

SA4U: Practical Static Analysis for Unit Type Error Detection

Max Taylor
taylor.2751@osu.edu
The Ohio State University
United States

Xiaorui Wang
wang.3596@osu.edu
The Ohio State University
United States

Johnathon Aurand
aurand.15@osu.edu
The Ohio State University
United States

Brandon Henry
brandon.henry@tangramflex.com
Tangram Flex
United States

Feng Qin
qin.34@osu.edu
The Ohio State University
United States

Xiangyu Zhang
xyzhang@cs.purdue.edu
Purdue University
United States

ABSTRACT

Unit type errors, where values with physical unit types (e.g., meters, hours) are used incorrectly in a computation, are common in today's unmanned aerial system (UAS) firmware. Recent studies show that unit type errors represent over 10% of bugs in UAS firmware. Moreover, the consequences of unit type errors are severe. Over 30% of unit type errors cause UAS crashes. This paper proposes SA4U: a practical system for detecting unit type errors in real-world UAS firmware. SA4U requires no modifications to firmware or developer annotations. It *deduces* the unit types of program variables by analyzing simulation traces and protocol definitions. SA4U uses the deduced unit types to identify when unit type errors occur. SA4U is effective: it identified 14 previously undetected bugs in two popular open-source firmware (ArduPilot & PX4.)

CCS CONCEPTS

• **Software and its Engineering** → **Abstract Data Types**; • **Software Defect Analysis**; • **Mathematics of computing** → **Mathematical Analysis**;

KEYWORDS

abstract data type inference; physical units; physical unit mining

ACM Reference Format:

Max Taylor, Johnathon Aurand, Feng Qin, Xiaorui Wang, Brandon Henry, and Xiangyu Zhang. 2022. SA4U: Practical Static Analysis for Unit Type Error Detection. In *37th IEEE/ACM International Conference on Automated Software Engineering (ASE '22)*, October 10–14, 2022, Rochester, MI, USA. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3551349.3556937>

1 INTRODUCTION

Unit type errors (UTEs) occur when developers mistakenly use incorrect physical units in a computation. For example, developers may accidentally store a value that represents a physical quantity measured in centimeters into a variable that is meant to store a value measured in meters. This occurs because the types of program variables (e.g., `int`, `double`) do not convey the physical units they are expected to store.

Unfortunately, UTEs are common in unmanned aerial systems (UAS) firmware. Recent work has shown that UTEs account for over 10% of bugs in UAS firmware [26]. Moreover, the consequences of UTEs are severe: Over 30% cause UAS crashes. There are several infamous disasters caused by UTEs: a \$125 million spacecraft was lost when developers failed to convert between the metric system and imperial units, and a satellite was unavailable for one week, among others [24] [29].

A unit type has two components: a *dimension* (e.g., distance and time) and a *frame of reference*. Dimensions are measured in physical units such as meters and hours. Frames of reference describe how a measurement was obtained. For example, a distance is relative to a starting location, which could be the front of the vehicle or somewhere else.

Today, developers have two choices in automated methods to handle UTEs in their source code. First, they may use a *unit library* such as C++'s BoostUnits [25] or Java's CaliperSharp [2]. A unit library wraps primitive types (e.g., `int`) in a semantic wrapper (e.g., `Meter`). Similarly, developers may choose to implement checks using the pattern demonstrated in [8]. Second, developers may use a dimensional analysis tool (e.g., Phys [18], PhysFrame [17], and [22]) to detect conversion errors.

There is a major drawback associated with unit libraries: Developers must annotate all program variables with their unit types. This puts a burden on developers and offers little help for existing firmware. Moreover, unit libraries either introduce runtime overhead (e.g., CaliperSharp) or only check limited expressions (e.g., BoostUnits). Today, unit libraries are unused in popular open-source UAS firmware.

While addressing the aforementioned problems, existing tools using dimensional analysis have their own limitations. Phriky-units [23] helps reduce the burden on developers whose projects use Robotic Operating System (ROS) [5]. Phriky-units annotates ROS's APIs, uses type inference to propagate the unit types of program variables, and then checks that the unit types are used correctly. But the annotation burden for developers who do not use ROS is huge. Phys [18] reduces the annotation burden by extracting unit type information from variable names. However, large code bases often use inconsistent units internally. For example, a variable called `altitude` could store a value measured in centimeters or meters. PhysFrame [17] extends Phriky-units and Phys to also check the reference system of measurements. Once again, developers who do not use ROS assume an annotation burden. Moreover, prior work only considers the *dimension* of a measurement (e.g., distance and

ASE '22, October 10–14, 2022, Rochester, MI, USA

© 2022 Association for Computing Machinery.

This is the author's version of the work. It is posted here for your personal use. Not for redistribution. The definitive Version of Record was published in *37th IEEE/ACM International Conference on Automated Software Engineering (ASE '22)*, October 10–14, 2022, Rochester, MI, USA, <https://doi.org/10.1145/3551349.3556937>.

time). These approaches fail to recognize when imperial and metric units are mixed.

These facts motivate us to develop SA4U (pronounced “safe for you”): Static Analysis for UAS. SA4U analyzes C and C++ programs to detect UTEs. The status quo of firmware development and existing solutions motivates us to design SA4U to:

- (1) **Work without developer annotations:** research suggests annotation burden remains a dominant factor in the lack of adoption of existing work, despite the state-of-the-art’s effort to reduce annotation burden [20].
- (2) **Check units:** existing tools only analyze the dimension of a measurement, and not its specific unit.
- (3) **Check reference frames:** in our experience, reference frames are often harder for developers to reason about than unit types.

To work without developer annotations, SA4U obtains type information from two sources: (1) protocol files, and (2) program traces. Protocol files define messages exchanged between the pilot’s control computer and the UAS. Common protocols (MAVLink [3] and LMCP [13]) distribute files that define the unit types of fields in messages. These files are used by existing firmware to generate data structures. We find that they are good sources for type checking too.

It is not always possible to propagate unit types from protocol files to all program variables because some expressions cannot be typed precisely (e.g., `vector` and `set`). To handle this challenge, SA4U uses program traces to deduce the types of program variables. SA4U instruments firmware to sample writes to program variables. Inspired by the success of simulations in related work (e.g., [28] and [27]), SA4U executes the firmware in a simulation to obtain a trace file. SA4U’s type deduction engine analyzes the trace file to identify the unit types of program variables.

The types obtained from the type deduction engine and protocol files are provided to SA4U. SA4U parses the firmware’s source code and applies a set of inference rules to generate a set of constraints. SA4U feeds these constraints to the Z3 theorem prover [12], then reports a UTE if the constraints are unsatisfiable.

SA4U partially interprets the source code as it parses to handle scenarios where the value of a variable controls the unit type of another variable, as is common with communication protocols such as MAVLink. SA4U over-approximates possible runtime values of these so-called *control variables* to precisely diagnose UTEs. SA4U keeps its run-time manageable by only approximating the runtime values of control variables.

To summarize, our contributions are:

- **A type deduction engine** that mines data from firmware simulations to deduce both unit types and reference frames.
- **SA4U**, a practical prototype for detecting UTEs in real-world source code. We also created a prototype editor plugin to provide developers with online feedback.
- **Experimental results** on two popular open-source firmware: ArduPilot & PX4. Our results show it is possible to accurately perform type deduction on program variables in UAS. They also show SA4U is effective, identifying 14 previously undetected UTEs.

```

1 float closest_z(Location l, Location ol,
2   Velocity v, Velocity ov, u32 time) {
3   // Velocity is measured in m/s, so
4   // delta_vel_d stores m/s.
5   float delta_vel_d = ov.z - v.z;
6
7   // pos is measured in cm, so delta_pos_d
8   // stores cm.
9   float delta_pos_d = ol.z - l.z;
10
11  // ERROR: cm - m.
12  return fabsf(delta_pos_d - delta_vel_d *
13    time) / 100.0f;
14 }

```

Figure 1: APM-20286: A previously reported UTE.

