

Integrating Runtime Verification into an Automated UAS Traffic Management System

Matthew Cauwels^[0000-0002-2450-812X], Abigail Hammer^[0000-0002-0340-0160],
Benjamin Hertz^[0000-0003-1627-9715], Phillip H. Jones^[0000-0002-8220-7552], and
Kristin Y. Rozier^[0000-0002-6718-2828] *

Iowa State University, Ames IA 50010, USA

{mcauwels, arhammer, benhertz, phjones, kyrozier} @iastate.edu

Abstract. Unmanned Aerial Systems (UAS) are quickly integrating into the National Air Space (NAS). With the number of registered small (under 55 pounds) UAS in the USA alone at over 1.5 million, and projected to expand rapidly, according to the Federal Aviation Administration (FAA), safety is a pressing consideration. Safe UAS integration into the NAS requires an intelligent, automated system for UAS Traffic Management (UTM). Even more than for manned aircraft, UTM must integrate runtime checks to ensure system safety, at the very least to make up for the lack of humans on board to employ the common-sense safety checks ingrained into the culture of human aviation.

We overview a candidate automated, intelligent UTM system and propose multiple integration points for runtime verification (RV) to ensure that each part of the UTM adheres to safety requirements during operation. We write, validate, and present patterns for formal requirements across multiple subsystems of this UTM framework. After encoding our requirements as flight-certifiable runtime observers in the R2U2 RV engine, we execute them in simulation across multiple real-life test flights supplemented with simulated data to cover additional cases that did not occur in flight. Lessons learned accompany an analysis of the efficacy and performance of RV integration into the UTM framework.

Keywords: UAS, UTM, Runtime Verification, R2U2.

1 Introduction

The Federal Aviation Administration (FAA) forecasts Unmanned Aerial System (UAS) numbers to continue to “expand rapidly” over the next 20 years with over 90% of the growth from consumer-grade or professional-grade (non-model) UAS used for commercial or research purposes [5]. Given the considerable traffic this will generate and the pressing concern for safe integration into the National Air Space (NAS), additional traffic management is required on top of current safety regulations [6]. A recent candidate for an intelligent, automated UAS Traffic Management (UTM) system addresses these concerns [25].

* Supported by NSF CAREER Award CNS-1552934 and NSF PFI:BIC grant CNS-1257011.
Reproducibility artifacts: <http://temporallogic.org/research/DETECT2020/>

One important consideration in such an automated system is how, and where, to integrate checks *during system operation* that continuously monitor for violations of system safety requirements, e.g., due to unexpected environmental conditions or other scenarios that could not be predicted and tested for during system design. This is especially critical given the automated nature of the systems involved: pilots and human ground controllers make numerous decisions in the control of commercial aircraft that serve as a foundation for their traffic management systems but are missing from UTM. For example, pilots regularly identify and dismiss off-nominal sensor readings and ground controllers operate under unstated assumptions, such as that the flight plans of two aircraft should never contain unsafe overlaps.

Runtime Verification (RV) provides checks that cyber-physical systems adhere to their safety requirements during operation. However, much of the research into RV has focused on increasing expressivity of monitored properties and operational reach of RV engines. The on-board resources, overhead, operational delays, and intrusive system instrumentation required to run these tools are incompatible with flight certification [12]. In response, the Responsive, Realizable, Unobtrusive Unit (R2U2), was designed to monitor sufficiently expressive properties, in real time, under hard resource constraints, with low-to-no overhead, and without system instrumentation that would violate flight certification [17]. Only three RV tools are flight-certifiable: R2U2, Lola [22], and Co-Pilot [16]; R2U2's flexible architecture was the easiest to adapt to our UTM system.

We examine the candidate UTM system [25], overviewing its design, implementation, and initial tests, e.g., with University of Iowa's (U of I's) Operational Performance Laboratory's (OPL's) Vapor 55 UAS flying over small, nearby airspace. We map out three subsystems where RV could be embedded within this UTM framework: on-board the Vapor 55, on-board each Ground Control System (GCS), and within the UTM cloud-based framework. However, the biggest bottleneck to the successful deployment of formal methods like RV is specification of the requirements under verification [19]. Building upon the runtime specification pattern categories of [19], we detail patterns for formal requirements specification across these subsystems and write, debug, and validate a covering set of temporal logic specifications. Using R2U2 to create runtime observers from this specification set, we deploy in simulation real-time RV over a set of real-life flight tests, expanding our data set to include realistic scenarios that were not able to be flown in real life. We examine the outputs from R2U2 and provide a roadmap for utilizing this data to robustify the UTM framework. Our case study details the process of RV integration for future adopters of systems like UTM.

Our contributions are as follows: (1) patterns useful for RV specifications across a real distributed UTM implementation; (2) a method for adding a single first-order operator to Mission-time Linear Temporal Logic (MLTL) specifications in an RV engine; (3) an open set of RV benchmarks from real-world UAS/GCS telemetry data; (4) an extensive experimental evaluation (124 specifications) of a distributed RV implementation in real-time; and (5) lessons learned from distributed RV specifications validation and refinement for a UTM system.

The remainder of this paper is organized as follows. Section 2 gives background information on MLTL and R2U2. Section 3 overviews the candidate UTM framework. Our formal specifications fill Section 4, including specifications specific to the on-board

an *R2U2 specification observation tree*. Redundant branches of the tree can be combined through a pre-flight optimization step for efficiency and reducing encoding size. For example, suppose R2U2 is implemented on a fixed-wing UAS and has two separate specifications: (1) the UAS’s landing gear will be stowed when it is above 1,000 ft, and (2) the UAS’s speed will be within 300mph to 400mph when above 1,000 ft. Since both of these specifications require the altimeter reading to exceed 1,000 ft, a single Boolean operator can be passed to both temporal logic observers.

