Data-driven Mutation Analysis for Cyber-Physical Systems

Enrico Viganò, Oscar Cornejo, Fabrizio Pastore, *Member, IEEE,* and Lionel C. Briand, *Fellow, IEEE*

Abstract—Cyber-physical systems (CPSs) typically consist of a wide set of integrated, heterogeneous components; consequently, most of their critical failures relate to the interoperability of such components. Unfortunately, most CPS test automation techniques are preliminary and industry still heavily relies on manual testing. With potentially incomplete, manually-generated test suites, it is of paramount importance to assess their quality. Though mutation analysis has demonstrated to be an effective means to assess test suite quality in some specific contexts, we lack approaches for CPSs. Indeed, existing approaches do not target interoperability problems and cannot be executed in the presence of black-box or simulated components, a typical situation with CPSs. In this paper, we introduce *data-driven mutation analysis*, an approach that consists in assessing test suite quality by verifying if it detects interoperability faults simulated by mutating the data exchanged by software components. To this end, we describe a data-driven mutation analysis technique (*DaMAT*) that automatically alters the data exchanged through data buffers. Our technique is driven by fault models in tabular form where engineers specify how to mutate data items by selecting and configuring a set of mutation operators. We have evaluated *DaMAT* with CPSs in the space domain; specifically, the test suites for the software systems of a microsatellite and nanosatellites launched on orbit last year. Our results show that the approach effectively detects test suite shortcomings, is not affected by equivalent and redundant mutants, and entails acceptable costs.

Index Terms—Mutation analysis, Cyber-Physical Systems, CPS Interoperability, Integration testing

1 Introduction

Cyber Physical Systems (CPSs) are heterogeneous systems that integrate computation, networking, and physical processes that are deeply interlaced [1]. In CPSs, conformance with requirements is verified through test cases executed at different development stages, based on available development artifacts [2]. In this paper, we focus on the identification of faults in the executable software to be deployed on the CPS, and thus target software-in-the-loop (SIL) and hardware-in-the-loop (HIL) testing.

When software systems are large and integrate a diverse set of components, it is difficult to ensure that the test suite can detect any latent severe fault. To ensure test suite quality, standards for safety-critical software provide methodological guidance, for example through structural coverage adequacy; however, those strategies do not directly measure the fault detection capability of a test suite. A more direct solution to evaluate test suite quality is *mutation analysis* [3], [4]. It consists of automatically generating faulty software versions and computing the mutation score, that is, the percentage of faulty software versions leading to a test failure. Mutation analysis is a good candidate to assess test suite quality because there is a strong association between high mutation scores and high fault revealing power for test suites [5].

Most mutation analysis techniques rely on the generation of faulty software through mutation operators that modify the software implementation (either the source or the executable code). Unfortunately, such techniques suffer from two major limitations, which are critical in the CPS context: (1) they cannot identify problems related to the interoperability of integrated components (integration testing) and (2) they can be applied only to components that can be executed in the development environment. In CPSs, major problems typically arise because of the lack of *interoperability of integrated components* [6], [7], mainly due to the wide variety and heterogeneity of the technologies and standards adopted. Also, there is limited work on integration testing automation [8], thus forcing companies to largely rely on manual approaches, which are error prone and likely to lead to incomplete test suites[1]. It is thus of fundamental importance to ensure the effectiveness of test suites with respect to detecting interoperability issues, for example by making sure test cases fully exercise the exchange of all possible data items and report failures when erroneous data is being exchanged by software components. For example, the test suite for the control software of a satellite shall identify failures due to components working with different measurement systems [10]. Unfortunately, well known, code-driven mutation operators (e.g., the sufficient set [11], [12]) simulate algorithmic faults by introducing small changes into code and are thus unlikely to simulate interoperability problems resulting in exchanges of erroneous data.

The second limitation of code-driven mutation analysis approaches concerns *the incapability of injecting faults into*

- *E. Viganò, O. Cornejo, F. Pastore, and L. Briand are affiliated with SnT Centre for Security, Reliability and Trust, University of Luxembourg, Luxembourg.*
  E-mail:{enrico.vigano,oscar.cornejo,fabrizio.pastore,lionel.briand}@uni.lu
- *L. Briand also holds a faculty appointment with school of EECS, University of Ottawa, Canada.*

Manuscript received April 01, 2022; revised September 09, 2022.

[1] In this paper, we rely on input domain partitioning to determine if a test suite is complete [9].

black-box components whose implementation is not tested within the development environment (e.g., because it is simulated or executed on the target hardware). For example, in a satellite system, such components include the control software of the Attitude Determination and Control System (ADCS), the GPS, and the Payload Data Handling Unit (PDHU). During SIL testing, the results generated by such components (e.g., the GPS position) are produced by a simulator. As for HIL testing, these components are directly executed on the target hardware and cannot be mutated, either because they are off-the-shelf components or to avoid damages potentially introduced by the mutation.

An alternative to code-driven mutation analysis approaches are model-based ones, which mutate models of the software under test (SUT). Unfortunately, existing approaches do not include strategies to simulate interoperability problems; also, their primary objective is to support test generation not the evaluation of test suites [13], [14], [15], [16]. Furthermore, model-based test generation may not be cost-effective if detailed models of the system under test are not available—which is often the case, especially in early development stages—and lack key information required for testing (e.g., which telecommands shall trigger a state transition in a satellite system).

To address the above-mentioned limitations, we propose *data-driven mutation analysis*, a new mutation analysis paradigm that alters the data exchanged by software components in a CPS to evaluate the capability of a test suite to detect interoperability faults. To this end, we present a technique, *data-driven mutation analysis with tables* (*DaMAT*), to automate data-driven mutation analysis by relying on a fault model that captures, for a specific set of components, both the characteristics of the data to mutate (e.g., the size and structure of the messages generated by the ADCS) and the types of fault that may affect such data (e.g., a value out of the nominal range). The latter is formalized as a set of parameterizable mutation operators. Based on discussions with practitioners, to simplify adoption, we rely on fault models in tabular form where each row specifies, for a given data item, what mutation operator (along with its corresponding parameter values) to apply to which elements of the data item. At runtime, *DaMAT* modifies the data exchanged by components according to the provided fault model (e.g., replaces a nominal voltage value with a value out of the nominal range).

DaMAT identifies test suite shortcomings that consist of message types, software states, and input partitions not being exercised. It also addresses the lack of adequate test oracles, i.e., oracles capable of detecting observable failures caused by data (exchanged by SUT components) not being equivalent to the data assumed by test cases. We have defined three analysis metrics enabling engineers to distinguish between such shortcomings and thus guiding the improvement of test suites; they are *fault model coverage*, *mutation operation coverage*, and *covered mutation score* (see Section 3.7).