2 SA4U IN A NUTSHELL

This section presents three key challenges that existing approaches face when detecting UTEs. Then we show how SA4U’s new unit type representation combined with its mining-assisted static analysis approach addresses these challenges.

2.1 Three Key Challenges

Challenge 1: Imprecise Unit Types. Figure 1 shows a procedure from ArduPilot that calculates the expected distance along the z axis between the UAS and an obstacle after `time` seconds. The procedure computes how fast the vehicle is approaching the obstacle at line 3. Then, it calculates the distance between the vehicle and the obstacle at line 6. Finally, the procedure returns the final distance along the z axis at line 9. However, line 9 contains a subtle UTE. Velocity is measured in meters per second, but position is measured in centimeters. Thus, line 9 mistakenly subtracts meters from centimeters. This error causes the firmware to fail to recognize that the UAS will pass too close to an obstacle.

Prior work fails to diagnose the UTE in Figure 1. Phys represents physical unit types as b_i^j , where b_i is a base unit in one dimension defined by the International System of Units (SI) [10] and $j \in \mathbb{Z}$. So, the type of `delta_pos_d` is represented as `meters^1`. The type of `delta_vel_d` is represented as `meters^1 * seconds^-1`. The computation at line 9 is therefore considered legal, since subtracting meters from meters is allowed. Phys cannot detect UTEs that involve different measurement units in the same dimension due to this design decision.

This type of error is likely to occur in practice. MAVLink [3], measures distance in meters in over 450 message fields. However, it also measures distance in centimeters in 35 other message fields. As another example, time is measured in microseconds in over 100 message fields, but is also measured in milliseconds in 60 others. It is easy to forget to convert between units within the same dimension. Tooling is needed to catch this type of error.

Challenge 2: Missing Frames of Reference. Figure 2 shows a simplified version of ArduPilot’s message handler for obstacle distances. A separate system onboard the UAS detects obstacles and communicates their distances to the firmware with this message. Execution of the firmware’s message handler starts at line 1. The

```

1 void handle_obstacle_distance_msg(const
  mavlink_obstacle_distance_t &msg) {
2   ...
3 + if (msg.frame != MAV_FRAME_BODY_FRD) {
4 +   log("Unsupported frame");
5 +   return;
6 + }
7   set_obstacle_boundary(msg.angle, msg.
    distance / 100.0);
8   ...
9 }
10
11 void set_obstacle_boundary(float angle,
    float distance) {
12   ...
13   _angle = angle;
14   _distance = distance;
15   ...
16 }

```

Figure 2: APM-16903: A previously undetected UTE discovered by SA4U.

```

1 vector<mav_mission_item> waypoints;
2
3 void handle_waypoint(mav_mission_item &m) {
4   // Logic to validate m not shown.
5   ...
6   waypoints.append(m);
7 }

```

Figure 3: ArduPilot’s procedure to handle waypoints.

handler performs some processing of the message and then invokes the `set_obstacle_boundary` function at line 7 with the message’s `angle` and `distance` fields. The body of `set_obstacle_boundary` assigns these fields to the `_angle` and `_distance` fields at lines 13 and 14. This is a UTE, since the program variable `_angle` is measured in the vehicle’s reference frame (i.e., `MAV_FRAME_BODY_FRD`). Meanwhile, `msg`’s members could be measured in *any* reference frame, e.g., relative to the front of the obstacle sensor instead of the vehicle’s reference frame. This error can cause the UAS to fail to recognize that its current trajectory leads to a collision, since `_angle`’s frame of reference is managed incorrectly. After consulting with ArduPilot’s developers, we created the patch shown at line 3. ArduPilot does not mean to support coordinate systems other than the vehicle’s body frame for this message, so the patch simply discards messages with unexpected frames.

Both Phys and Phriky do not consider a measurement’s frame of reference, so they are unable to identify the UTE in Figure 2. PhysFrame was designed to catch errors involving frames of reference. But it targets systems that use ROS, and cannot be trivially modified for use on this system.

Challenge 3: Inaccurate Unit Typing via Static Inference. One possible way to infer unit types for program variables is based on pure static inference, as Phriky does. Specifically, one can start

from the unit types that are precisely defined in the communication protocols between the firmware and the UAS or the controllers (e.g., MAVLink in ArduPilot and PX4), then propagate the unit types using dataflow analysis, and finally detect inconsistency of unit types using dimensional analysis. However, the presence of complex variable types (e.g., `vector`, `set`) in C++ prevents the propagation of unit types to many variables in the source code of firmware via static analysis.

Figure 3 shows such an example, a simplified code snippet from ArduPilot. This procedure handles waypoints a user uploads for the UAS to navigate to. The unit type of the fields in `mav_mission_item` can vary, depending on the message’s frames of reference. As a result, the container variable `waypoints` could have many elements with different unit types. When assigning an element out of this container to a variable `nextwaypoint`, a static inference tool must conservatively assign all possible unit types to the variable, which easily leads to inference explosion.

Due to similar reasons, prior work has shown that only a small fraction of variables can be assigned with unique unit types in ROS projects [18]. To address this challenge, Phys leverages the hints from variable names as a source of unit types. However, variable names in UAS firmware often provide misleading information about the unit type of a variable. For example, Phys uses the rule that variables whose name ends with `position` stores values in meters. But this is not always true in UAS firmware.

2.2 How does SA4U Help?

SA4U introduces three critical ideas. First, SA4U improves the precision of the representation of unit types. Specifically, SA4U differentiates between types *in the same dimension*. So, meters and centimeters are different types. Second, SA4U enhances the unit type representation with the frames of reference so that it can detect the inconsistency between variables with different frames of reference. Third, SA4U infers likely types of program variables through profiled values at run time, in addition to the unit types defined in the message fields in the communication protocol files.

SA4U’s Unit Type Representation. A *unit type* in SA4U is defined as a tuple (*Unit*, *Frame*). In contrast to prior work, SA4U encodes precise measurement unit information in its unit representation. SA4U represents the unit of a measurement as $s * b_i^j$, where b_i is an SI unit [10] (meter, second, mole, ampere, kelvin, candela, or gram) and $s \in \mathbb{R}$. s is the log10 of the measurement unit’s scalar multiple, so that type checking is decidable. All physical units (e.g., volts) can be expressed as a combination of SI base units. The frame is a constant (e.g., `MAV_FRAME_BODY_FRD`) or `ANY` that identifies the frame of reference for the measurement.

How does SA4U address challenge 1? Consider the bug shown in Figure 1. SA4U represents the unit of `delta_vel_d` as $0 * \text{meters} * \text{seconds}^{-1}$. The type of `time` is `seconds`. However, the type of `delta_pos_d` is represented as $-2 * \text{meters}$. SA4U reports an error at line 9 since `delta_vel_d * time` and `delta_pos_d` are subtracted, but their types are different.

This unit type representation also allows SA4U to elegantly distinguish imperial and metric units. It enables SA4U to diagnose UTEs like the one mentioned in §1. For example, SA4U represents yards as $\log_{10}(0.91) * \text{meters}$. In contrast, prior work proposes

handling this case by introducing a separate unit type for each imperial unit. But this direction is not practical: False positives would be reported even if developers correctly converted between unit systems since the representation of different units is orthogonal.

