

Automated RF Attenuation Managing System

*Furkan Kamil TUT¹, İrem ULUSOY KANDEMİR¹, Mustafa Melik GECGEL¹,
Mahmut Kursad ARPACIOGLU¹*

*¹Turkish Aerospace Inc. Fethiye Mah. Havacilik Bulvari No:16 Kahramankazan/Ankara/Türkiye
furkankamil.tut@tai.com.tr, irem.ulusoykandemir@tai.com.tr, melik.gecgel@tai.com.tr,
mkarpacioglu@tai.com.tr*

Abstract

During flight tests, sudden RF power spikes from close-range tracking or multipath effects risk saturating and permanently damaging sensitive receiver front-ends. To mitigate this without compromising long-range signal integrity, an automated RF protection system was developed. This architecture monitors dual telemetry channels using logarithmic RF detectors and a 16-bit DAQ module. Applying a calibrated voltage-to-dBm model, the system estimates RF power in real time. A deterministic control logic then automatically actuates RF switches to select between attenuated and direct signal paths, ensuring fail-safe operation during network anomalies. Laboratory validation demonstrated a conversion accuracy of ± 0.9 dBm across the -60 to 0 dBm range, with a protective response time under 300ms. This adaptive approach successfully prevents receiver saturation in close-proximity scenarios while preserving the critical link budget needed for extended-range telemetry operations.

Key words: telemetry, RF protection, automatic attenuation, fail-safe, link budget optimization

Introduction

Turkish Aerospace Inc. (TA) is the leading center for the design and production of air platforms in Türkiye. The company develops a diverse fleet across three main classes. Its fixed-wing portfolio features KAAN, HÜRJET, and HÜRKUŞ/HÜRKUŞ II. The rotary-wing lineup comprises ATAK, ATAK-II, and GÖKBAY, while the unmanned aerial vehicle (UAV) segment includes ANKA, ANKA-III, AKSUNGUR, and KİHA. Telemetry systems are essential for actively monitoring these platforms during their flight test phases. The study presented in this paper focuses on a prototype of HÜRJET, TA's prominent indigenous jet trainer.

HÜRJET is one of TA's flagship programs, developed for both lead-in fighter training and light attack missions. Its flight test campaign covers a broad operational envelope, with test sorties including high-speed low-altitude passes, touch-and-go maneuvers, and long-range cruise flights. Each of these phases generates demanding requirements for the telemetry link. Initially, the supporting ground station was equipped with a fixed RF attenuator at the receiver input to protect the front-end during close-range passes. However, this static configuration also reduced the available link margin during long-range tracking. This trade-off motivated the development of an automated system that dynamically inserts attenuation only when measured RF levels dictate, and bypasses

it as soon as the signal drops back to safe values.

A telemetry receiver chain is typically built around a low-noise amplifier (LNA), followed by gain, filtering, and downconversion stages. The receiver chain is designed to recover weak telemetry signals during long-range operation. Therefore, receiver front ends are optimized for maximum sensitivity to meet the critical requirements of the telemetry link budget [1]. That same sensitivity becomes a liability at close range. As the target aircraft approaches the ground antenna, the input level can climb rapidly. Multipath returns adding coherently at the feed produce the same effect. During close-range passes or sustained multipath conditions, the receiver input can rise from typical weak-signal levels toward the compression region. Our design manages these mission-level power excursions. It actively forces the receiver into a protected state the moment input levels exceed a configured threshold.

Receiver front-end components are usually optimized for weak-signal sensitivity rather than high input power. Depending on the selected LNA and receiver front-end design, excessive input levels may drive the first active stage into compression or expose it to overstress. Therefore, the allowable input level must be treated as a receiver-specific design limit. [2]. At higher input levels, the first active stage may enter nonlinear operation, resulting in gain

compression, distortion, and intermodulation products. If the excursion is large or sustained, the device may suffer permanent damage. Even brief saturation costs several seconds of clean data. Such data loss is unacceptable during critical envelope expansion phases. Low-altitude passes, touch-and-go maneuvers, and close-range orbits demand uninterrupted telemetry.