We performed an empirical evaluation of *DaMAT* to determine the effectiveness, feasibility, and applicability of data-driven mutation analysis for evaluating test suites. Our benchmark consists of software for CPSs in the space domain provided by our industry partners, which are the European Space Agency [17], GomSpace Luxembourg, a world-renowned manufacturer and supplier of nanosatellites, and LuxSpace, a European developer of infrastructure products (e.g., microsatellites) and solutions for space. More specifically, the benchmark includes (1) the on-board embedded software system for *ESAIL* [18], a maritime microsatellite recently launched into space, and (2) a configuration library used in constellations of nanosatellites [19]. Our empirical results show that *DaMAT* (1) successfully identifies different types of shortcomings in test suites, (2) prevents the introduction of equivalent and redundant mutants, and (3) is practically applicable in the CPS context.

To summarize, our contributions include:

- Data-driven mutation analysis, a new mutation analysis paradigm to assess how effectively a test suite detects interoperability faults. This paradigm is supported by an approach and a toolset detailed below.
- A set of mutation operators that simulate interoperability problems in CPSs. Our mutation operators have been designed and validated with a group of space software experts that includes the heads of the software and testing departments of our industry partners, which are GomSpace, LuxSpace, and ESA.
- A methodology to define tabular fault models that characterize faults possibly affecting the data exchanged by CPS components at runtime.
- A set of adequacy metrics that enable engineers to distinguish between the kinds of problems potentially affecting their test suites, to better guide test suite improvements: data types not being exercised, input partitions not being covered, inadequate oracles, and application states not reached.
- *DaMAT*, a technique that implements data-driven mutation analysis by mutating the data exchanged through data buffers in the SUT. *DaMAT* simulates interoperability problems by applying a chosen set of mutation operators configured in a tabular fault model. It reports mutation analysis results using the newly introduced adequacy metrics.
- An empirical evaluation conducted with an industrial benchmark including space software currently on orbit.
- A replicability package [20] and the source code of our toolset [21] with a tutorial [22].

This paper proceeds as follows. Section 2 describes background and related work. Section 3 describes the general principles of data-driven mutation analysis and our data-driven mutation technique, *DaMAT*. Section 4 reports on the design and results of our empirical evaluation. Section 5 concludes the paper.

2 Background and Related Work

Data-driven mutation analysis evaluates the effectiveness of a test suite in detecting **interoperability faults**. The CPS literature reports on four different interoperability types [6]: technical (which concerns communication protocols and infrastructure), syntactic (which concerns data format), semantic (which concerns the exchanged information, that is, errors in the processing of exchanged data), and cross-domain interoperability (which concerns interaction through business process languages such as BPEL [23]). Technical and syntactic interoperability are provided by off-the-shelf hardware and libraries (not tested by CPS developers) while cross-domain interoperability concerns systems integrated in online services (e.g., energy plants) but is out of scope for the type of CPSs we target in this work, which are safety-critical CPSs like flight systems, robots, and automotive systems. In this paper, we thus focus on *semantic interoperability* faults, that is, faults that affect CPS components integration and are triggered (i.e., lead to failures) in the presence of specific subsets of the data that might be exchanged by CPS components. We thus aim to ensure that a test suite fails when the data exchanged by CPS components is not the one specified by test cases (e.g., through simulator configurations). Related work includes mutation analysis [3], [4] and fault injection [24] techniques.

Mutation analysis concerns the automated generation of faulty software versions (i.e., mutants) through automated procedures called mutation operators [3], [4]. The effectiveness of a test suite is measured by computing the mutation score, which is the percentage of mutants leading to failures when exercised by the test suite.

Mutation operators introduce syntactical changes into the code of the SUT. The *sufficient set of operators* is implemented by most mutation analysis toolsets [25], [26], [27], [28]. Unfortunately, these operators simulate faults concerning the implementation of algorithms (e.g., a wrong logical connector), which is usually tested in unit test suites that, by definition, do not exercise the communication among components, our target in this paper. Also, as stated in the Introduction, such operators cannot be used to generate faulty data with simulated or off-the-shelf components. *Higher-order* mutation analysis [29], which simply combines multiple operators, has the same limitations.

Recent work has introduced mutation operators for cyber-physical systems [30]; they simulate low-level faults on hardware devices (e.g., by modifying the pin identifier of a general-purpose input/output integrated circuit) thus not addressing interoperability issues, our main focus in this paper. Further, they mutate the software implementation, thus presenting the same limitations as the approaches above.

Components integration is targeted by interface [31], integration [32], contract-based [33], and system-level mutation analysis [34]. The former three assess the quality of integration test suites by introducing changes that concern function invocations (e.g., switch function arguments) and inter-procedural data-flow (e.g., alter assignments to variables returned to other components); they can simulate integration faults in units integrated with API invocations but not interoperability problems concerning larger components communicating through channels (e.g., network). System-level mutation relies on operators for GUI components, which are out of our scope, and configuration files, by applying simple mutations, such as deleting a line of text, and are unlikely to lead to interoperability problems.

Fault injection techniques simulate the effect of faults by altering, at runtime, the data processed by the SUT [24]. Faults are introduced according to a fault model that describes the type of fault to inject, the timing of the injection, and the part of the system targeted by the injection. Different from data-driven mutation analysis, fault injection techniques aim to stress the robustness of the software, not assess the quality of its test suites.

Faults affecting components' communication, CPU, or memory can be simulated by performing bit flips [35], [36], [37], [38]. Communication faults are simulated also by duplicating or deleting packets, altering their sequence, or introducing incorrect identifiers, checksums, or counters [39], [40]. Faults affecting signals can be simulated by shifting the signal or increasing the number of signal segments [41]. The largest set of faults affecting data exchanged through files or byte streams is simulated by Peach [42], which includes also protocol-specific fault injection procedures such as replacing host names with randomly generated ones. In general, although existing techniques may simulate a large set of faults they do not cover all the CPS interoperability faults (see Section 3.2).

Approaches performing fault injections other than bit flips require a model of the data to modify. The modelling formalisms adopted for this purpose are grammars [43], [44], [45], [46], UML class diagrams [39], [40], or block models [42], [47]. Grammars are used to model textual data (e.g., XML), which is seldom exchanged by CPS components because of parsing cost. Block models enable specifying the representation to be used for consecutive blocks of bytes, which makes them applicable to a large set of systems; however, existing block model formalisms rely on the XML format, which is expensive to process and thus not usable with real-time systems [42], [47]. The UML class diagram is a formalism that enables the specification of complex data structures and data dependencies [39], [40]; however, it requires loading the data as UML class diagram instances, which is too expensive for real-time systems.