How does SA4U address challenge 2? SA4U constrains the frames of variables based on the definition it extracts from protocol files. Consider the example shown in Figure 2 without the patch. SA4U learns `msg.frame = frame(msg.angle)` from MAVLink’s protocol file. SA4U initially approximates the value of `msg.frame` to be `Any`. Since `_angle` is assigned with `msg.angle`, and `frame(_angle) != Any`, SA4U reports the mismatch error.

Now, consider the behavior of SA4U with the patch. SA4U witnesses the conditional statement at line 3, so SA4U refines its approximation of the value `msg.frame` as follows. After observing the return statement at line 5, SA4U applies the complement of the refined estimate in the remaining function body. Thus, SA4U recognizes that `frame(mavlink_obstacle_distance_t.angle) = MAV_FRAME_BODY_FRD`. SA4U does not report a false positive in the patched version.

How does SA4U address challenge 3? SA4U mines likely unit types from program traces obtained from simulated UAS execution. Specifically, SA4U inserts instrumentation to sample the runtime values of program variables. SA4U also samples the physical states (e.g., position, acceleration, velocity) of objects in the simulation. SA4U knows the unit types of simulated objects a priori (with minimal annotations from us). SA4U compares the values of program variables with the simulation’s physical states to deduce the variable’s likely unit type. In total, SA4U samples 15 values per simulated object. These values correspond to entries in the MAVLink `HIL_STATE_QUATERNION` message (i.e., the minimum information required to perform a simulation). This instrumentation is only required once per *simulator*, and allows all firmware that support the simulator to benefit from SA4U.

For example, ArduPilot contains a variable called `target_altitude`. This variable stores the altitude that the user commanded the UAS to navigate to. SA4U compares `target_altitude` with the simulation’s physical state, including the UAS’ current altitude stored in the field `altitude`. SA4U deduces `target_altitude` and `altitude` share the same unit type, since the UAS’ altitude is always eventually approximate to the value of `target_altitude`.

3 SA4U DESIGN

Figure 4 shows a high-level overview of SA4U. Developers run SA4U’s instrumentation program on firmware source code to insert instrumentation that tracks the runtime values of program variables. Then, developers run the instrumented binary in a simulator. In this step, the instrumentation creates a trace file that contains the runtime values of program variables. The instrumentation also writes quantities of interest (i.e. the values of physical quantities with known unit types) to the trace file. Next, the type deduction program mines the trace file to produce the type database. The type database tracks the types of some program variables. Finally, the static analysis stage employs unit type information from the type database and protocol file, and conducts type inference on the firmware source code to check types are used correctly.

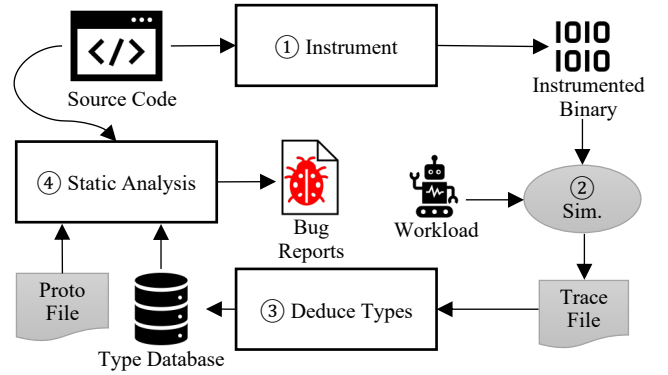


Figure 4: SA4U’s workflow. The numbered boxes show each stage of SA4U’s execution. Components shown in gray are not part of SA4U. Icons show each stage’s inputs and outputs.

3.1 Instrumentation

SA4U’s instrumentation tool inserts code to record the values of the firmware’s program variables. One potential problem with instrumentation is the disturbance of the timing during the firmware execution. As a result, the firmware detects starvation and triggers fail-safes that prevent the workload from executing normally. To prevent this scenario, we exploit the following two observations:

- (1) **We do not need all program variables.** The types of local variables can often be inferred from the types of non-local variables used within the same function based on their interactions. Therefore, we only instrument the accesses (loads and stores) of variables with non-local lifetimes.
- (2) **We do not need all the values of a program variable.** The instrumentation limits the number of values recorded to a fixed rate (in our case, 1 per second). Since the same values of a variable could be stored up to tens of thousands of times each second, this greatly reduces the amount of data that must be recorded and analyzed.

The instrumentation tool builds a partial function $ID : \text{Names} \rightarrow \mathbb{N}$ that maps variable names to integer IDs. When the instrumentation tool encounters an unseen name n , it defines $ID(n)$ to be the next available ID. The instrumentation tool outputs the definition of ID for subsequent phases of SA4U.

3.2 Type Deduction

The goal of the type deduction stage is to assign unit types to the program variables that appear in the trace file produced by the simulation stage. To accomplish this, we derive the unit types from several *quantities of interest*. Quantities of interest are measurements with types known a priori (e.g., vehicle altitude in meters). We trace the values of quantities of interest in the simulator. Then, we check if one of several relationships hold between program variables and each quantity of interest to deduce types.

To be precise, the input to the type deduction stage is a set of observations T obtained from the simulation stage, and the ID function from the instrumentation stage. An observation is a tuple $(\mathbb{Z}^+, \mathbb{N}, \mathbb{R})$, where the first component is a timestamp, the second

Rule	Intuition
$q_1 : U, \Box var_1 \approx q_1 \models var_1 : U$	Always approximately equal values share a type. Values that share a linear relation share a type. Prophecies tell types.
$q_1 : U, \Box var_1 \approx C_0 \times q_1 \models var_1 : C_0 \times U$	
$q_1 : U, \Box var_1 \approx x \implies \Diamond q_1 \approx x \models var_1 : U$ (and the reverse)	

Table 1: Type deduction rules used by SA4U.

component is a unique variable ID (i.e. from the *ID* function), and the third component is an observed value.

Before performing type deduction, we apply several filters to program variables to reduce false positives. First, we remove all program variables whose values are a constant < 10 or whose static type is an **enum**. Small constants are unlikely to be related to physical quantities, and enums definitely are not. Second, we remove variables that are only written once. This is because the written-once variables are not related to quantities of interest, which represent ever-changing aspects of the environment.

Table 1 uses linear temporal logic (LTL) to describe the rules SA4U uses to assign types to some (but not all) of the program variables that appear in T . In the table, var_1 and var_2 refer to arbitrary variable IDs that appear in T , and q_1 refers to an arbitrary quantity of interest. In LTL, the \Box operator means “always” and the \Diamond operator means “eventually.” Notice that our rules accept *approximate* equality. We say that $var \approx q$ if $|var - q|/|q| < \epsilon$. We permit approximate equality because we do not track every value written to both var and q . Section 5.2.6 discusses the sensitivity of our rules with respect to ϵ .

To summarize the rules in Table 1 in plain-speak:

- (1) Rule 1 says that if var_1 is always approximately q_1 , then var_1 and q_1 share a unit type. This allows SA4U to deduce the type of `_angle` in Figure 2, since `_angle`’s value is close to the angle between the UAS and its closest obstacle.
- (2) Rule 2 says that if var_1 has a linear dependence on q_1 , then the linear dependence can be applied across their types.
- (3) Rule 3 says that if var_1 predicts a value that a quantity q_1 eventually takes on, then var_1 and q_1 share a type. This allows SA4U to learn types like `target_altitude: m` since the UAS always reaches its target altitude.