3 UTM System Definition

In parallel with NASA’s third UTM Technical Capability Level [14], a hybrid university-industrial team proposed an intelligent, centralized UTM for low-altitude urban environments to coordinate UAS traffic in a safe and efficient way [25]. A high-level diagram of the proposed UTM system appears in Fig. 1.

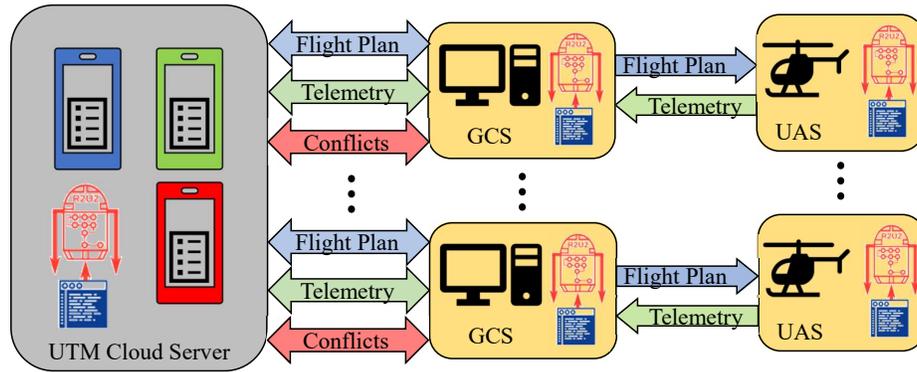


Fig. 1: An overview of the NSF funded cloud-based UTM [25].

Ground Control Stations (GCS) connect to the UTM Cloud Server and upload their proposed flight plan for approval. The UTM Cloud server performs pre-flight plan conflict detection using a dynamic geofencing algorithm [27]. The UTM then notifies the GCS if the flight plan is rejected or approved. If rejected, the GCS should submit a new flight plan until one is approved. When approved, the GCS streams the UAS’s telemetry data to the server, which then performs an en route conflict prediction. If an en route conflict is predicted, the server will alert all GCS involved of the conflict, so that they may have enough time to submit a new flight plan and perform an avoidance maneuver.

There are many challenges to overcome before such a UTM would be incorporated into our NAS [3, 18]. For example, an ongoing research question is how to handle uncooperative and hostile UAS in the UTM’s airspace. One assumption of this UTM is that all UAS are non-hostile, i.e., no UAS is purposefully flying an unapproved flight plan. However, this UTM was designed to receive telemetry data from anyone who connects to it, regardless of flight plan status. While the details of how to maintain communication with both cooperative and uncooperative UAS is still ongoing research [24], RV can be used within this UTM to alert the operator to the presence of uncooperative UAS.

Another ongoing research question for UTMs is whether low-altitude airspace should be structured, e.g., with similar traffic patterns and rules as ground transportation [9].

Regardless of which approach is used, RV can be incorporated to alert users of dangerous or undesirable circumstances. For example, this UTM was developed for unstructured airspace, so it has more general operating range specifications, such as those that make sure that all UAS are within the UTM’s airspace. Conversely, if a structured airspace was chosen, the structured ruleset can be formally verified using RV.

4 UTM Runtime Specifications

We first decide the types of interfaces to each UTM sub-system, then how R2U2 can be implemented into each subsystem, to drive specification elicitation.¹

4.1 UTM Sub-system I/O

UAS The UAS follows a flight plan provided by the GCS and is responsible for collecting and streaming its telemetry data to the GCS. Real flight data from OPL’s Vapor 55 UAV helicopter’s [1] internal log provides the data used for analysis and evaluation. The subset we chose is based on which signals were most useful for performing RV; see Table 1.

Table 1: Selected Output Signals from the UAS.

Signal	Description	Units
Pos{N, E, D}	Relative positional vector (North, East, Downward) from the home point.	{m, m, m}
Lat, Lon, Alt	GPS coordinate positions.	{DD, DD, MSL}
Roll, Pitch, Yaw	Euler angles of the UAS.	{deg, deg, deg}
P, Q, R	Euler angle-rates of the UAS.	{deg/s, deg/s, deg/s}
Vel{N, E, D}	Velocity vector of the UAS.	{m/s, m/s, m/s}
Acc{N, E, D}	Acceleration vector of the UAS.	{m/s ² , m/s ² , m/s ² }
Temp, TempE1/2	Temperature of the air and motors.	C
Pres	Atmospheric pressure.	hPa
Phase	Set of strings corresponding to preset phases of flight.	{<undefined>, Test actuators, Stationary, Hover, Cruise, Go to, Stop at, In flight, Landed}
Subphase	Set of strings corresponding to preset subphases of flight.	{Ready, Test, Takeoff, Manual, Waypoints, Home, Landing}
FlightMode	Set of strings corresponding to automatic and manual control.	{Automatic, Home}
RPM	RPM of the main motor.	–

For each UAS in the system, the number of inputs to an on-board R2U2 implementation remains constant over the entire run and is predetermined prior to runtime. This makes implementations of R2U2 equivalent across all UAS, meaning that the time spent creating specifications for an individual UAS remains constant. This is assuming all UAS in a system are the same class, i.e., all single-rotor helicopter-style UAS with similar parameters.

¹ Note that the list presented is not a comprehensive list of all our specifications; the full list can be found at <http://temporallogic.org/research/DETECT2020/>.