Bai et al. introduced a set of fault injection operators that alter the data generated by the simulators used in CPS testing [48]. They aim to reduce test input selection costs and do not target mutation analysis; indeed, their operators enable engineers to test exceptional scenarios by reusing simulator configurations for nominal cases. Different from us, they support robustness testing, not test suite assessment; further, their approach does not rely on any data modeling solution and thus requires the re-implementation of each mutation operator for each SUT, an error-prone activity.

To summarize, the modification of the data exchanged by software components enables the simulation of communication and, therefore, semantic interoperability faults. Test suites can thus be assessed by relying on fault injection techniques to mutate data. However, existing fault injection techniques do not target mutation analysis; consequently, we lack methods for the specification of fault models and metrics for the assessment of test suites. Also, a larger set of procedures for the modification of data is needed. Finally, block models can effectively capture the structure of the data to modify but formalisms not relying on XML are needed. Our paper addresses such limitations.

3 Approach

In this section we provide an overview of *DaMAT* (Section 3.1), followed by a description of its fault model (Sec-

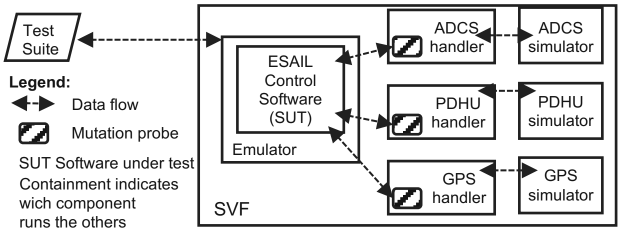

Fig. 1. Data mutation probes integrated into ESAIL.

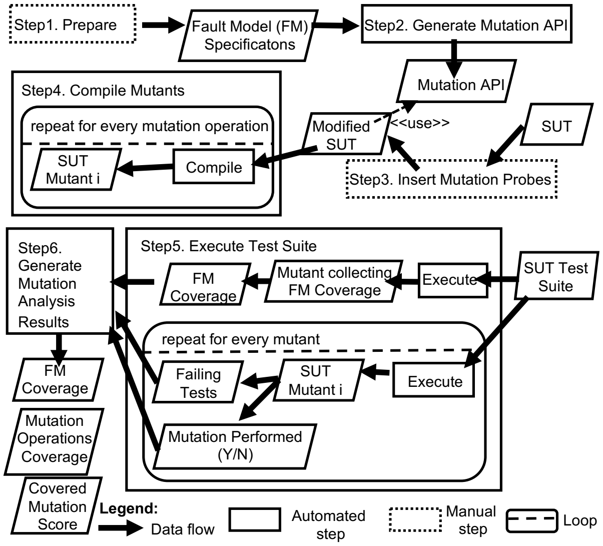

Fig. 2. The *DaMAT* process.

tion 3.2) and each of its steps (Sections 3.3 to 3.7).

3.1 Overview

Data-driven mutation analysis aims to evaluate the effectiveness of a test suite in detecting semantic interoperability faults. It is achieved by modifying (i.e., mutating) the data exchanged by CPS components. It generates *mutated data* that is representative of what might be observed at runtime in the presence of a component that behaves differently than expected in the test case. The mutated data shall lead to different execution behavior than the original data to trigger test failures. For these reasons, data mutation is driven by a fault model specified by the engineers based on domain knowledge.

Although different types of fault models might be envisioned, in this paper we propose a technique (*data-driven mutation analysis with tables*, *DaMAT*), which automates data-driven mutation analysis by relying on a tabular block model, itself tailored to the SUT through predefined mutation operators. To concretely perform data mutation at runtime, *DaMAT* relies on a set of *mutation probes* that shall be integrated by software engineers into the software layer that handles the communication between components. The runtime behaviour of mutation probes (i.e, what data shall be mutated and how) is driven by the fault model. Thus, *DaMAT* can automatically generate the implementation of mutation probes from the provided fault model. Depending on the CPS, probes might be inserted either into the SUT, into the simulator infrastructure, or both. Figure 1 shows the architecture of the ESAIL satellite system (one of the subjects considered in our empirical evaluation) with mutation probes integrated into the SVF[2] functions that handle communication with external components (PDHU, GPS, and ADCS in this case).

DaMAT works in six steps, which are shown in Figure 2. In Step 1, based on the provided methodology and predefined mutation operators, the engineer prepares a fault model specification tailored to the SUT. In Step 2, *DaMAT* generates a mutation API with the functions that modify the data according to the provided fault model. In Step 3, the engineer modifies the SUT by introducing mutation probes (i.e., invocations to the mutation API) into it. In Step 4, *DaMAT* generates and compiles mutants. Since the *DaMAT* mutation operators may generate mutated data by applying multiple mutation procedures, *DaMAT* may generate several mutants, one for each mutation operation (i.e., a mutation procedure configured for a data item, according to our terminology, see Section 3.5). In Step 5, *DaMAT* executes the test suite with all the mutants including a mutant (i.e., the coverage mutant) which does not modify the data but traces the coverage of the fault model (*FM coverage* in Figure 2). In Step 6, *DaMAT* generates mutation analysis results.

The manual steps of *DaMAT* (i.e., Step 1 and Step 3) cannot be automated because they require information that cannot be extracted automatically. The definition of the fault model in Step 1 requires a data model of the software, including valid and invalid input partitions, based on software requirement specification documents, a task that cannot be automated (see Section 3.3). Step 3 requires the identification of program statements handling the communication between components, which also cannot be automated (see Section 3.4).

3.2 Fault Model Structure

The *DaMAT* fault model enables the specification of the format of the data exchanged between components along with the type of faults that may affect such data. In this paper, we refer to the data exchanged by two components as *message*; also, each CPS component may generate or receive different *message types*. For a single CPS, more than one fault model can be specified. For example, in the case of ESAIL we have defined one fault model for every message type that could be exchanged by the three components under test (i.e., ADCS, PDHU, and GPS). In total, for ESAIL, we have 14 fault models, 10 for the communication concerning ADCS (we have 10 different message types), 3 for PDHU, and 1 for GPS. Our methodology for defining the SUT fault models is described in Section 3.3.

The *DaMAT* fault model enables the modelling of data that is exchanged through a specific data structure: the data buffer. This was decided because it is a simple and widely adopted data structure for data exchanges between components in CPS. Also, more complex data structures (e.g., hierarchical ones like trees) are often flattened into

[2] Software Validation Facility [49]; it usually includes one or more simulators, an emulator to run the code compiled for the target hardware, and test harnesses.