The traditional approach is to place a fixed attenuator between the antenna feed and the receiver. Attenuation values of 20 to 30 dB are commonly used to reduce the risk of receiver front-end saturation during close-range operation. However, this protection is applied continuously and therefore introduces a permanent link-budget penalty. At long ranges, free-space path loss has already pulled the signal close to the noise floor. An extra 30 dB of attenuation is often the difference between a healthy link and an unusable one. The fixed-pad approach therefore caps the station's effective tracking distance. It also degrades data quality on the longer legs of every mission.

The system presented in this paper resolves that compromise. Two logarithmic detectors sample the RF level on the main channels. Detector voltage is mapped to dBm through a calibrated polynomial. The system evaluates the switching condition at approximately 200 ms intervals. It then uses the real-time RF power estimates to determine whether the attenuators should be inserted or bypassed. A fail-safe layer takes care of every failure that might leave the receiver exposed. If the network, the host software, or the host PC itself fails, the receiver ends up in the protected configuration. The system was qualified in the laboratory and then flown on operational test missions. In the field, it recovered the lost link margin without giving up close-range protection.

Problem Definition and Motivation

In a typical pass, the ground station antenna tracks the aircraft using autotrack error signals derived from the telemetry carrier. Slant range varies continuously as the aircraft maneuvers. The received power at the antenna output therefore varies as well. On the longest legs, the receiver input may sit at -70 to -60 dBm. During takeoff, landing, or low-altitude passes, the same input can climb to -10 dBm or higher.

A receiver that has to handle a 60dB swing without losing the signal at one end or saturating at the other is asking too much of any single fixed configuration. The conventional fixed-pad approach simply chooses the safer end. It sacrifices long-range performance to guarantee close-range protection. With a 30dB attenuator in line, a -60 dBm signal at the antenna feed is reduced to -90 dBm at the receiver input. This reduction decreases the predicted receiver input level and therefore reduces the available SNR/link margin. Since telemetry data quality is commonly evaluated against the SNR or E_b/N_0 required for a target BER, this loss can degrade extended-range telemetry performance [3].

Multipath further increases the uncertainty in the received RF level. Test ranges often include terrain, buildings, or reflective ground surfaces. Reflected copies of the direct signal can add at the feed. The result is transient spikes that exceed the free-space prediction by several dB. These spikes are unpredictable in timing and amplitude. A static attenuator cannot react to them. Setting the pad high enough to absorb the worst case means accepting a chronic link penalty during normal operation.

Our design objective followed directly from this picture. We needed a system that inserts and removes attenuation based on the actual received power. The accuracy had to be sufficient to separate safe from hazardous levels. The response time had to be short enough to keep the front-end intact during a fast power excursion. The fail-safe behavior had to cover every credible fault. Threshold values had to be configurable, since different receivers and different missions need different operating points.

System Architecture

Fig. 1 shows the architecture. A remote I/O controller sits at the center of the system. It provides four 16-bit analog inputs, four digital outputs, and an Ethernet interface that speaks Modbus TCP. We chose this controller for its industrial reliability. Its built-in network watchdog and configurable safe-state outputs were equally important for our design. Both are central to the fail-safe strategy described later.




Fig. 1. Overall system architecture showing the RF signal paths, power detectors, switches, attenuators, and the control unit integrated on the antenna pedestal.

Two RF power detectors are connected to coupler outputs on the main telemetry channels (RF1 and RF2). Each detector outputs a DC voltage proportional to the logarithm of the input RF power. That voltage is sampled by the controller's analog inputs. The controller's digital outputs drive four RF switches: two on the main channels and two on their associated autotrack channels. Each switch selects between a direct path and a path that passes through a fixed attenuator. In the normal state, all four switches select the direct RF path. The receivers see the full feed signal. When the measured RF power on either main channel passes a configurable threshold, all four switches go to the attenuated path together. This applies the same protection command to the main receiver paths and the associated autotrack paths.

The host application is written in C#. The application queries the controller over Modbus TCP at approximately 200 ms intervals. On each cycle, it reads the analog inputs, converts the voltages to dBm, evaluates the threshold, and writes the digital outputs accordingly. Voltage readings, dBm estimates, switch states, and timestamps are all logged to CSV for post-flight review.