The type deduction stage applies its rules in the order shown in Table 1. As types are mined, a type database is assembled. The type database relates variable IDs to a unit type. Formally, the type database is a mapping $DB : \mathbb{N} \rightarrow Type$. In the case of rules 1 and 3, the type is the same as the associated quantity of interest’s. In the case of rule 2, the type is named by the mined linear relationship.

3.3 Static Analysis for Type Inference and UTE Detection

SA4U’s static analysis stage uses *protocol files* that describe messages exchanged between the user’s control computer and the UAS for type information. Static analysis also uses the type database generated in the type deduction stage. Combined, this allows SA4U to infer the types of a large number of program variables.

Before discussing the static analysis component, we need to introduce our definitions of the subtype relation \sqsubseteq and the operator $op \in \{+, -, \times, /\}$ between two unit types.

```

1 <msg id="103" name="VISION_SPEED_ESTIMATE">
2   <description>Speed estimate.</description>
3   <field type="frame"
4     name="frame">
5     Frame
6   </field>
7   <field name="usec"
8     units="us">
9     Timestamp
10  </field>
11  <field name="x"
12    units="m/s">
13    Global X speed
14  </field>
15 </msg>

```

Figure 5: Part of MAVLink’s protocol file.

- (1) Let unit types $U_1 = (D_1, F_1)$ and $U_2 = (D_2, F_2)$. We say that $U_1 \sqsubseteq U_2 \iff (D_1 = D_2) \wedge (F_1 = F_2 \vee F_2 = Any)$.
- (2) Consider $U_1 = (D_1, F)$ and $U_2 = (D_2, F)$, where $D_1 = c_1 * b_1^{i_1,1} * \dots * b_n^{i_1,n}$ and $D_2 = c_2 * b_1^{i_2,1} * \dots * b_n^{i_2,n}$. We define $U_1 \times U_2 = ((c_1 + c_2) * (b_1^{i_1,1+i_2,1}) * \dots * (b_n^{i_1,n+i_2,n}), F)$. We use a similar definition for division.
- (3) Consider $U_1 + U_2$. Without loss of generality, assume that $U_1 \sqsubseteq U_2$. We define $U_1 + U_2 = U_2$. We use a similar definition for subtraction.

Note that we add rather than multiply the scalar coefficients of units in the definition of the \times and $/$ operators. This is because we represent scalar coefficients as the \log_{10} of the actual scalar coefficient. For example, we represent the unit `centimeter` as $-2 * meter$ instead of $\frac{1}{100} * meter$. We avoid multiplying scalar coefficients because the type checker would need to solve a system of non-linear equations. This is Hilbert’s 10th problem, which is famously undecidable. Since we restrict the scalar coefficient of the unit component to be a rational number, checking if the scalar coefficients are compatible is decidable in linear time. This is because the problem is equivalent to deciding a system of linear Diophantine equations. We select the \log_{10} representation of scalar coefficients because most practical units have a scalar coefficient that is a power of ten.

3.3.1 Protocol Files. Figure 5 shows an example message from MAVLink’s protocol file. Firmwares such as ArduPilot and PX4 often contain generators that process the protocol file and create structs to represent each message. For example, `mavlink_vision_speed_estimate_t` in Figure 5. Message fields (e.g., `usec`) are struct members. Protocol files precisely define hundreds of messages and thousands of fields, and thus serve as a great source of unit types.

Name	Rule
Binary Operator Judgement	$\frac{\Gamma \vdash e_1 : U_1, e_2 : U_2}{\Gamma \vdash e_1 \text{ op } e_2 : U_1 \text{ op } U_2}$
Assignment Judgement	$\frac{\Gamma \vdash e : U_1, v : U_2 \quad \vdash U_1 \sqsubseteq U_2}{\Gamma \vdash v = e : U_2, \Gamma\{v \rightarrow U_2\}}$
Conditional Refinement	$\frac{\Gamma, \Sigma \vdash \text{var}_1 = \text{val} \implies \text{var}_2 : U \quad \vdash \bar{s} : \Gamma\{\text{var}_2 \rightarrow U\}, \Sigma}{\Gamma, \Sigma \vdash \text{if } (\text{var}_1 == \text{val}) \bar{s} : \Gamma\{\text{var}_2 \rightarrow U\}}$
Variable Type Inference	$\frac{\Gamma, \gamma \vdash T : U}{\Gamma, \gamma \vdash T v : U, \Gamma\{v \rightarrow U\}}$
Argument Type Inference	$\frac{\Gamma \vdash e_1 : U_1, \dots, e_n : U_n}{\Gamma \vdash f(e_1, \dots, e_n) \implies \text{ArgType}(f, 1) = U_1, \dots, \text{ArgType}(f, n) = U_n}$
Return Type Inference	$\frac{\Gamma \vdash \text{Return Type}(f) = U}{\Gamma \vdash f(\dots) : U}$

Table 2: The type inference rules used by SA4U in the static analysis stage.

We use a simple scheme to represent compound data types as variables. If a struct \mathbf{s} has a member \mathbf{a} , we treat $\mathbf{s.a}$ as its own variable. We make a simplified assumption that all elements of an array share a type. So, we treat an array access $\mathbf{array}[i]$ as an access to the variable \mathbf{array} . This helps us perform limited analysis on pointers and support arrays with variable lengths.

SA4U needs protocol files because its type deduction stage cannot learn the types of protocol messages. This is unfortunate, since MAVLink handlers are a dominant source of UTEs. Figure 5 illustrates the problem: the type of x depends on the value of the field frame. We introduce the constraint set Σ to account for this kind of relationship. Σ is a set of relationships in the template $\text{var}_1 = v_1 \implies \text{var}_2 : U$. In the case of the example in Figure 5, Σ contains the relationship $\text{vision_speed_estimate.frame} = \text{GLOBAL} \implies \text{vision_speed_estimate.x} : (1 \times m \times s^{-1}, \text{GLOBAL})$.

Protocol files provide two pieces of information used in the static analysis stage. The first piece of information is γ , which relates program variable types (e.g., $\text{vision_speed_estimate.t}$) to unit types. The second piece of information is the set of control relationships Σ .

3.3.2 Type Inference and UTE Detection. The inputs to the type inference system are the stack-frame model Γ , the relationship between program types and unit types γ , the control relationship between variables Σ , and the type database DB . Γ tracks the types of variables in the current stack-frame. Table 2 describes the type inference rules SA4U applies in the static analysis stage.

The binary operator judgement rule allows SA4U to perform dimensional analysis on expressions involving multiple types. Informally, it reads that if the current stack frame Γ assigns the type U_1 to e_1 and U_2 to e_2 , then the type of the expression $e_1 \text{ op } e_2$ is $U_1 \text{ op } U_2$. For example, $x : (\mathbf{m}, \text{GLOBAL}) / t : (\mathbf{s}, \text{GLOBAL})$ is assigned the type $(\mathbf{m} * \mathbf{s}^{-1}, \text{GLOBAL})$ according to this rule. Note that SA4U reports a type error when a binary operator cannot be

typed according to this rule. For example, SA4U will report a UTE error if it finds $x : (\mathbf{m}, \text{GLOBAL}) + t : (\mathbf{s}, \text{GLOBAL})$.

The assignment judgement rule prevents illegal stores to program variables. Informally, it says that if an expression e has the type U_1 in the current stack frame (Γ) and the expression v has the type U_2 in Γ , then $e = v$ can be assigned the type U_2 only if $U_1 \sqsubseteq U_2$. The unit type of a variable is initially unconstrained, and the frame is **Any**. This allows gradual type refinement through assignment. Similar to the case of the binary operator judgement rule, SA4U reports an error if an assignment expression is untypeable.

