the dipswitch (position 1 of 5) to "ON." Note that the NB knob must be changed after the dipswitch is changed in order to take effect—even if it's a momentary change to an adjacent position.

For Mauna Kea, we see typical ModeA/C/S transmissions at rates of 3–7 per second. In a 10 second interval, this becomes 30–70. By setting the knob to 3, for instance, one would lose up to one second of reaction time on a real beam-crosser. An airplane at 30,000 ft flying at 500 knots has an angular speed at zenith of 3°/s from Mauna Kea, so the beam intrusion would be cut by about 3° (12° → 9°, for instance). In busy areas, count rates exceed 100 Hz, in which case even a setting like 40 events per 10 second interval costs less than half a second in reaction time.

7 Glitch Immunity

TBAD's success comes from the fact that it is fundamentally performing a very simple job, asking the question: is there RF power at 1090 MHz that is either strong enough and in the main beam of the antenna, or so strong that we should simply shut down? But this simplicity invites false alarms from static discharges and the like. The TBAD firmware is fashioned to ignore signals deviating substantially from valid airplane-sourced patterns. Keep in mind that the leading edge of *any* received signal is interpreted by the system as either the first framing pulse (F1) from Mode-A/C transmissions or as the start of a Mode-S preamble (see Sec. 1.2). The decision tree goes as follows:

1. TBAD ignores any signals weaker than the DT threshold (or OT if enabled, as in Sec. 8; technically ignores signals only if also weaker than OS and DS, but it would be unusual for these to be set more sensitive than DT).

2. A single F1 pulse with nothing following (no following transitions) is ignored and not reported via serial transmission as long as the three signal flags are not set (see next point).

3. A single F1 pulse triggering the BT, OS, or DS thresholds will be reported, but with no shutter reaction (the Python parsing program calls this a “glitch”).

4. A single pulse that is above-threshold longer than a few microseconds not only is counted as F1, but also as the C1 bit, since the algorithm to differentiate a Mode-S preamble from Mode-A/C codes incorporates one sample—rather than the usual edge-sense (see Sec. 1.2). These instances are reported, but do not cause a shutter closure (the Python parsing program calls this a “pulse?”).

5. Even given these protections, in some environments the dominant source of false triggers appear to come from static discharge (for instance during activity within the telescope/TBAD enclosure) in the rare case when the discharge produces a second transition within the $\sim 20\,\mu$s Mode-A/C capture window. This can be defeated by engaging the decoder PCB dipswitch (position 2, to “ON”), in which case OS and DS conditions showing a single bit (in the four digits together with X and F2) are reported but no shutter action taken.

8 Data Mode Settings

There are two switches to control data generating/handling behavior of the TBAD system. The one on the decoder box (figure 13) decides whether non-threatening aircraft tripping the OT (omni) threshold will be recorded and transmitted to the computer for logging. Each log event takes about 1.4 ms at 115,200 baud during which time the device is busy, so that it may miss a critical event. This is hardly a problem in remote regions with sparse air traffic, but could become an issue in congested airspace. In such a case, one achieves a greater safety assurance when omni-logging is disabled.

The other switch is on the discriminator box, and it determines which data will be sent to the decoder for decoding the pulse pattern. Either the DT comparator is used (labeled DIRECT ONLY on box) or an OR combination of the DT and the OT comparators (labeled COMBINED) is used to generate the data pulse pattern. The decoder then processes the signal whenever a trigger criterion is activated. The rationale for having a choice is that the DT comparator may not always be activated by signals that we only see via the omni-log (OT) criterion. So we need to provide the OT signal for decoding in these cases. When triggering on an in-beam detection (or even either saturation criteria), we are virtually guaranteed that the DT signal will be strong enough to present a valid pulse sequence. But the omni signal is not guaranteed to be above the OT threshold in these cases, so we must always at least accept the DT pulses. A table can clarify the interaction:

Figure 14: Decoder front panel LEDs.