TABLE 1
Data-driven mutation operators.

Fault Class	Types	Parameters	Description
Value above threshold (VAT)	I,L,F,D,H	T: threshold Δ: delta, difference with respect to threshold	Replaces the current value with a value above the threshold T for a delta (Δ). It simulates a value that is out of the nominal case and shall trigger a response from the system that shall be verified by the test case (e.g., the system may continue working but an alarm shall be triggered). Not applied if the value is already above the threshold. *Mutation procedure:* $v' = \begin{cases} (T + \Delta) & if\, v \leq T \\ v & otherwise \end{cases}$
Value below threshold (VBT)	I,L,F,D,H	T: threshold Δ: delta, difference with respect to threshold	Replaces the current value with a value below the threshold T for a delta (Δ). It simulates a value that is out of the nominal case and shall trigger a response from the system that shall be verified by the test case (e.g., the system may continue working but an alarm shall be triggered). Not applied if the value is already below the threshold. *Mutation procedure:* $v' = \begin{cases} (T - \Delta) & if\, v \geq T \\ v & otherwise \end{cases}$
Value out of range (VOR)	I,L,F,D,H	MIN: minimum valid value MAX: maximum valid value Δ: delta, difference with respect to minimum/maximum valid value	Replaces the current value with a value out of the range $[MIN; MAX]$. It simulates a value that is out of the nominal range and shall trigger a response from the system that shall be verified by the test case (e.g., the system may continue working but an alarm shall be triggered). Not applied if the value is already out of range. This was inspired by the *ARBC* operator [40]; however, *DaMAT* enables engineers to explicitly specify the delta. *Mutation procedure 1:* $v' = \begin{cases} (MIN - \Delta) & if\, MIN \leq v \leq MAX \\ v & otherwise \end{cases}$ *Mutation procedure 2:* $v' = \begin{cases} (MAX + \Delta) & if\, MIN \leq v \leq MAX \\ v & otherwise \end{cases}$
Bit flip (BF)	B	MIN: lower bit MAX: higher bit STATE: mutate only if the bit is in the given state (i.e., 0 or 1). VALUE: number of bits to mutate	A number of bits randomly chosen in the positions between MIN and MAX (included) are flipped. If STATE is specified, the mutation is applied only if the bit is in the specified state; the value -1 indicates that any state shall be considered for mutation. Parameter VALUE specifies the number of bits to mutate. This was inspired by the *BitFlipperMutator* operator [42]; however, *DaMAT* introduces the STATE parameter, which is not supported by related work. *Mutation procedure:* the operator flips VALUE randomly selected bits if they are in the specified state.
Invalid numeric value (INV)	I,L,F,D,H	MIN: lower valid value MAX: higher valid value	Replace the current value with a mutated value that is legal (i.e., in the specified range) but different than current value. It simulates the exchange of data that is not consistent with the state of the system. It matches the *ARR* operator [40]. *Mutation procedure:* replace the current value with a different value randomly sampled in the specified range.
Illegal Value (IV)	I,L,F,D,H	VALUE: illegal value that is observed	Replace the current value with a value that is equal to the parameter *VALUE*. It matches the *ValidValuesMutator* operator [42]. *Mutation procedure:* $v' = \begin{cases} VALUE & if\, v \neq VALUE \\ v & otherwise \end{cases}$
Anomalous Signal Amplitude (ASA)	I,L,F,D,H	T: change point Δ: delta, value to add/remove VALUE: value to multiply	The mutated value is derived by amplifying the observed value by a factor *VALUE* and by adding/removing a constant value Δ from it. It is used to either amplify or reduce a signal in a constant manner to simulate unusual signals. The parameter *T* indicates the observed value below which instead of adding we subtract. *Mutation procedure:* $v' = \begin{cases} T + ((T - v) * VALUE) + \Delta & if\ v \geq T \\ T - ((v - T) * VALUE) - \Delta & if\ v < T \end{cases}$
Signal Shift (SS)	I,L,F,D,H	Δ: delta, value by which the signal should be shifted	The mutated value is derived by adding a value Δ to the observed value. It simulates an anomalous shift in the signal. This was inspired by work on signal mutation [41]; however, *DaMAT* also enables engineers to rely on SS to increment (or decrement) counters and identifiers. *Mutation procedure:* $v' = v + \Delta$
Hold Value (HV)	I,L,F,D,H	V: number of times to repeat the same value	This operator keeps repeating an observed value for V times. It emulates a constant signal replacing a signal supposed to vary. *Mutation procedure:* $v' = \begin{cases} previous\ v' & if\ counter \leq V \\ v & otherwise \end{cases}$
Fix value above threshold (FVAT)	I,L,F,D,H	T: threshold Δ: delta, difference with respect to threshold	It is the complement of VAT and implements the same mutation procedure as VBT but we named it differently because it has a different purpose. Indeed, it is used to verify that test cases exercising exceptional cases are verified correctly. In the presence of a value above the threshold, it replaces the current value with a value below the threshold T for a delta Δ. *Mutation procedure:* $v' = \begin{cases} (T - \Delta) & if\, v > T \\ v & otherwise \end{cases}$
Fix value below threshold (FVBT)	I,L,F,D,H	T: threshold Δ: delta, difference with respect to threshold	It is the counterpart of FVAT for the operator VBT. *Mutation procedure:* $v' = \begin{cases} (T + \Delta) & if\, v < T \\ v & otherwise \end{cases}$
Fix value out of range (FVOR)	I,L,F,D,H	MIN: minimum valid value MAX: maximum valid value	It is the complement of VOR and implements the same mutation procedure as INV but we named it differently because it has a different purpose. Indeed, it is used to verify that test cases exercising exceptional cases are verified correctly. *Mutation procedure:* $v' = \begin{cases} v & if\, MIN \leq v \leq MAX \\ random(MIN, MAX) & otherwise \end{cases}$

Legend: I: INT, L: LONG INT, F: FLOAT, D: DOUBLE, B: BIN, H: HEX data buffers in order to be exchanged by different components (e.g., through the network). When the CPS software is implemented in C or C++ (common CPS development languages) data buffers are implemented as arrays. Figure 3 shows three block diagrams representing (part of) the buffer structure used to exchange messages of type InterfaceHouseKeeping and InterfaceStatus in ESAIL.

A data buffer is characterized by a *unit size* that specifies the dimension, in bytes, of the single cell of the underlying array and a *buffer size*, which specifies the total number of units belonging to the buffer. Each data buffer can contain one or more *data items*; the size of data items may vary as they may span over multiple units. Also, each data item is interpreted by the CPS software according to a specific *representation* (e.g., integer, double, etc.). In ESAIL, the unit size is one byte and the data items may span over one or two buffer units (see Figure 3).