The conditional refinement rule helps SA4U use type information from protocol files. Since protocol files often use *control fields* in messages to indicate the frames of other message members, we must partially interpret the source code to approximate the values that may be present in control fields. The rule in Table 2 shows how SA4U uses information in conditional branches to refine the conditional type. Informally, it says that if there is a control relationship between var_1 and var_2 , and the conditional statement's body \bar{s} can be typed using the control relationship, then a conditional statement with suitably established value of var_1 infers that \bar{s} can be typed. Note that we do not consider the myriad of ways that a conditional statement could establish the value of var_1 . Instead, we observe that developers use control relationships in simple ways. They compare the value of control fields directly to constants in either if statements or switch statements.

The variable type inference rule uses the unit types extracted from the protocol files to type program variables. Informally, it says that if a static type (i.e., a type in the actual program) T has the unit type (i.e., the type used in dimensional analysis) U , then witnessing the declaration of a variable with type T updates the stack model Γ with v assigned U . This allows SA4U to handle the static types from entries in the protocol file.

The argument type inference rule checks type consistency across multiple call sites. We build a free function called `ArgTypes` that assigns a type to argument i of a function f . If we can find a model of `ArgTypes` then argument types are consistently provided throughout the source code. To reduce false positives, we ignore calls to functions in the standard library, and maintain a small list of function types to ignore.

Finally, the return type inference rule checks type consistency in return values between multiple call sites. Similar to argument type inference, we use a free function to model the return type of procedures. If return types are used in inconsistent ways then the constraint system will be unsatisfiable.

SA4U applies the inference rules shown in Table 2 to each function in the firmware’s source code to generate a type-model of the program. Specifically, when SA4U encounters an expression that matches the expression on the top of a proof bar in the table, SA4U generates the constraint shown on the bottom of the proof bar. After SA4U has parsed the entire source code, SA4U invokes the Z3 theorem prover [12] to check that the constraints are satisfiable. If they are satisfiable, then the type rules are satisfied by the subject source code. Otherwise, SA4U generates a bug report that summarizes the issue.

4 IMPLEMENTATION

This section discusses the implementation. SA4U contains three tools: one for instrumentation, one for type deduction, and one for static analysis. We also discuss workloads and simulation. All of SA4U is open-source, and is available at <https://github.com/obicons/sa4u>.

4.1 Instrumentation

SA4U’s instrumentation tool is a Clang tool [6] that rewrites assignment expressions. Clang tools provide a library to use the front end of the Clang C++ compiler, simplifying the task of source code transformer development. Assignment expressions are rewritten to invoke an instrumentation macro. The instrumentation macro logs the values of program variables to a CSV file.

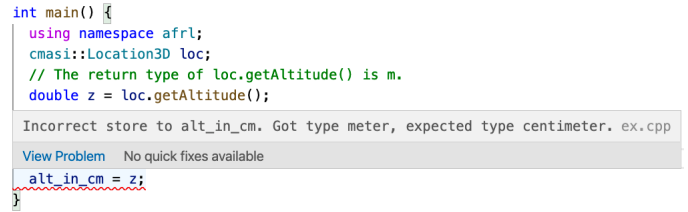
4.2 Simulation

Step 2 of the SA4U’s workflow requires the developer to execute the firmware in a physics simulator. We chose to use Gazebo [19] because it is robust and both ArduPilot and PX4 support it. Any simulator works for this job, but small code changes are required to track quantities of interest.

Users must provide a workload during simulation. A ideal workload contains enough complex behavior so that the simulation will cover most of the firmware code. In practice, obtaining suitable workloads is straight-forward. For example, both ArduPilot and PX4 contain end-to-end functional tests that we used as a workload.

4.3 Type Deduction

We implemented SA4U’s type deduction tool in Python. Initially, we tried to use Daikon [14], a popular invariant detector. However, we discovered that Daikon’s execution took far too long (on the order of days) even with our effort on restricting outputs. This is mainly



```
int main() {
    using namespace afl;
    cmasi::Location3D loc;
    // The return type of loc.getAltitude() is m.
    double z = loc.getAltitude();
}

Incorrect store to alt_in_cm. Got type meter, expected type centimeter. ex.cpp
View Problem No quick fixes available
alt_in_cm = z;
```

Figure 6: Screenshot of the SA4U Editor Plugin. SA4U reports an error based on the contents of the protocol file and the type database.

because Daikon mines general invariants that SA4U does not need. So we created a simple invariant miner for type deduction.

4.4 Static Analysis

Our static analysis engine is a Python tool that uses `libclang` [7]. `libclang` provides simple bindings to the Clang compiler’s front end. This allows our static analysis tool to analyze any code that Clang can compile. Similar to the instrumentation engine, the static analysis tool depends on widely available compilation databases. We use the Z3 theorem prover [12] to check the inference rules shown in Table 2. Z3 is an efficient formula solver, and is a popular choice for implementations similar to ours. The static analysis engine accounts for the majority of our implementation burden.

4.5 Editor Plugin

We created a prototype editor plugin to make SA4U easier to use. Figure 6 shows the SA4U editor plugin in action. We wrote a language server protocol (LSP) server that communicates with the developer’s text editor. Almost all text editors (even Emacs and Vim) can use LSP servers. We tested our LSP server with a VSCode plugin. Our LSP server is a thin wrapper around the SA4U analysis binary. The server simply runs SA4U and scans its output for issues. Overall, our implementation was painless. We only had to write 150 lines of code to create the server and extension.

5 EVALUATION

Our evaluation addresses three research questions: (1) How effective is SA4U at detecting UTEs in UAS firmware? (2) How does SA4U compare to the state-of-the-art? (3) How efficient is SA4U?

5.1 Experiment Methodology

5.1.1 Platform and Target Firmware. We conducted our experiments on an Ubuntu 18.04 server equipped with an Intel Core i5-2500 CPU and 8GB of memory. We evaluated SA4U on ArduPilot [1] and PX4 [4]. We chose these systems for several reasons. First, ArduPilot and PX4 are the most popular open source UAS firmware. Combined, they have more stars on Github than any other UAS project. Second, UAS firmware is a complex category of control software, whose safety is paramount. There are many onboard sensors, complex control equations, and remote message systems. UTEs in any component could lead to disastrous scenarios. Finally, we want to show that SA4U works independently of the firmware’s hardware abstraction mechanisms.

Bug ID	UTE Features			Found by	
	Same dimension?	Frame?	MAV?	SA4U	Phys
APM-16903	No	Yes	Yes	✓	✗
APM-19868	No	Yes	Yes	✓	✗
APM-21291	No	Yes	Yes	✓	✗
APM-21309	No	Yes	Yes	✓	✗
PX4-17337	No	Yes	Yes	✓	✗
PX4-17354	No	Yes	Yes	✓	✗
PX4-19973	No	Yes	Yes	✓	✗
PX4-19983	No	Yes	Yes	✓	✗
PX4-19995	No	Yes	Yes	✓	✗
PX4-20000	No	Yes	Yes	✓	✗
PX4-20018	No	Yes	Yes	✓	✗
PX4-20019	No	No	Yes	✓	✗
PX4-GPSDrivers-110	Yes	No	No	✓	✗
PX4-20023	Yes	Yes	No	✓	✗

Table 3: Previously unknown bugs detected by SA4U and state-of-the-art tool Phys.