GCS The GCS has many responsibilities within the UTM system. It is responsible for: (1) submitting flight plans to the UTM; (2) directing and receiving telemetry data from an inflight UAS; (3) pre-processing and transmitting any telemetry data received from its UAS to the UTM; and (4) monitoring for any conflict alerts from the UTM. For our case study, we look only at implementing RV to monitor (1), (2), and (4). Due to limitations on the way the UTM’s test data was produced, i.e., the Vapor 55 was only simulated during the UTM test, and because it would be identical to the UAS’s R2U2 implementation, we omitted (3) from the GCS’s R2U2 implementation.

Table 2: Input and Output Signals from the GCS to the UTM

Telemetry Signals			
Signal	GCS I/O	Description	Units
ID	O	The flight plan ID of the telemetry transmission.	int
Time	O	The time stamp when the GCS transmits the telemetry to the UTM.	UNIX
wp{Lon, Lat, Alt}	O	The latitude, longitude, and altitude of the waypoint the UAS is currently flying toward.	DD/MSL
Lon, Lat, Alt	O	The UAS’s measured longitude, latitude, and altitude.	DD/MSL
Vel	O	The UAS’s velocity measurement.	m
Ang	O	The UAS’s heading measurement.	deg.
Flight Plan Signals			
Signal	GCS I/O	Description	Units
fp_ID	I	The UTM’s assigned flight plan ID for the approved flight plan.	int
Status	I	The UTM’s response to the GCS’s flight plan	{ Approved, Rejected, Replaced }
Start	O	The start time of the flight plan.	UNIX
End	O	The estimated end time of the flight plan.	UNIX
Phase	O	The type of waypoint.	{ START, STOP, CRUISE, HOME }
fp{Lon, Lat, Alt}	O	The specific waypoint’s latitude, longitude, and altitude.	DD/MSL
Time_Filed	O	The time stamp when the GCS transmitted the flight plan to the UTM.	UNIX

A challenging aspect of the GCS is that the flight plan data must be continuously streamed to R2U2, since flight plans that are transmitted once across the GCS to the UTM are not saved anywhere in R2U2’s memory (see Table 2). This made formatting R2U2’s inputs from the GCS challenging; in particular, the number of waypoints within a GCS’s flight plan can vary. This led to $NumTelem + NumFP + (NumWPsFP)(NumWP)$ total inputs from a GCS to R2U2, where $NumTelem$ is the number of telemetry inputs, $NumFP$ is the number of inputs from the flight plan, $NumWpsFP$ is the number of signals associated with each waypoint, and $NumWP$ is the number of waypoints within the flight plan. For our specific system, $NumTelem = 9$, $NumFP = 4$, $NumWpsFP = 5$, and $NumWP$ varies between 4 and 10 waypoints.

This variance in the number of inputs from one GCS to another led us to develop specifications that validate across *all* instances of $NumWP$. We accomplished this by adjusting R2U2’s pre-processing layer to iterate across a loop of all instances of one variable (say, $Phase$) and determine if at least one violates a certain property. While this is not a full-fledged first-order logic [4, 8], it leads to a mapping of multiple inputs to a single Boolean atomic input to R2U2 and acts like a single first-order operator to MLTL. At the sacrifice of precision, i.e., rather than knowing which exact waypoint was violating a property, R2U2 reports if at least one input is violating a property, which allows for easier automation and generality when incorporating R2U2 across iterative types of inputs (i.e., varying number of waypoints).

UTM Cloud Server Since the UTM is implemented as a cloud-based, centralized server, it is in charge of consolidating all flight plan and telemetry information and determining whether any two UAS will conflict. Like the instances of R2U2 for the GCS, the number of inputs for the UTM varies: once with the number of waypoints in a flight plan and again with the number of UAS. Thus, the total number of inputs to an instance of R2U2 for the UTM can be calculated by $NumID(NumTelem + NumFP) + (NumWPsFP)(\sum_{i=0}^{NumID} NumWP[i])$, where $NumID$ represents the total number of flight plan IDs in the UTM and $NumWP[i]$ is the specific number of waypoints for flight plan i . This can lead to a large number of inputs for R2U2, i.e., 20 UAS with 4 waypoints each would be 580 inputs.

Similar to the GCS, to get traction on such a large number of inputs, we have designed our specifications similar to first order logic, i.e., *for all UAS* a certain property holds or *there exists a UAS* where a property is violated. Again, we trade expressiveness for performance: we retain real-time performance guarantees but only promise R2U2 will immediately alert the UTM operator of a violation, not identify the specific UAS responsible.

4.2 Coverage of Real-world Specification Types

To help organize our specifications, each one is categorized into one of six labels: (1) operating range, (2) sensor bounds, (3) rates of change, (4) control sequences, (5) physical model relationships, and (6) inter-sensor relationships. These categories resemble those of [19,26], though we add a level of granularity to several for ease of organization.

Operating Range Every sensor to, and variable within, a given system has an expected *operating range* and should it fall below or exceed a given threshold, this may indicate a hazardous system state. For example, the proposed centralized UTM will cover a predefined airspace. Should a UAS stray beyond these operating limits of the UTM, an alert will be sent to the UTM operator to inform the corresponding UAS’s GCS that they are reaching or exceeding a safety threshold of the system.

Sensor Bounds Sensors and variables also have well defined *bounds* on the values they can return. For example, a UAS should never see latitude values that are meaningless (i.e., latitude measurements less than -90° or greater than 90°). These types of specifications may be used in conjunction with *Operating Range* specifications to help diagnose whether there is a user error (accidentally operating outside their airspace) or hardware failure (sensor returning bad data to the system).