Decoder/Discrim.	behavior	comment
Omni Enable / COMBINED	logs omni; uses DT|OT data	proper setting for logging all activity
Omni Enable / DIRECT ONLY	logs omni; uses DT data	unusual setting; could lose data
Omni Disable / COMBINED	no omni log; uses DT|OT data	unusual setting, but harmless
Omni Disable / DIRECT ONLY	no omni log; uses DT data	proper setting for safety (not data)

Most TBAD operations have been in the Omni Enable and COMBINED data modes. This is appropriate when one is using TBAD in data-collection mode. In high-traffic areas, it may be advisable to operate with Omni logging disabled and using only directional data. Doing so will place less burden on the microcontroller and serial communication, opening up more time for the system to respond to real threats. But there is no reason for Mauna Kea to chose anything other than omni-enabled and combined data. TBAD used to operate at 9600 baud, in which case the blocking behavior of serial transmissions stood a much greater chance of preventing reaction to important signals. Now at 115,200 baud, this is much less worrisome. Indeed, the much higher chatter over Apache Point or even UCSD for that matter does not appear to impact TBAD's ability to hear the important signals amidst the chatter.

9 Front Panel Indicator Lights

The front panel of the decoder box has a series of status LEDs. These offer useful diagnostics, especially in testing mode, when all the equipment is sitting together.

Figure 14 shows the arrangement of LEDs, together with the system power switch. In normal, quiet operation, the power and no-fault LEDs should be green. No-fault indicates that the discriminator box is drawing the expected amount of current ($+20\%$, -10%), and that the communications lines between the discriminator and decoder are operational. The shutter-closed LED illuminates red when the shutter is closed, which generally happens on power-up. Any time the shutter is closed, it will stay so for a minimum of 10 seconds. The rightmost three LEDs blink when a closure condition is sensed by the antenna/discriminator. If the Directional or Omni (yellow) LEDs blink even once, the shutter will close. The Beam LED may blink without a shutter closure, according to the knob setting for how many beam events may be tolerated in the last ten seconds before shutting off the laser.

10 Communication

The output serial transmission is 115,200 bits per second, 8 data bits, no parity, one stop bit (115200 8N1). The RJ-45 jack is configured with ground on pin 6 and transmission on pin 5. At present the communication is one-way, but this may change in future upgrades. The pinout is shown in Fig. 15. A standard ethernet patch cable successfully connects to standard terminal servers possessing RJ-45 serial ports.

There are two principal types of TBAD serial transmissions:

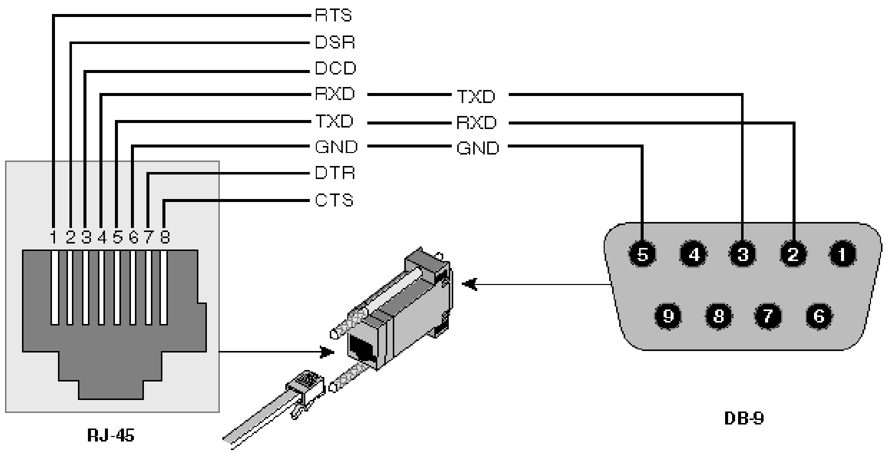

Figure 15: Pinout of RJ-45 socket, and conversion to standard DB-9 connector, if needed.