The *DaMAT* fault model enables engineers to specify (1) the *position* of each data item in the buffer, (2) their *span*, and (3) their *representation type*. Our current implementation supports six data representation types: int, long int, float, double, bin (i.e., data that should be treated in its binary form), hex (i.e., data that should be treated as hexadecimal). Further, for each data item, *DaMAT* enables engineers to specify one or more data faults using the mutation operator identifiers. For each operator, the engineer shall provide values for the required configuration parameters (e.g., the nominal range for a numeric data item).

Table 1 provides the list of mutation operators included in *DaMAT* along with their description and a description of their configuration parameters. The *DaMAT* mutation operators generate *mutated data item instances* through one or more *mutation procedures*, which are the functions that generate a mutated data item instance given a correct data item instance observed at runtime. For example, the *VAT* operator includes only one mutation procedure (i.e., setting the current value above the threshold) while the *VOR* operator includes two mutation procedures, which are (1) replacing the current value with a value above the specified valid range and (2) replacing the current value with a value below the valid range. The operators VOR, BF, INV, and SS have been inspired by related work [40], [41], [42]; instead, the operators VAT, VBT, FVAT, FVBT, FVOR, IV, ASA, and HV are a contribution of this paper and were conceptualised from discussions with engineers having leading roles in ESA, GomSpace, and LuxSpace, our partners in this research[3].

Most of our operators generate mutated data values that are deterministic. Such design choice aims to (1) maximize the likelihood of reproducing test outcomes, which is necessary for debugging, and (2) maximize the likelihood of altering software behaviour and thus causing a test failure, which is necessary to avoid false alarms. The operators generating mutated data values above or below thresholds (i.e., VAT, VBT, VOR, FVAT, FVBT) produce deterministic values differing from the threshold by a given delta. Such delta is selected by the engineers to produce deterministic mutated data values that should trigger a change in software behaviour (i.e., different output) and, therefore, be detected by the test suite. The same holds for operators that alter data item instances that are supposed to follow constant or periodic functions (i.e., ASA, HV, SS). Only three operators (i.e., BF, INV, and FVOR) generate random mutated data values. Operators that produce random data values belonging to a valid range (i.e., INV and FVOR) are used when any data value should trigger the desired change in behaviour while a predefined data value may lead to unexpected results; for example, when IP addresses are dynamically assigned, to cause a message loss, replacing a destination IP address with any another IP in the range is preferable to using a predefined IP value which may, for example, accidentally match the destination IP. Finally, the bit flip (BF) operator simulates the generation of noise, which, by definition, is random (i.e., BF mutates a subset of the bits selected by the engineers); however, to eliminate non-determinism, engineers can configure BF to mutate all the bits in the specified set (i.e., by setting VALUE equal to $MAX - MIN$). With non-deterministic operators, engineers can still debug executions by relying on execution traces with the mutated data values generated. Finally, deterministic operators may prevent the identification of unforeseen critical inputs not correctly processed by the SUT; however, such an objective goes beyond the purpose of *DaMAT*, which aims to assess test suites, not to perform robustness testing. Extensions of *DaMAT* to achieve this goal will be the object of future work.

Our partners confirmed the appropriateness of every operator in Table 1 to simulate plausible and critical interoperability faults in CPS software; our partners also confirmed that they could not recall other types of data modifications leading to interoperability problems. Although other data representation types (e.g., null terminated strings) and operators (e.g., replacement of a random char in a string) might be envisioned, in this paper, we focus on operators that are critical in the CPS context, based on our discussions with domain experts. For example, CPS components are unlikely to exchange strings.

Our mutation operators can be applied to other data structures than buffers as they apply to data values. In future work, additional mutation operators can be defined to address specific aspects of other data structures, e.g., changes to structures of trees or lists.

3.3 Fault Modelling Methodology (Step 1)

The fault model shall enable the specification of all possible interoperability problems in the SUT while minimizing equivalent and redundant mutants. Equivalent mutants have the same observable output as the original SUT. Instead, redundant mutants have the same observable output as other mutants. We use the term *observable output* to refer to any output that can be verified by the test suite. The equivalent or redundant nature of a mutant depends on the equivalence relation for observable outputs (i.e., how to determine if two outputs are the same). In a testing context, such equivalence relation depends on the type of testing being performed. For example, system test cases, different than unit test cases, are unlikely to verify the values of

[3] Our discussions involved the Head of Software for ESAIL, the Mission Lead for GomSpace Luxembourg, the Head of the Flight Software Systems Section at ESA ESTEC, and several software engineers working with GomSpace, LuxSpace, and ESA.

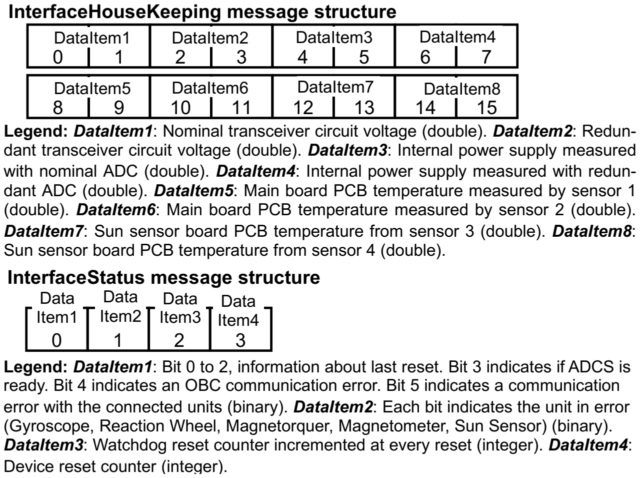

InterfaceHouseKeeping message structure

Legend: ***DataItem1***: Nominal transceiver circuit voltage (double). ***DataItem2***: Redundant transceiver circuit voltage (double). ***DataItem3***: Internal power supply measured with nominal ADC (double). ***DataItem4***: Internal power supply measured with redundant ADC (double). ***DataItem5***: Main board PCB temperature measured by sensor 1 (double). ***DataItem6***: Main board PCB temperature measured by sensor 2 (double). ***DataItem7***: Sun sensor board PCB temperature from sensor 3 (double). ***DataItem8***: Sun sensor board PCB temperature from sensor 4 (double).

InterfaceStatus message structure

Legend: ***DataItem1***: Bit 0 to 2, information about last reset. Bit 3 indicates if ADCS is ready. Bit 4 indicates an OBC communication error. Bit 5 indicates a communication error with the connected units (binary). ***DataItem2***: Each bit indicates the unit in error (Gyroscope, Reaction Wheel, Magnetorquer, Magnetometer, Sun Sensor) (binary). ***DataItem3***: Watchdog reset counter incremented at every reset (integer). ***DataItem4***: Device reset counter (integer).