Table 3: UAS, GCS, and UTM Specifications Investigated

Name	Description	MLTL Specification
UAS_RC_8	The difference between two consecutive pressure $Pres$ readings cannot exceed a maximum rate of climb $MaxPrevPres$.	$\neg(\Box_{[0,3]}\neg(Pres_leq_MaxPrevPres \wedge Pres_geq_MinPrevPres))$
UAS_IS_1	Since the altimeter and the barometer both derive the air pressure, the error between these two measurements of pressure will be less than the $MaxPresErr$ and greater than $MinPresErr$.	$(Pres_lt_MaxPresErr) \wedge (Pres_gt_MinPresErr)$
GCS_CS_7	The reference latitude Lat_{WP} and longitude Lon_{WP} will be contained within the set of waypoints given in the flight plan.	$wpLonLat_eq_fpLonLat$
GCS_PM_2	If a telemetry stream is reporting that the UAS's heading Ang is between 90° and 180° , then, if the UAS's velocity Vel is greater than 0 m/s, the UAS's latitude Lat should be decreasing while its longitude Lon should be increasing.	$\neg(Ang_eq_Quad4 \wedge Vel_gt_Zero) \vee (Lat_geq_PrevLat \wedge Lon_geq_PrevLon)$
UTM_OR_11	Every UAS's position will be bounded within the given airspace. All latitude Lat will be bounded between $(41.6000^\circ, 41.6720^\circ)$.	$\Box_{[0,3]}(Lat_leq_LatUB \wedge Lat_geq_LatLB)$
UTM_SB_3	Every UAS's position will exist on Earth GPS coordinates. All latitude Lat measurement's will be bounded by $(-90^\circ, 90^\circ)$ degrees.	$\Box_{[0,3]}(Lat_leq_MaxLatUB \wedge Lat_geq_MinLatLB)$

See <http://temporallogic.org/research/DETECT2020/> for a complete set of specifications.

Note that there is an implied \Box operator outside all of the specifications due to the *stream-based* nature of R2U2 runtime observers. That is: R2U2 outputs a stream of verdicts indicating whether each specification holds starting at every discretized execution time stamp. Formally, $\forall i$, R2U2 gives a verdict as to whether $\pi, i \models \varphi$ in the form of a stream $\langle i, verdict \rangle$. So, even the purely propositional formulas are still asserting that a relationship holds, e.g., throughout a flight.

Rates of Change Additionally, sensor's and variable's *rate of change* may also be bounded. For example, a UAS will have some maximum change in velocity between any two telemetry transmissions. Should it exceed this value, it may indicate that the UAS's transmission rate varied (e.g., a dropped transmission). Additionally, one could monitor to make sure there is change between two consecutive sensor measurements, or that the amount of variance between sensor measurements is not skewed in one direction or another, which could mean the UAS is under a cyber-attack, such as GPS spoofing.

Control Sequences Because this system follows a rigorous series of stages, several specifications monitor that the system is adhering to its specified *control sequence*. For example, the intended sequence of states for the UTM is to: (1) receive a flight plan from a GCS, (2) approve or reject the flight plan, (3) if approved, issue the GCS a corresponding flight plan ID, and (4) the GCS transmits the telemetry data of the UAS

with the corresponding flight plan ID. Many different hazardous situations can be made by removing or rearranging this intended sequence; thus, monitoring for any out-of-order sequences can help alert the system or the user to execute a mitigation action.

Physical Model Relationships In many systems, there exist *physical relationships* between one or more combinations of sensors and actuators when commanding the system. For example, if a UAS is commanded to accelerate, the motors should respond accordingly to execute that command. These types of relationships can detect sensor calibration errors and ensure that sensors agree about the system’s overall state.

Inter-sensor Relationships To help diagnose failures, some systems may be able to invoke specifications that use multiple sensors, either of the same or different type, to measure common values. For example, the relationship between barometric pressure (obtained from an on-board barometer) and altitude (obtained from the GPS) allows for more than one way to measure altitude. RV can use these types of specifications to determine if both sensors agree. If they do not agree, then polling, or other system health management techniques, could be used to determine the faulty sensor and switch the primary source for the UAS altitude measurements.

4.3 Specification Validation

Because specification creation is a circular process [19], we chose to validate our list of RV specifications in a variety of ways. The first was a Matlab-based approach where we incorporated logged data for each subsystem into Matlab and validated the ways in which the Boolean atomics were created. The second was by uploading our MLTL runtime specifications for each individual subsystem into an open-source MLTL satisfiability checker [11] to perform specification debugging via checking each specification, its negation, and the conjunction of all specifications for satisfiability [21]. The third way these specifications were validated was by running the pre-recorded data into the R2U2 tool chain and checking to see if the specification held true over the system trace. If it did, we injected faults into the pre-recorded data and monitored R2U2’s output to see if it correctly detected the faults. Of the list of 124 specifications we made for the UAS, GCS, and UTM, Table 3 presents six specifications that we feel encapsulate interesting properties about each subsystem.

5 Evaluation

The UTM test scenario consists of 20 UAS interacting with the UTM – OPL’s Vapor 55 hardware-in-the-loop simulation and our 19 physics-based simulated flights – with the goal of testing the UTM’s conflict detection logic. Of the 20 flight plans, 18 were conflict-free, one was designed to create a pre-departure conflict, and one deviated from the pre-approved flight plan, creating an en-route conflict. During the 42 minute test, the UTM correctly detected and alerted both GCSs of the en-route conflict, with OPL’s GCS submitting a new, conflict-free flight plan en-route.

Although we intended to have R2U2 embedded into the UTM system for this test, in practice this would have required enhancements to core functionalities and improving the networking capabilities of the UTM. However, all test data was recorded and put to

use offline in refining our specifications and implementations of R2U2 into each subsystem. We argue that since R2U2 has previously been embedded and used in several successful aerospace applications [7, 10, 13, 17, 23], our offline, real-time simulations of this embedding perform representatively to an actual implementation. Note we plan to incorporate R2U2 into the UTM system for the next test.