1. 4-byte “payloads” accommodating Mode-A, Mode-C, DME, spurious signals, and messages;
2. 14 or 28-byte payloads accommodating Mode-S and ADS-B messages, respectively

Each event is packaged with additional informational bytes, totaling $12+N$, where N is the payload size [4, 14, 28]. All bytes are ASCII code, and the fields are as follows:

bytes	field	values	interpretation
0	shutter condition	i, o, s	informational, open, shutter closed
1–N	squawk/alt code	0–9, 0–F	4-digit decimal code, or 14 or 28-digit hex code
$N+1$	omni saturation?	’.’, O	O if saturation
$N+2$	directional saturation?	’.’, D	D if saturation
$N+3$	in beam?	’.’, B	B if considered to be in central beam
$N+4$	knob/power condition	a–p; A–P	NB knob position; upper case if current is good
$N+5$	first framing pulse	’.’, F	should always be present
$N+6$	X-bit in code	’.’, X	don’t expect, except TSIM
$N+7$	final framing pulse	’.’, F	legitimate squawk code will have, TSIM not
$N+8,9$	checksum	0-F	2-digit hex, sum of bytes 0 to $N+7$
$N+10,11$	termination	\r\n	carriage return plus line feed

The parsing program, written in Python, does some interpretation of the code and packages it in a more verbose time-stamped sequence in the log file. An example:

```
2010-05-09 18:18:12.236 s0730.DBLF.FFB 1400
2010-05-09 18:18:12.263 s1200..BLF.FDE ----- VFR
2010-05-09 18:18:12.289 s0730..BLF.FE5 1400
2010-05-09 18:18:12.321 s1200ODBLF.F15 ----- VFR
```

These codes, spanning less than 0.1 seconds in time, indicate a nearby airplane at an altitude of 1400 feet (code = 0730) squawking 1200 (VFR default code). The shutter is closed, the “in-beam” criterion was triggered for all, and the directional and omni signals hovered around saturation. The “L” indicates that the knob was in the 12th position (20 events/10 s), and the fact that it is upper case indicates proper system current. The first and last framing pulses are present, but no "X" bit (normal).

The website: `http://www.airsport-corp.com/modec.htm` has a useful description of the transponder codes and specifications. The mapping of codes to altitude is at `http://www.airsport-corp.com/modecascii.txt`.

Additional informational codes are provided to indicate when the shutter re-opens and if the unit closes the shutter due to perceived power or high-background failures. Since these codes are not associated with real data, we fill the 4-digit code with 8's or 9's, since these cannot happen in the 3-bit Mode-A/C codes. A leading "o" or "s" indicates shutter open or closed in the normal sense. Examples of the three special codes are `o8888...D...A7` (shutter open event), `s9999...d...CF` (power/current bad), and `s9998...D...AE` (excessive background). The "D" in each case indicates the knob set to 3 events per 10 s.

The Mode-S and ADS-B codes are similar except that the X-bit and final framing pulse are guaranteed to be absent (just dots). The Python program provides some context: the first five bits of the transmission are the data format (DF), so a DF-XX code is listed on the line, where XX can range from 00 to 31. Depending on the Python code used, a line may follow with additional decoding of the transmission.

Finally, if the system is idle for a minute (will often happen in quiet skies), a keep-alive is transmitted, with the code `i0000...D...81`.

Future firmware modifications may extend the functionality of TBAD. In this case, it may be advisable to acquire a PIC microcontroller communication module to work with MPLab software so that firmware upgrades can be made on site.