Hardware Components

The hardware consists of three main parts. These are the RF power detectors, the RF switches, and the I/O controller. The I/O controller is installed in the same system

enclosure together with its power supply and a temperature sensor.

The detectors are logarithmic types. They cover roughly -60 dBm to 0 dBm at the input. The output ranges from approximately 0.4 V to 1.6 V across that range. The logarithmic transfer function compresses the voltage swing in a way that fits the 0 – 10 V analog input span without losing useful resolution at either end. We use one detector per main channel.

The four RF switches are SPDT latching units. Insertion loss is on the order of 0.1 dB. Isolation is high, and the operating range goes from DC up to 6 GHz. That is comfortably above the telemetry bands of interest. Two switches sit on the main channels, two on the autotrack paths. Each switch is actuated by a 24 VDC line driven by a relay tied to the controller's digital outputs.



Fig. 2. Internal view of the RF protection control box showing the RF switches, power detector, relay modules, and wiring.

The controller provides the main interface between the RF measurement hardware, the relay drivers, and the host application. Its four analog inputs (AI0–AI3) are 16-bit over a 0–10 V span. This gives a voltage resolution near 0.15 mV. That is sufficient for resolving the detector output over the calibrated operating range. The four digital outputs (DO0–DO3) drive the relay drivers. We also wired a PT1000 RTD into one of the spare analog inputs to monitor enclosure temperature. The box sits outdoors on the antenna pedestal, so we wanted long-term thermal data on the electronics.

RF Power Measurement and Calibration

How accurately the system protects the receiver depends entirely on how accurately it converts detector voltage to dBm. The detector's transfer characteristic is approximately logarithmic. However, it deviates noticeably at the ends of its range. Linear and low-order polynomial fits did not get us within the accuracy we wanted across the full –60 to 0 dBm span.

We ran the calibration in the lab. The setup used a signal generator and a spectrum analyzer as the reference. The generator was swept in 1 dB steps from –60 dBm to 0 dBm. At each step, we read the actual power on the analyzer and the detector voltage through the controller. This built a 61-point dataset.

A least-squares polynomial fit was then applied to the data. Third- and fourth-order fits left residuals over ± 2 dBm at the extremes. We iteratively increased the polynomial order. A ninth-order fit settled the maximum residual at about ± 0.9 dBm over the full span. The conversion takes the form of (1):

$$P_{\text{dBm}} = a_9V^9 + a_8V^8 + \dots + a_1V + a_0 \quad (1)$$

V is the voltage read by the controller. The coefficients a_0 through a_9 come from the regression. Fig. 3 plots the calibration data and the fitted curve. The agreement is close throughout the range. That is what justifies the choice of order.




Fig. 3. Calibration curve: measured detector output voltage versus RF power and the fitted ninth-order polynomial.

The polynomial is evaluated on every measurement cycle inside the C# application. A ninth-order evaluation is essentially free on a modern host CPU. The fit is markedly better than what we obtained from piecewise-linear interpolation or lookup tables. The improvement is most noticeable near the upper and lower edges of the detector's range.

Control Algorithm and Protection Logic

The control logic is intentionally conservative. A hazardous level on any monitored channel switches all channels to the attenuated path. This applies to both main and autotrack paths. This approach is dictated by the hardware configuration: the main and autotrack channels share the same feed. High power on one main channel almost certainly means high power on the corresponding autotrack channel.

On every polling cycle (about 200 ms), the application reads AI0 and AI1. It converts the voltages to dBm with the calibrated polynomial. Each value is then compared to a configurable threshold T. A typical setting is $T = -20$ dBm. If both P_{RF1} and P_{RF2} are below T, DO0 through DO3 are commanded to the normal (direct path) state. If either P_{RF1} or P_{RF2} is above T, all four outputs go to the protection (attenuated path) state. The OR-logic ensures a spike on a single channel pulls the whole RF chain into protection.

A hysteresis margin keeps the switching from chattering when the RF level fluctuates near T. The system only returns to normal mode after the measured power falls below T minus the hysteresis offset for several consecutive cycles. The behavior remains stable under normal fading and multipath conditions. Fig. 4 shows the full algorithm as a flowchart.




Fig. 4. Control algorithm flowchart showing the periodic measurement, threshold evaluation, hysteresis application, and output actuation cycle.