Fig. 3. Structure of data buffers in ESAIL.

TABLE 2
DaMAT fault modelling methodology

Data nature	Representation type	Dependencies	# of input partitions	Operators	Comments
numerical	I, L, F, D	stateless or stateful	2	[VAT,FVAT] or [VBT,FVBT]	Nominal below T Nominal above T
			3 or more	[VOR,FVOR]	
		stateful		INV	For valid range
				[VOR,FVOR]	For out of range
		signal		ASA, SS, HV	
categorical	I, H	N/A	N/A	IV	
	B	N/A	N/A	BF	
ordinal	I, H	N/A	N/A	ASA	
other	B	N/A	N/A	BF	

Legend: N/A not applicable. I: INT, L: LONG INT, F: FLOAT, D: DOUBLE, B: BIN, H: HEX. [] complementary pair of operators.

all the state variables of the system and thus mutants that are nonequivalent for unit test suites might be considered equivalent for system test suites. For example, in satellite systems, the correctness of the GPS triangulation algorithm output is verified by unit test cases; system test cases, instead, verify if the software takes appropriate actions when the satellite is out of orbit. Consequently, slight changes in the coordinates communicated by the GPS component may not lead to any change in the observable output verified by the test suite.

DaMAT assumes that the mutated data is not automatically corrected by software (e.g., through cyclic redundancy check codes), thus leading to different execution behaviors than expected in the test cases. To target data that is automatically corrected (i.e., to assess if automated data correction is accurately verified by the SUT test suite), it is necessary to work with test suites capable of detecting the presence of unexpected corrections (e.g., by verifying correction counters); otherwise, mutations would lead to equivalent test executions.

The *DaMAT* fault modelling methodology includes the following activities:

1) Identifying, within the SUT, the messages to be mutated; each message will be targeted by a different fault model.
2) Identifying, for each message, the data items that constitute the message and their characteristics, such as *the position* of each data item in the buffer, *their span*, and their *representation type* (e.g., integer, double, binary).
3) Selecting a subset of data-driven mutation operators to apply from Table 1; the selection shall be based on a set of guidelines provided below.
4) Configuring each selected mutation operator based on software specifications, according to our guidelines.
5) Writing a specification in tabular form (i.e., a *CSV* file), where each row represents the configuration of a mutation operator for a specific data item, from a specific fault model. All the fault models of the SUT can be specified in the same specification document. An example is provided in Table 3 and commented at the end of this Section.

We provide a set of guidelines for the selection and configuration of mutation operators that are summarized in Table 2. For guidance, we account for the nature of the data (i.e., numerical, categorical, ordinal, or binary) and their representation type. Also, for numerical data, we consider the data dependencies, that is how data values depend on previously observed values; we identified three categories: (1) *stateless* (i.e., there are no dependencies between consecutive values), (2) *stateful*, when values depend on previous ones (e.g., messages sequence identifier), and (3) *signal* when values are determined by a function of independent variables like time. Data dependencies determine the granularity of the mutation (i.e., with data dependencies, small differences shall be noticed, as explained below); for non-numerical data, since data dependencies are not observed, we do not provide mutation operators with different granularities.

For *stateless numerical data*, our guidelines are driven by input space partitioning concepts [9]. Indeed, given equivalence relations among outputs, it is unlikely that every change in *stateless numerical data* will result into nonequivalent mutants; however, we can partition the input domain into regions with equivalent values (partitions). Precisely, we rely on the *interface-based input domain modeling* approach [9]: for each data item we identify a number of input partitions (set of values or value ranges) according to the interface specifications of the interacting components. In our methodology, the number and type of mutation operators selected for stateless numerical data depend on the number of input partitions identified.

With *two input partitions* (e.g., nominal and exceptional data values), engineers can rely either on the pair [VBT,FVBT] or the pair [VAT,FVAT]. An example is provided in Figure 4; it concerns a data item with two input partitions for nominal (i.e., 0 to 604,800) and exceptional values (i.e., above 604,800). The VAT and FVAT operators are configured with the threshold parameter set to 604,800 and delta (Δ) set to 1 (often, we use Δ as the step granularity). At runtime, the mutant integrating the VAT operator will replace any value below 604,800 with 604,801 (i.e., $604,800 + 1$); the mutant with the FVAT operator will replace any value above 604,800 with 604,799 (i.e., $604,800 - 1$).

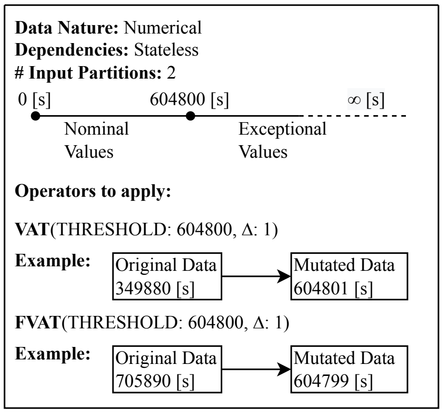

Fig. 4. Example of the *DaMAT* methodology applied to stateless numerical data with two input partitions.

With *three input partitions*, engineers must configure one VOR and one FVOR operator. Figure 5 shows the case of a data item used to transmit voltage information; it presents an input partition for nominal values (i.e., 8.5 [V] to 14.5 [V]) and two input partitions for exceptional values (i.e., below 8.5 and above 14.5). If a different delta (Δ) is considered for the upper and lower bounds, engineers may configure two pairs [VBT,FVBT] and [VAT,FVAT], for the lower and upper bounds, respectively (see items 1 to 4 in Table 3, described below).

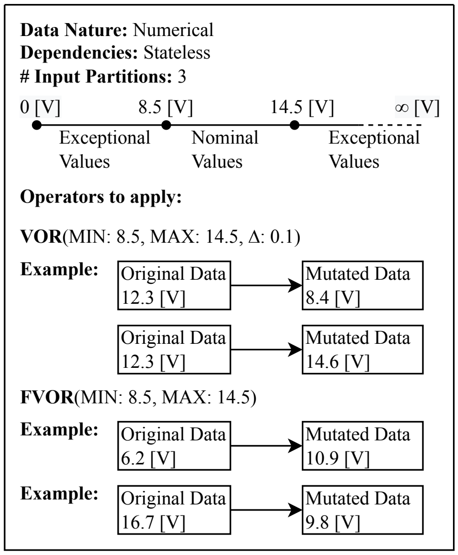

Fig. 5. Example of the *DaMAT* methodology applied to stateless numerical data with three input partitions.