11 Audio Output

Starting with Rev5 an audio output feature can emit a short (13 ms) burst of 1 kHz tone whenever a signal type of the user's choice is generated by the discriminator. A knob on the decoder front panel allows choices of OFF, BT, DS, and OS. The speaker is adjacent to the knob and projects out the front of the decoder. In OFF mode, the speaker is silent. In BT mode any time the IN BEAM LED flashes, the speaker chirps. Likewise for DS and OS selections. The effect is similar to that of a Geiger counter. At around 70 Hz repetition rate, the speaker tone becomes approximately continuous. This feature enhances understanding of how TBAD is reacting during sky tests or dome tests when a user is present. It is recommended that the speaker be in the OFF position when not in use, otherwise the power drain and overuse may result in longer-term problems.

12 Decoder Configuration

One jumper and one dipswitch are used to control the behavior of TBAD, described below.

12.1 Manual Mode

The microcontroller adds enough smarts to TBAD to be able to ignore glitches (single-pulse events) and to require multiple "in-beam" events before closing the shutter. Most of the signal-level decisions, however, are already made by analog stages at the discriminator. The decoder unit therefore has a jumper setting that allows shutter control based purely on the B (in beam), O (OMNI saturation), or D (DIREC saturation). This has never been used in practice, but is available in the event that one may want to bypass the microcontroller. The jumper is just to the upper left of center in Figure 17.

12.2 Dipswitch Control

Starting with Rev5, the TBAD decoder unit incorporates a bank of dipswitches that can be used to set levels of five otherwise unused microcontroller digital input pins. As of this writing, the first two positions have assigned functions.

The first (position 1), as discussed in Sec. 6, optionally multiplies the number of beam events per 10 second interval by ten, so that, for instance a knob setting of 8 results in 80 events per 10 seconds. This scaling applies until knob position 30, at which point the count saturates at 255.

The second (position 2), as discussed in Sec. 7, allows TBAD to ignore single-bit (double-pulse) saturation events to reduce immunity to static discharge events. Legitimate aircraft transmissions do not have this characteristic, so will be processed normally. The only exception is when aircraft deliberately transmit "zeros" in Mode A/C, as sometimes happens for military craft. TBAD win this case will still count "in-beam" instances, but will ignore saturation conditions in these cases. The other notable condition is DME pulses. But because the C1 bit is a composite of time-domain sampling plus edge-trigger, a typical DME pulse has both B1 (or D1) and C1 bits engaged, in addition to the F1 (first framing pulse; see Sec. 1.2). So DME saturations are still permitted by this mode.

13 Dense Traffic Recommendations

When operating in areas experiencing high air traffic, various settings/options may be employed to reduce false alarms, and prevent excessive log activity (which may preoccupy TBAD and impede safety functions).

Note that in very dense environments TBAD spends substantial time sending serial transmissions. During transmission, TBAD is insensitive to new signals. At 1.4 ms per Mode A/C transmission, the theoretical limit to TBAD response is 700 Hz, but this would require conspiratorial spacing of signals. In practice, 300 Hz is an observed limit. Another possibly unexpected side effect occurs: the action of ticking away time is interrupted as well, so that TBAD's notion of how much time has passed is impacted. Time flies when TBAD is "having fun." This means that the 10 s interval used to count beam events *and* to hold the shutter closed may stretch to 12 s, 17 s, or even longer in extreme cases. Fortunately, this "failure" is failsafe on both counts: it is easier to slip *NB* "in-beam" events into a longer interval, and the shutter is held closed for longer. In many cases, most of the activity is from OMNI log activity, rather than from closure threats.

The following actions can improve the reliability and safety of TBAD in dense environments:

1. Turn OT to a lower voltage so that OMNI logging is restricted to closer traffic. Each 0.15 V corresponds to a factor of two in distance sensitivity.

2. Turn the NB knob to a high number to reduce "in-beam" false alarms. The decoder PCB dipswitch (pos. 1 to ON) will boost the numbers by a factor of 10 (saturating at 255) for increased range. Evaluate transmission rates at your location (note diurnal variation) in deciding how much delay you can tolerate for higher NB settings.