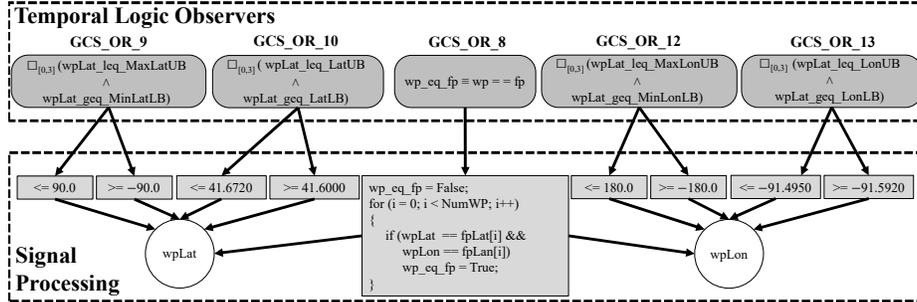


Fig. 2: A small observation tree from the GCS’s R2U2 implementation. Two sensor values, $wpLat$ and $wpLon$, are inputs to the signal processing layer, which pre-processes them into Boolean atomics for the temporal logic observers.

R2U2 was hosted on a Ubuntu 18.04 LTS operating system on an Intel Core i7-4810MQ CPU with a 2.80GHz clock and 16GB of RAM. Each subsystem of R2U2 was run independently, i.e., each subsystem was run with its own instance of R2U2 across its own input trace and no cross-platform communication was performed. Fig. 2 shows an example of how specifications are encoded into R2U2’s observations trees.

Operating Range As seen from Fig. 3, the UTM’s R2U2 monitors and reports if the operating range bounds are satisfied for all of UAS’s latitude measurements. As the original test data was fault-free, we injected a fault, which revealed a sudden spike in R2U2’s output during the injected fault. This corresponds to a dropped transmission in the original data. Thus, we refine our specification to include an overarching $\square_{[0,3]}$ operator, which acts as a sliding window temporal filter, to suppresses such output bouncing.

Sensor Bounds Similar to Fig. 3, Fig 4 shows the UTM’s R2U2 monitoring and reporting if any of the UAS’s latitude measurements exceed the sensor bound threshold of $(-90^\circ, 90^\circ)$. Similarly, the original data was fault free, so we injected a fault into one of the UAS’s latitude measurements. Again, testing revealed transmission losses, so we added a $\square_{[0,3]}$ filter to suppress any false positives triggered by missing data.

Rates of Change The pressure recorded by a UAS’s on-board barometer changes as it ascends and descends. Thus, we developed a specification to monitor change in pressure: the difference between two consecutive pressure readings are limited to ± 0.4 hPa (derived from the maximum rates of climb and descent [15]). Unlike our other specifications, Fig. 5 shows that we needed to include a conjunction of two $\square_{[0,3]}$ filters to remove all output bouncing: one filters outlying violating verdicts and one filters outlying satisfying verdicts.

Control Sequence The UTM’s test scenario included one UAS deviating from its pre-approved flight plan. Figure 7 shows R2U2 correctly detecting this real-world deviation in real time.

Inter-sensor Relationship The difference between the barometer’s and GPS’s pressure should be bounded within acceptable error. A comparison of the two sensors can help diagnose sensor failures (see [15] for more details). For example, Fig. 6 shows a side-by-side comparison of two pressure traces: an unmodified and a modified version with a fault injected from $t = 1500$ to $t = 1750$.

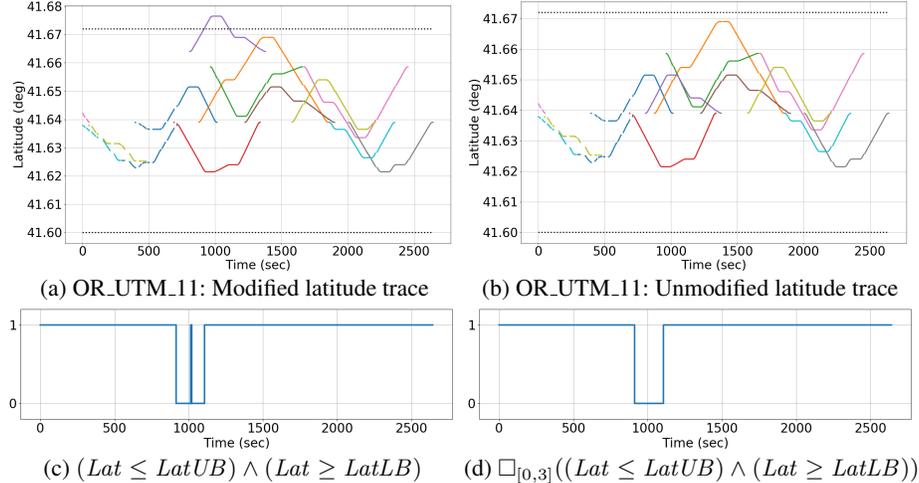


Fig. 3: Two instances of the UTM’s R2U2 monitoring: a modified run (a) where one UAS (purple) temporarily exceeds the operating range bounds, and an unmodified run (b) where all UAS lie within the operating range (dashed lines). Both fault-injected runs show R2U2 identifies the corresponding violation of the specification; however, the output of the purely Boolean formula (c) bounces due to a missed telemetry transmission. To avoid a false positive, due to missing data, we add a temporal logic filter (d) that monitors for multiple subsequent nominal data sequences.