Fail-Safe Architecture

A central design rule was that no single failure should leave the receiver unprotected. The primary fail-safe mechanism is the controller's host-connection watchdog. The controller monitors the Modbus TCP session continuously. If no valid request arrives within the configured timeout, it declares Host Connection Lost. It then drives every digital output to its preconfigured safe state. We set the safe state to protection mode. All RF switches go to the attenuated path automatically.

Whatever causes the timeout, the receiver ends up in the fully attenuated configuration within the watchdog interval. The fail-safe triggers under various fault conditions, such as physical network disconnections, host PC crashes, or software hangs.

A second layer comes from the controller's peer-to-peer (P2P) I/O mapping. P2P maps analog inputs directly to digital outputs against preconfigured thresholds. This happens entirely inside the controller and without any host involvement. The P2P path does not have access to the polynomial calibration, so its

accuracy is lower. However, it remains functional even if the host is completely offline. Treating P2P as a backup rather than the primary path keeps the calibration accuracy when the host is up. It also provides an additional protection layer when the host is unavailable.

Noise Mitigation at Low RF Levels

During field deployment, we encountered an unexpected anomaly that the laboratory data had not predicted. Below roughly -35 dBm of input power, the detector voltage stopped tracking the input. The reading sat at about 0.77 V regardless of the actual signal level. This was clearly inconsistent with the detector's logarithmic transfer curve. Probing the rail traced it to switching-mode power-supply (SMPS) noise on the 24 VDC line that fed the system off the 220 VAC mains.

The SMPS ripple and high-frequency switching artifacts were coupling into the detector's DC output through the shared distribution. The result was an effective noise floor that masked the detector's response at low signal levels. We addressed it on four fronts. First, we placed a low-dropout (LDO) linear regulator on the detector supply rail to isolate it from the SMPS. Second, we reworked the grounding into a star configuration. The analog detector ground and the digital relay-driver ground would no longer share return current. Third, we rerouted the harness to keep the low-level analog wires physically away from the power and digital control wiring. Fourth, hysteresis logic was added in the host application to prevent unstable switching near the threshold.

After these changes, the noise floor dropped substantially. The detector tracked its specified curve down to the bottom of its operating range.

Software Architecture

The host application is organized into five C# modules. ConnectionManager handles the Modbus TCP session: connect, heartbeat, reconnect, and timeout. Everything below it talks to the controller through this single layer.

RFMeasurementEngine pulls raw ADC values from AI0 and AI1. It converts them to voltages and applies the ninth-order calibration polynomial to produce dBm estimates. DecisionEngine takes those estimates, applies the threshold and the hysteresis state, and produces the binary protection-or-normal decision for each output channel.

TemperatureMonitor reads the PT1000 RTD on a third analog input. It converts the resistance to temperature using the standard PT1000 curve and logs the value. The temperature stream gives us long-term reliability data on the

pedestal-mounted electronics across seasons and operating conditions.

CsvLogger records everything in a single stream: timestamp, raw voltages, dBm values, digital output states, and enclosure temperature. The logger uses for decimal formatting. This keeps the CSV files portable across machines with different locale settings.

Laboratory Validation

Three lab tests qualified the system before we put it in the field. These covered measurement accuracy, end-to-end response time, and fail-safe behavior.

For accuracy, the signal generator was swept across -60 to 0 dBm in 1 dB steps with the system running. We compared the dBm values reported by the application against the spectrum analyzer reference at every step. The worst-case deviation was ± 0.9 dBm. This matches the calibration residual and confirms that the protection decision has enough margin to be reliable.

For response time, the generator output was stepped from -25 dBm to -15 dBm. This crosses the -20 dBm threshold. The time from the step to the digital output transition was measured. The total latency, end to end, came in under 300 ms across many trials. This includes ADC conversion, polynomial evaluation, decision logic, the Modbus write, and relay actuation. The number is dominated by the polling interval. It is acceptable given how slowly aircraft slant range changes in practical scenarios.

For fail-safe, we ran the system in normal mode (attenuators bypassed) and pulled the Ethernet cable between the host PC and the controller. The controller switched all four digital outputs to the protection state within the configured watchdog timeout, exactly as designed. We repeated the test by killing the application and by

powering down the host PC. The result was the same in both cases.