3. Employ the lookback jumpers: one jumper at 100 ns is good, and adding another at 80 ns (in a different column, see Fig. 12) is even more stringent.

4. If in-beam false alarms are frequent, the BT threshold can be increased to make this criterion more selective. The effective beam size on the sky will be reduced. Too large a change may begin to eliminate bona-fide in-beam events, so a recommended practice is to increase 0.05 V and evaluate performance with respect to in-beam detections and false alarms, and repeat until the false alarms disappear or the legitimate in-beam detections are impacted.

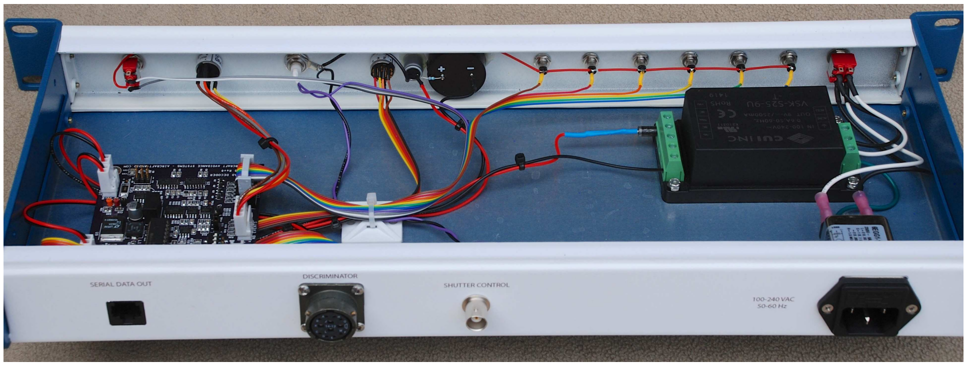

Figure 16: Inside the decoder box is a power supply (right) and the PCB containing the microcontroller (left).

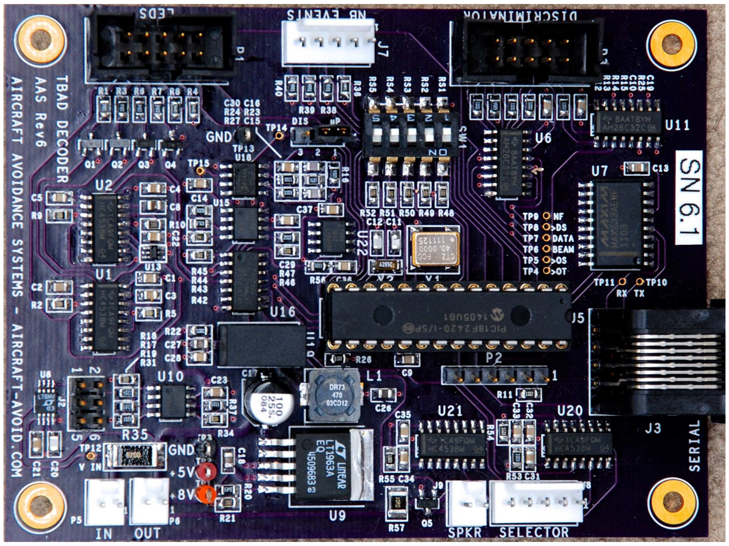

Figure 17: Decoder PCB. Note the jumper at top center establishing whether the microcontroller or the unprocessed discriminator signals control the shutter. Adjacent to the jumper is a 5-position dipswitch—two of which are utilized at present, the others available for expansion options.

5. False alarms may also be mitigated by reducing the BEAM WINDOW width (Sec. 4) in increments of 10 ns, following an iterative procedure as in recommendation #4 above.
6. Disable OMNI logging on the decoder front panel and move the data switch to DIRECT ONLY on the discriminator to completely eliminate OMNI logging for maximum safety assurance.

14 Electronics

For completeness, here are pictures of the electronics implementations and populated printed circuit boards (PCBs) in the two boxes. First, the decoder internals are shown in Figures 16 and 17.