In the presence of *more than three* input partitions, engineers shall configure one [VOR,FVOR] pair for each extra partition above three (e.g., two pairs in the case of five partitions). The parameter Δ is used to determine the partition to which the mutated data belongs.

The example in Figure 6 shows four partitions for the Voltage of a *TCTM* (TeleMetry and TeleCommand) switch in the *ADCS* of ESAIL. Voltage values below 0.55 $[V]$ or above 2.75 $[V]$ are considered non-nominal values. Among nominal values, the voltage is between 0.55 $[V]$ and 1.65 $[V]$ for Position A of the switch and between 1.66 $[V]$ and 2.75 $[V]$ for Position B. The first mutation operator (i.e., VOR with MIN set to 0.55, MAX set to 1.65, and Δ set to 0.01) leads to replacing nominal values for Position A with non-nominal values below 0.55 $[V]$ (i.e., 0.54) and nominal values belonging to Position B (i.e., 1.66). The second mutation operator (i.e., FVOR with MIN set to 0.55, MAX set to 1.65, and Δ set to 0.01) leads to replacing values below 0.55 $[V]$ and values belonging to Position B (i.e., 2.16) with random nominal values for Position A (e.g., 0.98 and 1.17). The third mutation operator (i.e., VOR with MIN set to 1.65, MAX set to 2.75, and Δ set to 0.01) leads to replacing nominal values for Position B with values belonging to Position A (i.e., 1.64) and exceptional values above 2.75 $[V]$ (i.e., 2.76). The fourth mutation operator (i.e., FVOR with MIN set to 1.65 and MAX set to 2.75) leads to replacing values below 1.65 $[V]$ (e.g., belonging to Position A) and non-nominal values above 2.75 $[V]$ with random nominal values for Position B.

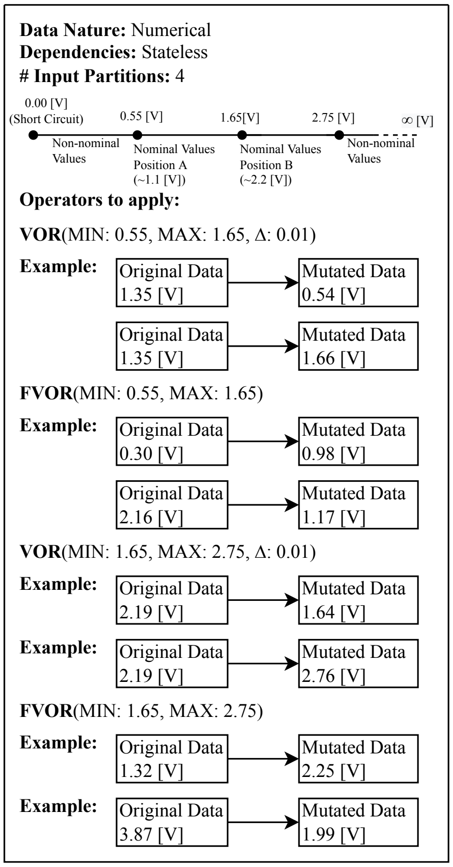

Fig. 6. Example of the *DaMAT* methodology applied to stateless numerical data with four input partitions.

In the presence of *stateful data*, replacement with random values in the valid range (i.e., the INV operator) will lead to nonequivalent mutants (e.g., because it leads to data values that are systematically different than the values expected for the current system state). For example, since the sequence ID of a message depends on the system state (e.g., the number of sent messages), the INV operator shall be used to replace a correct identifier with another one, leading to a failure (i.e.,

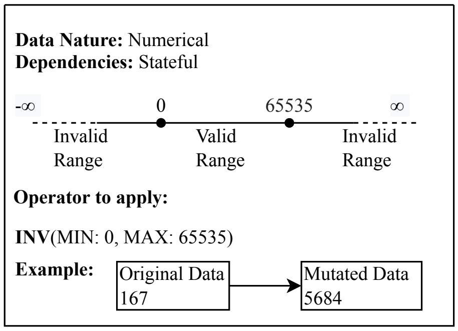

Data Nature: Numerical
Dependencies: Stateful

Operator to apply:

INV(MIN: 0, MAX: 65535)

Example:

Original Data	Mutated Data
167	5684

Fig. 7. Example of the *DaMAT* methodology applied to stateful numerical data.

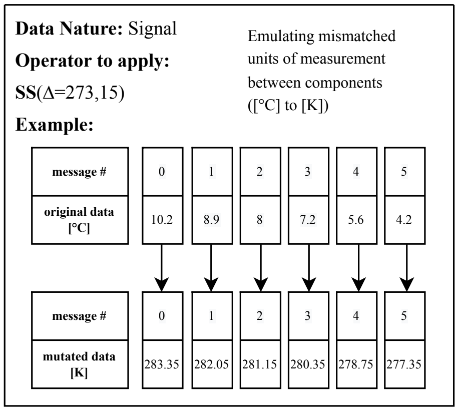

Data Nature: Signal

Operator to apply:

SS(Δ=273,15)

Example:

Emulating mismatched units of measurement between components ([°C] to [K])

message #	0	1	2	3	4	5
original data [°C]	10.2	8.9	8	7.2	5.6	4.2

message #	0	1	2	3	4	5
mutated data [K]	283.35	282.05	281.15	280.35	278.75	277.35

Fig. 8. Example of the *DaMAT* methodology applied to a signal.

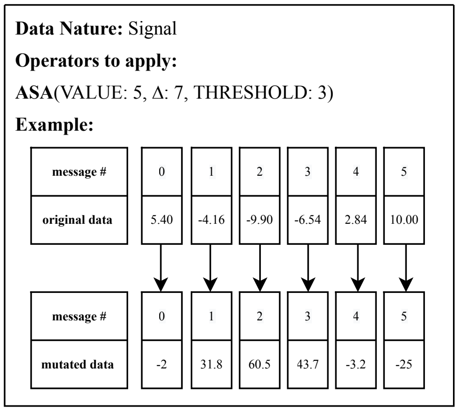

Data Nature: Signal

Operators to apply:

ASA(VALUE: 5, Δ: 7, THRESHOLD: 3)

Example:

message #	0	1	2	3	4	5
original data	5.40	-4.16	-9.90	-6.54	2.84	10.00

message #	0	1	2	3	4	5
mutated data	-2	31.8	60.5	43.7	-3.2	-25

Fig. 9. Example of the ASA operator applied to a signal.