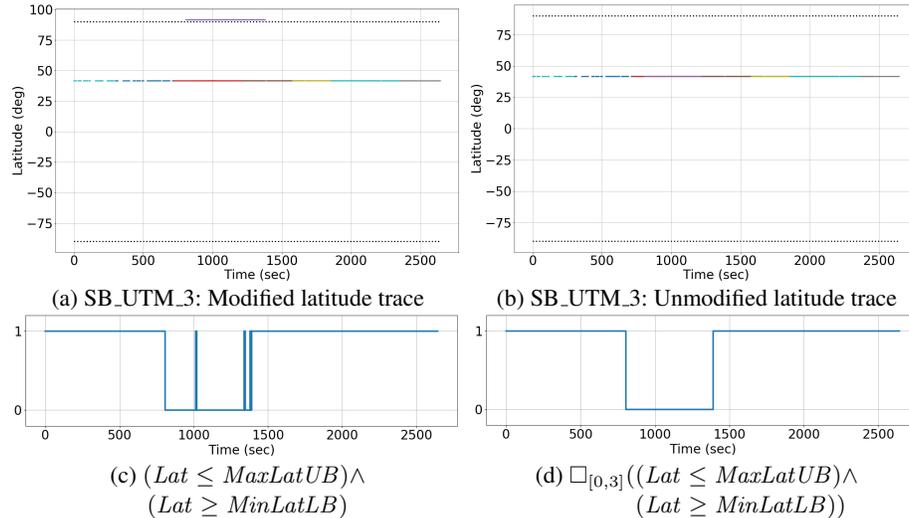


Fig. 4: Like Fig. 3, the top graphs show modified (a) and unmodified (b) input traces. Similarly, dropped telemetry transmissions cause output bouncing (c), so a $\square_{[0,3]}$ filter is applied (d).

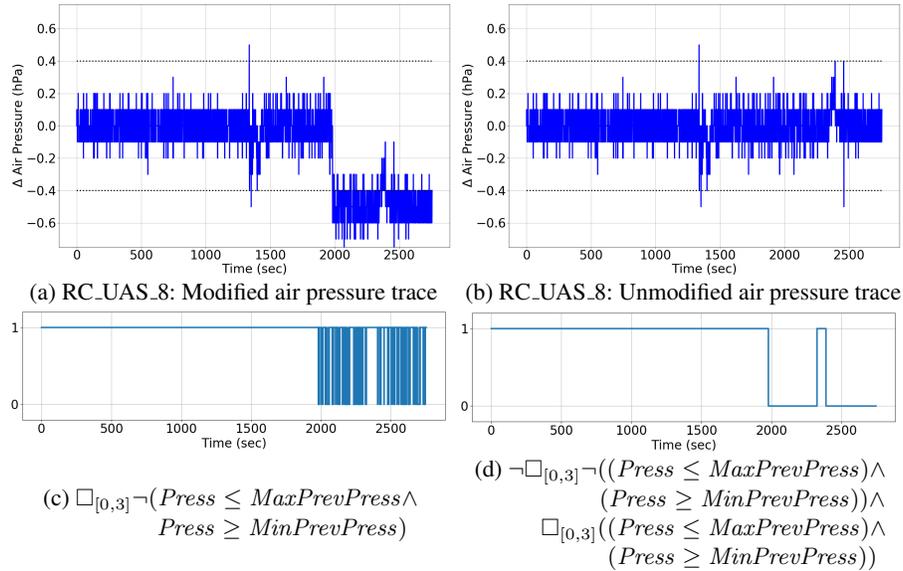


Fig. 5: Two instances of the UAS’s R2U2 monitoring: (a) a modified trace where we injected a shift in the air pressure’s rate of change, and (b) an unmodified trace where a few anomalies exceed the pressure rate of change bounds (dashed lines). Both outputs of the fault-injected run from R2U2 are shown; however, the output of the original formula (c) bounces due to noisy input jumping back within the margins. To remove this bouncing, we added another $\square_{[0,3]}$ filter (d) to keep the current state until all outliers are filtered and the state has unquestionably changed.

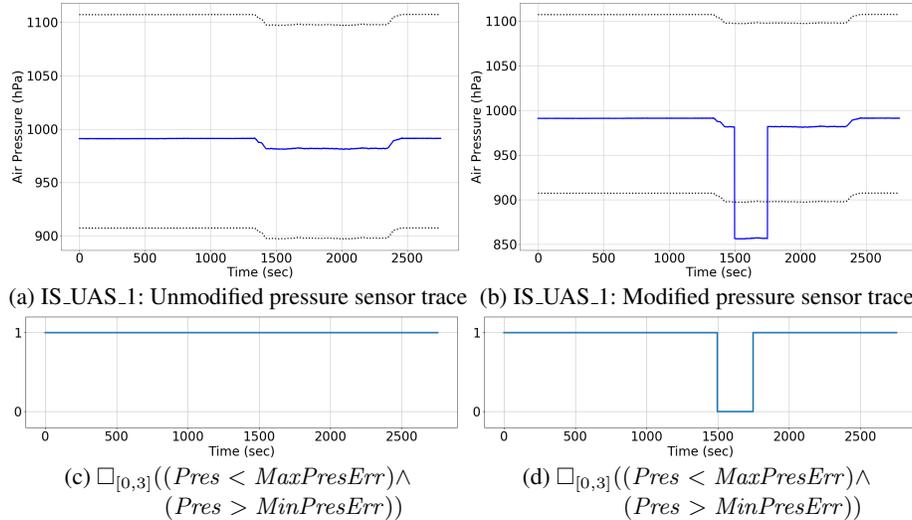


Fig. 6: Two instances of the UAS’s R2U2 monitoring: an unmodified run (a) where the pressure from the barometer remains within the error margins of the GPS’s calculated atmospheric pressure (dashed lines), and a modified run (b) where the same data was injected with a fault by subtracting 100hPa from the barometer’s atmospheric pressure reading. R2U2’s output (c) acknowledges the error-free trace of (a), and (d) shows that R2U2 detects the violation from (b).