Field Test Results

After lab qualification, the system was installed on the antenna pedestal of an operational telemetry ground station. It was used during flight test missions. Fig. 5 shows the installation. The control box is mounted directly on the pedestal. This keeps the RF cables between the antenna feed and the switches and detectors short.



Fig. 5. The RF protection system control box installed on the antenna pedestal of the telemetry ground station.

The field campaign covered a range of distances and altitudes, including extended-range cruise legs. Fig. 6 plots a representative segment. The top panel shows antenna-to-aircraft distance. The middle panel shows aircraft altitude. The bottom panel shows the measured RF level on both channels.




Fig. 6. Field test data recorded during flight operations: antenna-to-aircraft distance (top), aircraft altitude (middle), and received RF power levels on channels 1-2 and Switch State (bottom).

The bottom panel of Fig. 6 makes the system's behavior easy to read. Both channels show clean step transitions of about 30 dB at distinct moments through the mission. Each step corresponds to a switching event. When the received RF power crosses the protection threshold, the attenuator drops in. The measured level then drops by roughly 30 dB. This typically happens during a close-range pass or a multipath spike. Once the level falls back below the threshold, with the hysteresis margin satisfied, the attenuator is removed. The level rises by the same step. The behavior is the same on Channel 1 and Channel 2. This confirms that all monitored paths react together.

With the automated system running, the average received signal at the receiver input during long-range tracking sat around -45 dBm. Under comparable flight profiles and atmospheric conditions, the same station with the legacy fixed 30 dB pad would have averaged around -75 dBm. That is a 30 dB deficit. It translates directly into worse SNR and more dropouts. The automated approach therefore recovered the full 30 dB of link margin at long range. It preserved exactly the same close-range protection the fixed pad provided.

Across the entire field campaign, we did not observe a single false trigger or missed event. The hysteresis stopped the system from chattering during fades. The fail-safe layer behaved correctly during the network-disconnect tests we ran between missions.

Conclusions

This paper described the design, validation, and field deployment of an automated RF attenuation management system for telemetry ground stations. The system reduces the practical trade-off between fixed front-end protection and long-range link margin by selecting the attenuation state according to the measured RF level. It does this by switching the attenuator in and out of the chain in real time, based on what the receiver is actually seeing. The main results are a calibration accuracy of ± 0.9 dBm over the -60 to 0 dBm range, an end-to-end protection response below 300 ms, and a fail-safe layer that defaults to the protected state under every fault we tested.

In the field, the system recovered roughly 30 dB of link margin at extended range compared with the fixed-attenuator baseline. Close-range protection remained intact. The whole solution is built from off-the-shelf parts. This provides a practical and modular alternative to the

dedicated automatic-protection products on the market.

Future Works

Future work will focus on adding an RF pre-filtering stage matched to the 2.2–2.4 GHz frequency band used in our telemetry operations. The receiver front-end is currently exposed not only to the desired telemetry carrier but also to out-of-band emissions from nearby radio sources, such as cellular base stations, ISM-band devices, and other co-located ground systems. Without selective filtering ahead of the protection chain, these unwanted signals contribute to the total power seen by the detector and can drive the LNA toward compression even when the telemetry carrier itself remains at a safe level. The planned solution is a tightly tuned bandpass filter inserted immediately after the antenna feed, so that only in-band energy reaches the detector and the receiver while interferers outside this band are rejected before they can degrade link performance or trigger unnecessary protection switching. Once validated on the present ground station, the same filtering strategy will be deployed on the other antenna installations where similar interference patterns have been observed during operations, providing a consistent and band-specific defense layer across the wider telemetry infrastructure.

References

- [1] Range Commanders Council, *Telemetry Standards (IRIG Standard 106-23)*, Secretariat Range Commanders Council, White Sands Missile Range, NM, 2023.
- [2] D. M. Pozar, *Microwave Engineering*, 4th ed., John Wiley & Sons, Hoboken, NJ, 2011.
- [3] Range Commanders Council, *Telemetry Applications Handbook*, RCC Document 119-06, Secretariat, Range Commanders Council, White Sands Missile Range, NM, May 2006.