5.1.2 Workload for Type Deduction. We used the same workload for ArduPilot and PX4 in the type deduction phase. The workload guides the UAS to takeoff, navigates to 8 waypoints, and then lands. We obtained this workload from ArduPilot’s existing test files.

5.1.3 SA4U’s Effectiveness. We applied SA4U to our target systems to evaluate how effectively it detects both known and unknown bugs. For known bugs, we used 6 bugs from a recent bug study [26] to see whether SA4U can detect them. For unknown bugs, we evaluated SA4U on the latest versions of ArduPilot and PX4.

We compared SA4U with the state-of-the-art tool Phys for detecting unknown bugs. To make Phys work on ArduPilot and PX4, we used SA4U’s code to annotate MAVLink variables. Additionally, we evaluated the false positive rates of SA4U using 5 random files from ArduPilot and PX4.

5.1.4 SA4U’s Efficiency and Sensitivity. We conducted experiments to evaluate SA4U’s execution time when running it with the latest versions of ArduPilot and PX4. The execution time include SA4U’s analysis and Z3’s type inference. Also we compare SA4U’s execution time with Phys’s execution time over the same firmware. Furthermore, we evaluated the sensitivity of SA4U’s type deduction engine on the approximation threshold ϵ .

5.2 Experimental Results

Bug ID	Frame?	Found by SA4U?
APM-3542	Yes	✓
APM-4105	Yes	✓
APM-4550	Yes	✗
PX4-12517	Yes	✓
PX4-12532	Yes	✓
PX4-13180	Yes	✓

Table 4: Known bugs used to evaluate SA4U.

5.2.1 Detecting Known Bugs. Table 4 shows the detection result of SA4U on 6 known bugs from a recent bug study [26]. We selected these 6 bugs because other bugs were found in old versions of the firmware that we could not easily build in our environment. SA4U must be able to build the firmware in order to analyze it due to its dependency on libclang. Note that it is impossible to report a traditional recall statistic because the number of UTEs in the evaluation subjects is unknown. In the table, the bug ID column identifies the bug in the repository’s bug management system. The frame column shows if a bug involves a frame of reference. SA4U was able to diagnose 5 of the 6 bugs. APM-4550 was not diagnosed because SA4U incorrectly infers that a conversion of reference frame takes place when a variable is multiplied, however the multiplication was not a conversion.

5.2.2 Detecting Unknown Bugs. Table 3 shows the previously unknown bugs detected by SA4U. The unknown UTEs reported by SA4U have various features, errors within a dimension or between different dimensions, errors involving different frames of reference or within the same frame, and errors involving MAVLink message handlers. This is mainly because SA4U’s interprocedural analysis and its type representation precisely capturing the unit information within a dimension and the frame information of the unit.

We noticed that most of the UTEs listed in Table 3 span multiple procedures. We believe this is because developers often manually test new code or code revision, and even simple manual tests are likely to expose UTEs contained in a single function. Moreover, developers can easily remember the types of local variables while they write a procedure.

However, interprocedural flows are tricky. Developers are likely to misunderstand function contracts because they are not formally documented. This leads to UTEs. Another possible reason is that developers could refactor a function and violate its contract, thus breaking previously working code. Finally, developers must keep up with the latest changes to protocol specifications (e.g., MAVLink often adds new message fields). If firmware falls behind, the values sent by pilot control computers will be misinterpreted by the

firmware. For example, the bug shown in Figure 8 was caused by developers failing to notice a change in MAVLink’s specification.

5.2.3 Comparing with the state-of-the-art. We compared SA4U with the state-of-the-art tool in this space: Phys [18]. SA4U was able to find bugs in the subject systems that Phys could not detect for two reasons. First, because Phys does not consider the reference frame in its type representation. Observe that the reference frame was involved in 13 of the 14 bugs. Second, Phys is unable to detect 2 of the 14 bugs because it lacks the precise unit type information that SA4U uses. For example, in PX4-GPSDrivers-110, a driver fails to convert from milliseconds to microseconds.

Direct comparison with PhysFrame is not possible because PhysFrame is designed specifically for ROS programs. However, PhysFrame struggles with the bugs shown in Table 3 at a conceptual level. Notice that 12 of the 14 bugs involve MAVLink. SA4U introduces the *conditional refinement* rule in Table 2 to assign types to expression whose type depends on the value of a control field. Since PhysFrame lacks a corresponding rule, it conceptually struggles to detect these bugs.

5.2.4 False Positives and False Negatives. To evaluate false positives, we looked at the outputs of SA4U for 5 random files. SA4U found 32 errors in the 5 files. While this seems high, there are duplicates in the reports. For example, if a bug occurs in a function defined in a header file and the function is called in multiple source files, then there will be multiple bug reports. Among those 32 errors, we can confirm that there are 20 true errors (including duplicates). This yields a false positive rate of approximately 38%.

There are multiple sources of these false positives. First, the type deduction phase can derive an incorrect type for a program variable. For example, the program variable `emergency_mode_alt` could be assigned the type `centimeter`, when the actual type is `meter`. When this occurs, SA4U’s static analysis phase reports a false positive. Second, the static analysis phase itself is not sound. For example, developers use conversion functions to transform the types of values. SA4U is not aware of all conversion functions, so it misdiagnoses values whose types that are transformed this way.

It is worth noting that SA4U is neither sound (i.e., SA4U has false positives) nor complete (i.e., SA4U has false negatives). Once again, the type deduction phase could incorrectly label the type of a program variable. If this occurs, the static analysis phase could fail to diagnose a UTE. Moreover, the static analysis phase can only detect errors if a sufficiently large number of types have been inferred. Finally, sometimes programmers temporarily store values with incorrect types in variables. SA4U reports false positives in this case.

5.2.5 SA4U’s Efficiency. Table 5 shows SA4U’s execution time on its evaluation subjects. SA4U is efficient to detect UTEs in large code bases. For example, it took SA4U about 37 minutes to scan 848,562 lines of code in 7,320 files in ArduPilot.

SA4U has a speedup of up to 2.16x over the state-of-the-art Phys. Note that SA4U’s analysis is almost identical to Phys’. This speedup is mainly attributed to implementation differences in SA4U. Specifically, SA4U uses Z3, a high-performance theorem prover, to perform type inference. Z3 performs type inference in parallel, thus giving SA4U a nice speedup.

ϵ	Approximate	Linear	Eventually
100%	16823	343	58
97.5%	N/A	1331	8613
95%	N/A	1569	19841
90%	3625	2248	32754
80%	3610	3707	41287
70%	3594	5916	53003
60%	3501	9097	63867
50%	3493	14539	87498
40%	3470	20857	107657
30%	3463	32131	130617
20%	3444	51887	159957
10%	3444	90694	252289
5%	3382	N/A	318665
2.5%	3373	N/A	322034
1%	3363	N/A	322056

Table 6: The effect of ϵ on type deduction sensitivity. “Approximate” refers to rule 1 in Table 1, “Linear” refers to rule 2, and “Eventually” refers to rule 3.

UTE Tool	Firmware	LoC	Runtime (seconds)
SA4U	ArduPilot	848,562	2215
	PX4	197,795	689
Phys	ArduPilot	848,562	4792
	PX4	197,795	3951

Table 5: Runtime of UTE detection tools..