The PCB in the discriminator is shown in Figure 18. Note that the back (unseen) side of the board contains potentiometers, BNC connectors, and switches that are accessed through the front panel of the discriminator box (Figure 5).

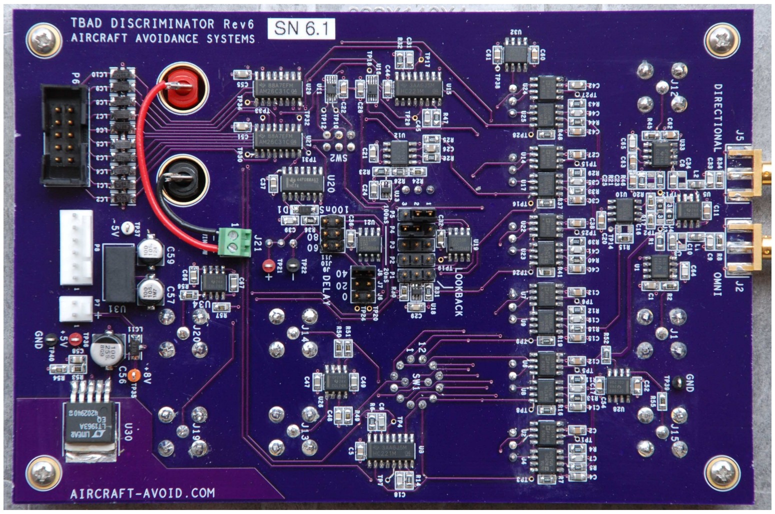

Figure 18: Discriminator PCB. OMNI and DIREC signals come in at left. The P1–P5 jumpers for the lookback function (Figure 12) are in the center of the board (here showing 80 ns and 100 ns lookback jumpers engaged). The window delay is the broken/split jumper bank just to the left of this (here set to 40 ns, or J8).

Also, the inside of the decoder box contains the “RF Legos” to convert 1090 MHz signals from the antennas into power levels, shown in Figure 19.

15 Physical Specifications

15.1 Antenna Assembly

- Mass: 3.93 kg
- Outer dimensions: 24.82 in (630.4 mm) wide; 24.29 in (617.0 mm) tall; 1.375 in (34.9 mm)
- Mounting plate is 0.125 in (3.17 mm) thick, with extensions 0.15 in (3.8 mm) forward and 1.1 in (27.9 mm) rearward
- Two SMA female connections presented for connection to discriminator

15.2 Discriminator

- Mass: 4.35 kg
- Box dimensions (without lid): 8 × 10 × 4 in (203 × 254 × 102 mm) width by length by depth
- Box dimensions with lid: 8.45 × 10.45 × 4.35 in (215 × 265 × 110 mm); hinges add 0.425 in (11 mm); latch adds 0.2 in (6 mm); connector protrudes 0.2 in (6 mm)
- Box mounting flange: essentially 11 × 4.75 × 0.08 in (279 × 121 × 2 mm) plate; Four 0.2 in (5.1 mm) holes on symmetrically centered rectangular pattern 10.5 × 4.125 in (267 × 105 mm) in size.
- Two SMA female conenctions presented for connection to antenna; 12-pin circular connector for connection to decoder