Physical Model Relationship As shown in Fig. 8, when a UAS’s heading is between 90° and 180° and its velocity is non-zero, then the UAS’s latitude should be decreasing while its longitude is increasing.

Lessons Learned Many of our specifications are rather simplistic, i.e., $\square_{[0,3]}(\varphi_1 \wedge \varphi_2)$; however, their simplicity allows for easy validation and verification. They are easy to validate through discussion with system designers. Additionally, we used temporal filters, e.g., the $\square_{[0,3]}$ sliding window filter, extensively to mitigate false-positives. As false-positives can cause mistrust of the RV monitor, we built our specifications to err on the side of missing a fault. As seen in Section 5, if R2U2 sent a fault alert, the fault was clear for the human operators receiving the alert. Many of our specifications encapsulate intuitive bounds and re-

relationships for sensor values and variables that humans implicitly assume about a given system, i.e., latitude coordinates are bounded between $(-90^\circ, 90^\circ)$ and that events cannot end before they start. These “common-sense” specifications are often overlooked, yet they catch real faults, e.g., from variable overflow and underflow, sensor or wiring failures, and excessive noise. Our coverage categorization for specifications allowed us to enumerate many such sanity checks about the UTM system, which helped us achieve a reasonable covering set of specifications for the UTM’s three sub-systems. In practice, this led to R2U2 identifying a real-life fault where a data-translation error caused the UTM to register flight plans that ended before they started. Such an error would be obvious to human controllers but automated systems require RV to flag this impossibility. Future work is aimed toward creating automated tools for specification elicitation.

6 Conclusion

Before UAS can integrate into the NAS, we need to establish a provably safe, intelligent, and automated UTM system. To help facilitate this, we have integrated the state-of-the-art runtime verification tool R2U2 across the three different layers of an actual UTM implementation: on-board the individual UAS, in conjunction with each operator’s GCS, and embedded into a centralized, cloud-based UTM server. By validating and releasing over 100 runtime MLTL specifications, two sets of recorded traces from test flights of a real-life UTM implementation, and the results of checking those formulas, we contribute a large benchmark suite. This suite is useful for verification of

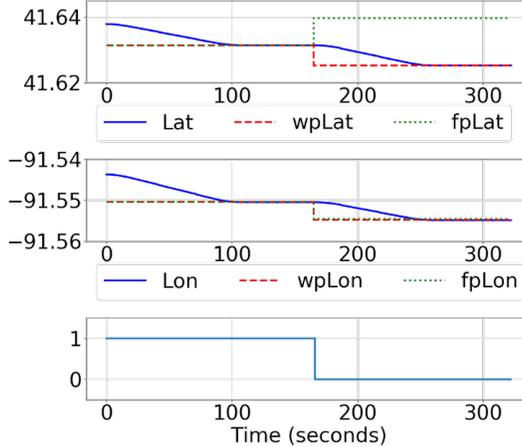


Fig. 7: The latitude (top) and longitude (middle) traces for an adversarial UAS, showing that the GCS is commanding it to a different waypoint (red, dashed line) instead of one from its approved flight plan (green, dotted line). Corresponding to the violation of CS_GCS.7 (Table 3), R2U2’s output (bottom) shows it successfully detects this real-world fault.

the algorithms and implementations of future RV tools, providing both nominal and faulty traces and realistic sensor noise and outlier readings that challenge RV engines.

Additionally, we exemplify the real-world challenges of implementing RV into a centralized, high-traffic UTM. We demonstrate real-time performance of extending MLTL formulas with a single first-order operator, where we validate whether a specification holds *for all* UAS or if *there exists* a UAS that violates a specification. When refining our specification set, we found sensor noise and outliers triggered false positives and that a simple $\square_{[0,3]}$ around each critical sensor check eliminated these while only slightly delaying the trigger of actual faults. Of our 124 specifications, two-thirds contain this construct. This modification can be automatically inserted into specifications for real-life systems where false positives cannot be tolerated. Though we verified a short (42 minute) relatively small real-life system (26, 33-64, and 634 sensor inputs for the UAS, GCS, and UTM, respectively) we still found it hard to manually write a sufficiently covering set of specifications. To ensure we did not miss covering unstated assumptions, we used coverage metrics to brainstorm our list of 124 specifications: variable coverage (every variable appears in at least one specification) and pattern coverage (specifications follow each pattern from [19]). Our experience informs an on-going project to enable more automated specification elicitation.

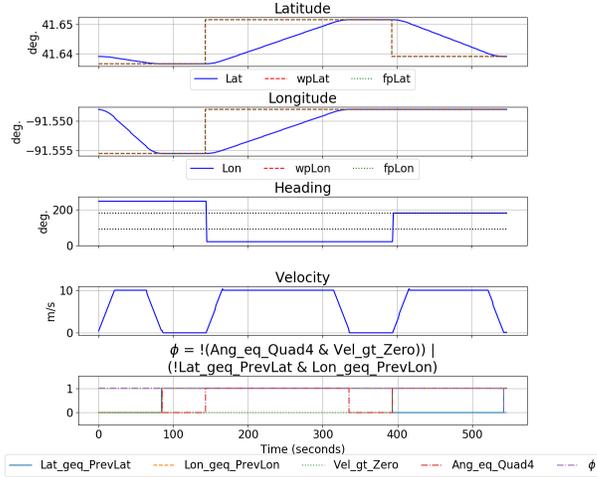


Fig. 8: Single instance of R2U2 on a simulated UAS showing the latitude, longitude, heading, and velocity. With assumptions of the UAS operating in North America and there is a relationship between heading and trajectory, then a relationship between velocity, heading, and position can be verified.