5.2.6 Invariant Mining Sensitivity. Table 6 shows the effect of the choice of ϵ on the number of types mined by the type deduction engine. ϵ ’s meaning is overloaded for each invariant template. Recall that ϵ allows us to tolerate the measurement errors we introduced by only periodically sampling the values of variables. In the case of **Approximate**, ϵ controls the relative error between the measurements. So, as the permissible relative error becomes smaller, fewer invariants are mined. In the case of **Linear**, ϵ controls the absolute value of the Pearson correlation coefficient where we accept the relationship. As the Pearson correlation approaches 0, the relationship between two variables is less linear. Finally, in the case of **Eventually**, ϵ controls the confidence threshold where the mined invariant is accepted.

We used the values in Table 6 to select ϵ when we ran SA4U’s type deduction engine. We selected $\epsilon = 5\%$ for **Approximate** invariants because the extra mined relationships were useful versus the 2.5% and 1% levels. For **Linear** and **Eventually** invariants we used 0.975.

6 CASE STUDIES

We discuss the following two UTE cases reported by SA4U in more details to illustrate its detection capability.

6.0.1 APM-19868. Figure 7 shows a UTE that occurred in ArduPilot. MAVLink defines the `vision_speed_estimate_t` message to communicate the estimated speed of the UAS based on the input from an external visual navigation system. The `vision_speed_estimate_t`

```

1 void on_vsn_spd_est(vision_speed_estimate_t
2   e) {
3   ...
4   e.usec = corrector.correct_time(
5     e.usec,
6     lcl_time()
7   );
8   ...
9 }
10 // Accounts for transport delay for the
11 // timestamp rmt received at time local.
12 int Corr::correct_time(int rmt, int lcl) {
13   ...
14   int diff = lcl - link_offset;
15   est_rmt_time = diff;
16   link_offset = est_rmt_time - rmt;
17   ...
18 }

```

Figure 7: APM-19868: A UTE affecting ArduPilot.

message contains the field `usec` that communicates the time when the estimate was generated. This field is measured in microseconds, storing either the time since the system booted, or the time since the UNIX epoch (i.e. January 1, 1970). The receiver is expected to use the magnitude of the field to determine the timestamp’s format.

Line 12 of Figure 7 shows ArduPilot’s handler for the `vision_speed_estimate_t` message. SA4U’s instrumentation program adds instrumentation to line 15 since it contains a store to a non-local variable. Next, we simulate the UAS flight to obtain a trace file. Then, the type deduction program analyzes the values stored in `est_rmt_time`. The quantity of interest that `est_rmt` best matches is microseconds since system boot. This type information is recorded to the type database.

Then, SA4U runs the static analysis stage and parses the source code that contains the functions shown in Figure 7. When parsing line 3, SA4U knows that the type of `e.usec` can be in either UNIX epoch format (i.e. `TIME_UNIX`) or time since system boot (i.e. `TIME_BOOT`) because of the contents of the protocol file. So, SA4U generates the constraint `ArgType(Corr::correct_time, 1) = (1e-6 * s, {TIME_BOOT, TIME_UNIX})`. Later, When encountering line 14, SA4U applies its binary operator judgement rule and generates the constraints `type(local) = type(link_offset)` and `type(diff) = type(lcl)`. Next, SA4U generates the constraint `type(est_rmt_time) = type(diff)` when it parses line 15. Notice that this simplifies to `type(est_rmt_time) = type(lcl)`. Finally, SA4U once again applies the binary operator judgement rule to generate the constraint that `type(est_rmt_time) = ArgType(Corr::correct_time, 1)` when it parses line 16. But this constraint is unsatisfiable due to a conflict between the constraints from the type database and first constraint SA4U introduced: `ArgType(Corr::correct_time, 1) = (1e-6 * s, {TIME_BOOT, TIME_UNIX})`. Since the constraint cannot be satisfied, SA4U creates a UTE bug report.

```

1 ImageTargetHandler handler;
2 void on_landing_target(landing_target &t) {
3   if (t.has_pos && t.frame == LOCAL) {
4     handle_pos_target(t);
5 +  } else if (t.has_pos) {
6 +     log("Unsupported frame.");
7   } else {
8     handler.set_target(t);
9   }
10 }
11
12 void ImageTargetHandler::set_target(
13   landing_target &t) {
14   ...
15   _x = t.img_angle_x;
16   _y = t.img_angle_y;
17 }

```

Figure 8: PX4-17354: A complex UTE SA4U found in PX4.

We reported this bug to the developers of ArduPilot. They confirmed that this is indeed a UTE, however no patch has been accepted at the time of writing. Developers are still unsure about the best way to change `Corr::correct_time` to remove the bug.

6.0.2 *PX4-17354*. Figure 8 shows a UTE that occurred in PX4. The `LANDING_TARGET` MAVLink message communicates a landing location to the UAS. This message supports two different uses. First, pilots can command the UAS to land in a location in a reference picture. Second, pilots can command the UAS to land at a position specified as GPS coordinates. In the second case, the landing target can be specified in any coordinate system.

Line 4 of Figure 8 checks if the coordinate system of the landing target is local since PX4 only supports the local coordinate system for landing messages. However, if the frame is not local, the else branch on line 8 is taken. This causes read accesses to two image fields (`img_angle_x` and `img_angle_y`) of `t` at lines 15 and 16, respectively. However, the message may contain a landing position specified in GLOBAL coordinate system. Thus, the two image fields at lines 15 and 16 are undefined and these accesses are incorrect.

SA4U identifies this bug, albeit somewhat crudely. First, SA4U correctly learns the types of `ImageTargetHandler::_x` and `ImageTargetHandler::_y`. Then, SA4U performs static analysis. In the `else` branch on line 8 of Figure 8, SA4U represents the coordinate system of each of `t`’s members as any possible coordinate system, minus local. Then, SA4U sees the function call on line 9. This call must be incorrect since the possible frames of `t` are not compatible with the stores at lines 15 and 16.

We patched the bug by introducing the code shown on lines 6 and 7. When we reported the bug and submitted our patch to PX4’s developers, they confirmed the bug and accepted our change.

7 RELATED WORK

Dimensional Analysis. Dimensional analysis is a widely used technique to validate equations. Early work provided programming language support and packages, e.g. [16]. Osprey [15] used dimensional analysis to validate C programs, but required manual

annotations. Later work (e.g. [18], [17], [22]) tried to reduce annotation burden several ways. Phys [18] uses variable names to infer unit types. PhysFrame [17] extends Phys to also consider the frame of the unit. PhysFrame is a well-designed tool, but it is built specifically for programs that use ROS. It cannot be trivially applied to the systems studied here. Phriky-units [23] and [22] reduces the annotation burden by pre-annotating shared libraries, and then using type inference to deduce the types of program variables.

Invariant Mining. Mining program invariants is a mature idea. Daikon [14] is the most prolific example. Later work (e.g. [9]) extends invariant mining to capture temporal invariants, an idea we use in our eventually-equal mining rule. Researchers often use mined invariants in program analysis. For example, CoFi [11] mines invariants in distributed system and injects faults at points where invariants do not hold. Similar to our work, MonkeyType [21] runs Python unit tests to discover likely program variable types.

8 CONCLUSION

We presented SA4U, a tool for finding UTEs in real UAS firmware. SA4U obtains type information from traces of program variables, and protocol files. SA4U partially interprets the source code of UAS to constrain the unit types of program variables. This allows SA4U to analyze message handlers of common protocols. Then, SA4U applies dimensional analysis to infer the types of variables not known from other sources. In the future, we wish to extend SA4U in two ways. First, we plan to ease developer burden by developing repair tools to generate patches for UTEs. Second, we hope to integrate other source of type information (e.g., Phys) with SA4U.