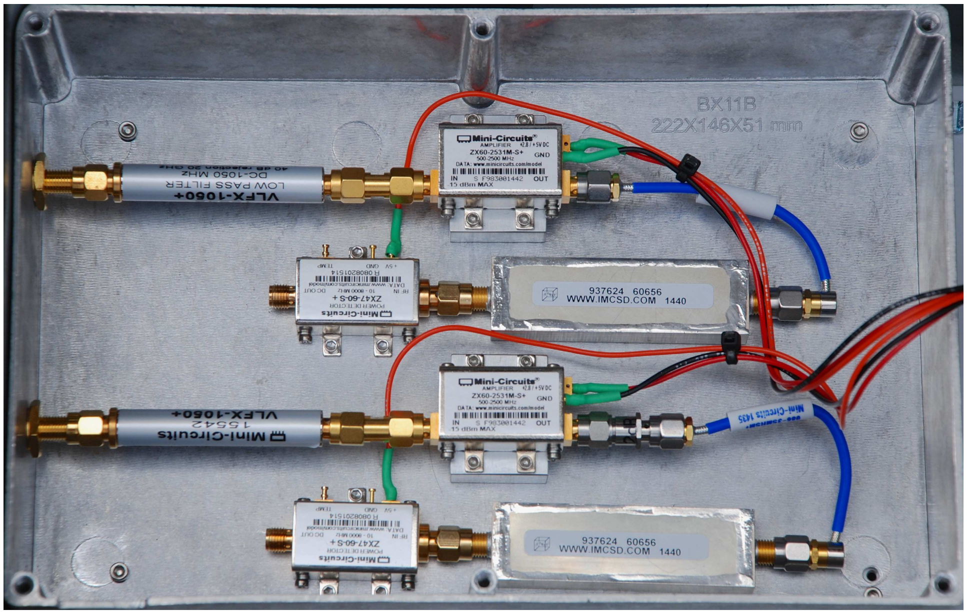

Figure 19: Identical RF chains for the omni and directional antenna signals. Each chain consists of a low-pass filter, a 37.5 dB amplifier, a narrow-band filter, and the power detector.

15.3 Decoder

- Mass: 1.52 kg
- 1U (1.75 in) 19-inch rack mount unit, 7.5 in deep. Body 17 × 7 × 1.75 in (432 × 178 × 44 mm)
- Connectors: 12-pin circular connector for connection to discriminator; 2 BNC female for shutter control; RJ-45 serial output
- IEC-C14 AC input with integrated glass fuse (250 mA)
- 100–240 VAC input; 50–60 Hz
- 7.8 W power consumption; 13.2 VA (when discriminator connected)

15.4 Interconnect Cable

- 6 pairs, AWG22, twisted, foil-shielded, with drain wire
- power pair (red/black) carries 0.73 A and develops voltage drop 0.0117 V/ft (0.0385 V/m)
- stright-through pin-to-pin connection, according to table below

color	function	pin
brown	OS+	A
black	OS–	B
green	BT+	C
black	BT–	D
white	DATA+	E
black	DATA–	F
yellow	DS+	G
black	DS–	H
blue	OT+	J
black	OT–	K
red	+8 V	L
black	ground	M

15.5 TSIM

- Mass: 3.07 kg
- Box dimensions (without lid): 6 × 8 × 4 in (152 × 203 × 102 mm) width by length by depth
- Box dimensions with lid: 6.45 × 8.45 × 4.30 in (164 × 215 × 109 mm); hinges add 0.425 in (11 mm); latch adds 0.2 in (6 mm); BNC connector protrudes 0.51 in (13 mm) further.
- Box mounting flange: essentially 9 × 4.75 × 0.08 in (229 × 121 × 2 mm) plate; Four 0.2 in (5.1 mm) holes on symmetrically centered rectangular pattern 8.5 × 4.125 in (216 × 105 mm) in size.
- IEC-C14 AC input; on-board fuse, 200 mA
- 105–125 VAC input, 47–420 Hz, single phase
- 4 W power consumption; 6.5 VA

15.6 TSIM Antenna Assembly

- Mass: 460 g
- Outer dimensions: 8 × 8 × 1.525 in (203 × 203 × 38.7 mm)
- Mounting plate is 0.125 in (3.17 mm) thick, with extensions 0.15 in (3.8 mm) forward and 1.25 in (31.8 mm) rearward
- Mounting holes around periphery 0.169 in (4.29 mm) diameter, 0.25 in (6.34 mm) in from edge, ±1.75 in (±44.4 mm)

installation requirements; verifiable/calibratable measures