References

1. AeroViroment: VAPOR All-electric Helicopter UAS. <https://www.avinc.com/uas/view/vapor-vtol>, last accessed on Dec. 17, 2019
2. Alur, R., Henzinger, T.A.: Real-time logics: Complexity and expressiveness. *Information and Computation* **104**(1), 35–77 (1993)
3. Aweiss, A.S., Owens, B.D., Rios, J.L., Homola, J.R., Mohlenbrink, C.P.: UAS Traffic Management National Campaign II. In: 2018 AIAA SciTech. pp. 1–16 (Jan 2018)
4. Bakhirkin, A., Ferrère, T., Henzinger, T., Nickovic, D.: The First-Order Logic of Signals. In: EMSOFT (2018)
5. Federal Aviation Administration (FAA): FAA Aerospace Forecast – Fiscal Years 2019–2039. Online: https://www.faa.gov/data_research/aviation/aerospace_forecasts/media/FY2019-39_FAA_Aerospace_Forecast.pdf (2019)

6. Federal Aviation Administration (FAA): Unmanned Aerial Systems (UAS). Online: <https://www.faa.gov/uas/> (2020)
7. Geist, J., Rozier, K.Y., Schumann, J.: Runtime Observer Pairs and Bayesian Network Reasoners On-board FPGAs: Flight-Certifiable System Health Management for Embedded Systems. In: RV14. vol. 8734, pp. 215–230. Springer-Verlag (September 2014)
8. Havelund, K., Peled, D., Ulus, D.: First order temporal logic monitoring with bdds. In: FM-CAD. pp. 116–123 (2017)
9. Hunter, G., Wei, P.: Service-oriented separation assurance for small uas traffic management. In: INCS19. pp. 1–11 (2019)
10. Kempa, B., Zhang, P., Jones, P.H., Zambreno, J., Rozier, K.Y.: Embedding Online Runtime Verification for Fault Disambiguation on Robonaut2. In: FORMATS. LNCS, Springer (2020)
11. Li, J., Vardi, M.Y., Rozier, K.Y.: Satisfiability Checking for Mission-Time LTL. In: CAV. LNCS, vol. 11562, pp. 3–22. Springer, New York, NY, USA (July 2019)
12. The international conference on runtime verification. Online: <https://www.runtime-verification.org/> (2001–present)
13. Moosbrugger, P., Rozier, K.Y., Schumann, J.: R2U2: Monitoring and Diagnosis of Security Threats for Unmanned Aerial Systems. FMSD pp. 1–31 (April 2017)
14. NASA: Unmanned Aircraft System (UAS) Traffic Management (UTM). <https://utm.arc.nasa.gov/index.shtml>, last accessed on Mar. 12, 2020
15. NASA: Earth atmosphere model. Online: <https://www.grc.nasa.gov/WWW/K-12/airplane/atmosmet.html> (May 2015)
16. Pike, L., Wegmann, N., Niller, S., Goodloe, A.: Copilot: monitoring embedded systems. *Innovations in Systems and Software Engineering* **9**(4), 235–255 (2013)
17. Reinbacher, T., Rozier, K.Y., Schumann, J.: Temporal-Logic Based Runtime Observer Pairs for System Health Management of Real-Time Systems. In: TACAS. pp. 357–372 (2014)
18. Rios, J., Mulfinger, D., Homola, J., Venkatesan, P.: NASA UAS traffic management national campaign: Operations across Six UAS Test Sites. In: DASC. pp. 1–6 (2016)
19. Rozier, K.Y.: Specification: The Biggest Bottleneck in Formal Methods and Autonomy, vol. 9971, pp. 8–26. Springer (Nov 2016)
20. Rozier, K.Y., Schumann, J.: R2U2: Tool Overview. In: RV-CUBES. vol. 3, pp. 138–156. Kalpa Publications, Seattle, WA, USA (September 2017)
21. Rozier, K.Y., Vardi, M.Y.: LTL satisfiability checking. *International Journal on Software Tools for Technology Transfer (STTT)* **12**(2), 123 – 137 (March 2010)
22. Schirmer, S.: Runtime Monitoring with LOLA. Master’s thesis, Saarland University (November 2016), <https://elib.dlr.de/113126/>
23. Schumann, J., Rozier, K.Y., Reinbacher, T., Mengshoel, O.J., Mbaya, T., Ippolito, C.: Towards Real-time, On-board, Hardware-supported Sensor and Software Health Management for Unmanned Aerial Systems. *IJPHM* **6**(1), 1–27 (June 2015)
24. Wargo, C.A., Glaneuski, J., Hunter, G., DiFelici, J., Blumer, T., Hasson, D., Carros, P., Kerczewski, R.J.: Ubiquitous surveillance notional architecture for system-wide daa capabilities in the nas. In: 2018 IEEE Aerospace Conference. pp. 1–14 (2018)
25. Wei, P., Atkins, E.M., Hunter, G., Rozier, K.Y., Schnell, T.: Pre-Departure Dynamic Geofencing, En-Route Traffic Alerting, Emergency Landing and Contingency Management for Intelligent Low-Altitude Airspace UAS Traffic Management. Online: https://www.nsf.gov/awardsearch/showAward?AWD_ID=1718420 (July 2017)
26. Zhao, Y., Rozier, K.Y.: Formal specification and verification of a coordination protocol for an automated air traffic control system. *Science of Computer Programming* **96** (May 2014)
27. Zhu, G., Wei, P.: Low-Altitude UAS Traffic Coordination with Dynamic Geofencing. 16th AIAA Aviation Technology, Integration, and Operations Conference (June 2016)