ACKNOWLEDGMENTS

The authors would like to thank the anonymous reviewers for their valuable feedback and thoughtful suggestions. This work is partially sponsored by the grants: AFRL FA864921P0206, NSF 1901242, and ONR N000142012733.

REFERENCES

- [1] 2021. ArduPilot. <https://ardupilot.org>.
- [2] 2021. CaliperSharp. <https://github.com/point85/CaliperSharp>.
- [3] 2021. MAVLink. <https://mavlink.io>.
- [4] 2021. PX4: Open Source Autopilot for Drone Developers. <https://px4.io>.
- [5] 2021. ROS - Robot Operating System. <https://www.ros.org>.
- [6] 2022. Clang Tools. <https://clang.llvm.org/docs/ClangTools.html>.
- [7] 2022. libclang. https://clang.llvm.org/doxygen/group__CINDEX.html.
- [8] Eric Allen, David Chase, Victor Luchangco, Jan-Willem Maessen, and Guy L. Steele Jr. 2004. Object-oriented units of measurement. In *Proceedings of the 19th annual ACM SIGPLAN conference on Object-oriented programming, systems, languages, and applications*. 384–403.
- [9] Ivan Beschastnikh, Yuriy Brun, Michael D. Ernst, Arvind Krishnamurthy, and Thomas E. Anderson. 2011. Mining Temporal Invariants from Partially Ordered Logs. In *Managing Large-Scale Systems via the Analysis of System Logs and the Application of Machine Learning Techniques* (Cascais, Portugal) (SLAML '11). Association for Computing Machinery, New York, NY, USA, Article 3, 10 pages. <https://doi.org/10.1145/2038633.2038636>
- [10] Bureau International des Poids et Mesures. 2019. The International System of Units (SI). <https://www.bipm.org/documents/20126/41483022/SI-Brochure-9-EN.pdf/2d2b50bf-f2b4-9661-f402-5f9d66e4b507>.
- [11] Haicheng Chen, Wensheng Dou, Dong Wang, and Feng Qin. 2020. CoFi: consistency-guided fault injection for cloud systems. In *Proceedings of the 35th IEEE/ACM International Conference on Automated Software Engineering*. 536–547.
- [12] Leonardo de Moura and Nikolaj Bjørner. 2008. Z3: An Efficient SMT Solver. In *Tools and Algorithms for the Construction and Analysis of Systems*, C. R. Ramakrishnan and Jakob Rehof (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 337–340.
- [13] Matthew Duquette. [n.d.]. *The Common Mission Automation Services Interface*. <https://doi.org/10.2514/6.2011-1542> arXiv:<https://arc.aiaa.org/doi/pdf/10.2514/6.2011-1542>
- [14] Michael D. Ernst, Adam Czeisler, William G. Griswold, and David Notkin. 2000. Quickly detecting relevant program invariants. In *ICSE 2000, Proceedings of the 22nd International Conference on Software Engineering*. Limerick, Ireland, 449–458.
- [15] Lingxiao Jiang and Zhendong Su. 2006. Osprey: a practical type system for validating dimensional unit correctness of C programs. In *Proceedings of the 28th international conference on Software engineering*. 262–271.
- [16] Michael Karr and David B. Loveman. 1978. Incorporation of Units into Programming Languages. *Commun. ACM* 21, 5 (may 1978), 385–391. <https://doi.org/10.1145/359488.359501>
- [17] Sayali Kate, Michael Chinn, Hongjun Choi, Xiangyu Zhang, and Sebastian Elbaum. 2021. PHYFRAME: Type Checking Physical Frames of Reference for Robotic Systems. In *Proceedings of the 29th ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering (Athens, Greece) (ESEC/FSE 2021)*. Association for Computing Machinery, New York, NY, USA, 45–56. <https://doi.org/10.1145/3468264.3468608>
- [18] Sayali Kate, John-Paul Ore, Xiangyu Zhang, Sebastian Elbaum, and Zhaogui Xu. 2018. Phys: Probabilistic Physical Unit Assignment and Inconsistency Detection. In *Proceedings of the 2018 26th ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering (Lake Buena Vista, FL, USA) (ESEC/FSE 2018)*. Association for Computing Machinery, New York, NY, USA, 563–573. <https://doi.org/10.1145/3236024.3236035>
- [19] Nathan Koenig and Andrew Howard. [n.d.]. Design and use paradigms for gazebo, an open-source multi-robot simulator. In *2004 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*(IEEE Cat. No. 04CH37566), Vol. 3. IEEE, 2149–2154.
- [20] Steve McKeever, Oscar Bennich-Björkman, and Omar-Alfred Salah. 2021. Unit of measurement libraries, their popularity and suitability. *Software: Practice and Experience* 51, 4 (2021), 711–734.
- [21] Meyer, Carl. 2017. Let your code type-hint itself: introducing open source MonkeyType. <https://instagram-engineering.com/let-your-code-type-hint-itself-introducing-open-source-monkeytype-a855c7284881>.
- [22] John-Paul Ore, Carrick Detweiler, and Sebastian Elbaum. 2017. Lightweight Detection of Physical Unit Inconsistencies without Program Annotations. In *Proceedings of the 26th ACM SIGSOFT International Symposium on Software Testing and Analysis (Santa Barbara, CA, USA) (ISSTA 2017)*. Association for Computing Machinery, New York, NY, USA, 341–351. <https://doi.org/10.1145/3092703.3092722>
- [23] John-Paul Ore, Carrick Detweiler, and Sebastian Elbaum. 2017. Phriky-Units: A Lightweight, Annotation-Free Physical Unit Inconsistency Detection Tool. In *Proceedings of the 26th ACM SIGSOFT International Symposium on Software Testing and Analysis (Santa Barbara, CA, USA) (ISSTA 2017)*. Association for Computing Machinery, New York, NY, USA, 352–355. <https://doi.org/10.1145/3092703.3098219>
- [24] Kathy Sawyer. 1999. Mystery of Orbiter Crash Solved. <https://www.washingtonpost.com/wp-srv/national/longterm/space/stories/orbiter100199.htm>.
- [25] Matthias Schabel and Steven Watanabe. 2008. Boost.Units 1.0.0. <http://boost.cowic.de/rc/pdf/units.pdf>.
- [26] Max Taylor, Jayson Boubin, Haicheng Chen, Christopher Stewart, and Feng Qin. 2021. A Study on Software Bugs in Unmanned Aircraft Systems. In *2021 International Conference on Unmanned Aircraft Systems (ICUAS)*. 1439–1448. <https://doi.org/10.1109/ICUAS51884.2021.9476844>
- [27] Max Taylor, Haicheng Chen, Feng Qin, and Christopher Stewart. 2021. Avis: In-Situ Model Checking for Unmanned Aerial Vehicles. In *2021 51st Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN)*. 471–483. <https://doi.org/10.1109/DSN48987.2021.00057>
- [28] Christopher Steven Timperley, Afsoon Afzal, Deborah S. Katz, Jam Marcos Hernandez, and Claire Le Goues. 2018. Crashing Simulated Planes is Cheap: Can Simulation Detect Robotics Bugs Early?. In *2018 IEEE 11th International Conference on Software Testing, Verification and Validation (ICST)*. 331–342. <https://doi.org/10.1109/ICST.2018.00040>
- [29] Massimo Trella, Ellen Herring, Richard Freeman, William Kilpatrick, Alan Reth, Michael Greenfield, John Credland, Robert Laine, Dino Machi, and Alan Smith. 1998. SOHO MISSION INTERRUPTION JOINT ESA/NASA INVESTIGATION. https://umbra.nascom.nasa.gov/soho/SOHO_final_report.html.