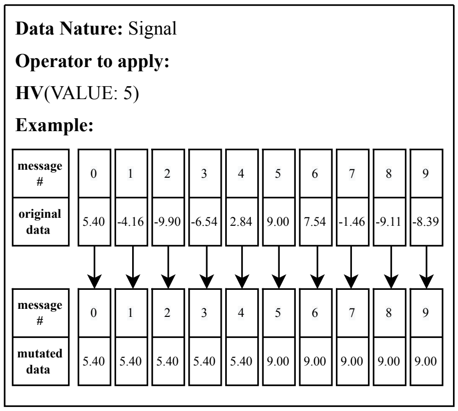

Data Nature: Signal

Operator to apply:

HV(VALUE: 5)

Example:

message #	0	1	2	3	4	5	6	7	8	9
original data	5.40	-4.16	-9.90	-6.54	2.84	9.00	7.54	-1.46	-9.11	-8.39

message #	0	1	2	3	4	5	6	7	8	9
mutated data	5.40	5.40	5.40	5.40	5.40	9.00	9.00	9.00	9.00	9.00

Fig. 10. Example of the HV operator applied to a signal.

the SUT shall notice that the message has an invalid ID and report an output that captures such anomaly, which shall then be detected by the test suite). Figure 7 illustrates this case.

When there is no guarantee that applying the INV operator makes the SUT behave differently (e.g., because the SUT extracts the max value in a sequence), the valid data range might be partitioned as for stateless data. However, to avoid redundant mutants, engineers should rely either on the INV operator or the partitioning of the valid data range. The effect of data outside the valid data range should instead be verified by means of the [VOR, FVOR] pair.

For *signal values*, depending on the shape of the expected signal, engineers should configure one operator among ASA, SS, and HV. The configuration of more than one of these operators may lead to redundant mutants (e.g., because each of them triggers the same warning in the SUT). In the example of Figure 8, the SS operator is used to simulate a mismatch in measurement units; indeed we add 273.15 to report Kelvin degrees instead of Celsius. Figures 9 and 10 exemplify the cases for the ASA and HV operators, respectively. In Figure 9, values above (below) 3 are incremented (decremented) according to the formula in Table 1. Figure 10 shows that the same value is repeated five times, then the newly observed value (i.e., 9.00) is repeated another five times, and so on.

With *categorical data* represented using *integers and hexadecimals*, engineers must configure one IV operator for each possible value; indeed, a change in the observed category shall trigger a different behaviour in the SUT. Figure 11 provides an example where the IV operator is used to alter the command to be executed by the component receiving the message; we configure three IV operators because the system under test includes three commands.

With categorical data in *binary form*, each bit indicates a specific class (e.g., the unit in error for the DataItem2 in the IFStatus message of Figure 3). To verify that the test suite can detect any possible category change, engineers must configure two BF operators for every bit (both MIN and MAX must coincide with the bit position), one operator must flip a bit when it is set (i.e., $STATE = 1$), and the other one when it is unset (i.e., $STATE = 0$); an example appears in Figure 12.

For *ordinal data*, which is represented by means of either integers or hexadecimals, we suggest to apply the ASA operator with *T* being set to the middle point of the ordinal scale and *VALUE* set to the step distance between consecutive data (usually 1) as exemplified in Figure 13.

For data in *binary form* (e.g., pictures), engineers must configure a BF operator to flip a number of bits that is sufficient to alter the semantics of the data (e.g., introduce sufficient noise in images). Figure 14 shows a BF operator that, at runtime, flips four randomly selected bits.

Table 3 provides a specification in tabular form (i.e, the format processed by *DaMAT*) of two fault models configured for the IfKH (i.e., Interface House Keeping) and

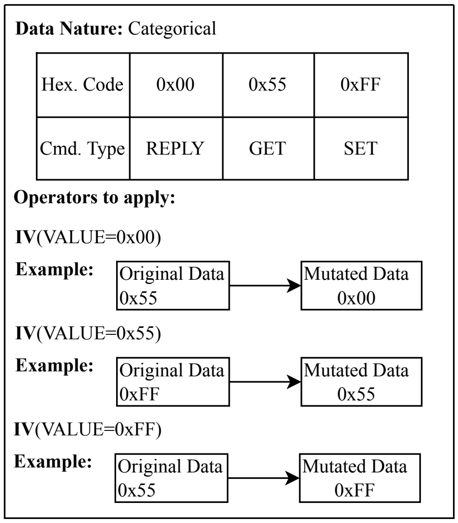

Fig. 11. Example of the *DaMAT* methodology applied to categorical data in hexadecimal format.

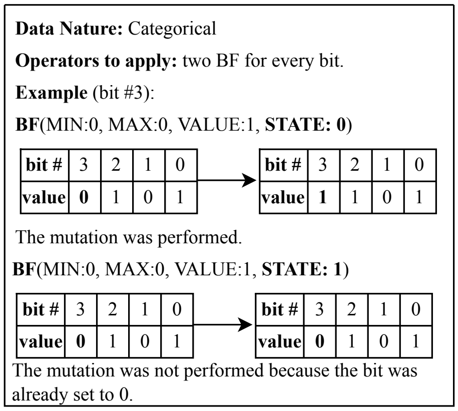
Data Nature: Categorical
Operators to apply: two BF for every bit.
Example (bit #3):
BF(MIN:0, MAX:0, VALUE:1, STATE: 0)
bit # 3 2 1 0
value 0 1 0 1
bit # 3 2 1 0
value 1 1 0 1
The mutation was performed.
BF(MIN:0, MAX:0, VALUE:1, STATE: 1)
bit # 3 2 1 0
value 0 1 0 1
bit # 3 2 1 0
value 0 1 0 1
The mutation was not performed because the bit was already set to 0.

Fig. 12. Example of the *DaMAT* methodology applied to categorical data in binary format.

IfStatus (i.e, Interface Status) messages. In the fault models, each row captures the configuration of a mutation operator for a specific data item. For example, row number 5 indicates that *DaMAT* interprets as double the data inside the two buffer units starting at position 10 (units 10 and 11) and applies the VAT operator. Rows 1 and 3 show that, for a same numerical data item (i.e., the one covering units 0 and 1), we can apply both the VAT and VBT operators, using a different delta for each. Rows 2 and 4 show the FVAT and FVBT operators complementing the VAT and VBT operators in rows 1 and 3. They simulate the case in which data for the nominal cases is observed instead of data for exceptional cases, as visible in Table 1. Rows 8 to 23 show that different bits of a same data item can be targeted by different BF operators. Rows 8 to 13 concern binary categorical data with two categories each, thus we configured two BF each. Rows 14 to 23 concern binary categorical data with five categories; consequently, they present ten BF operators configured for the five categories.

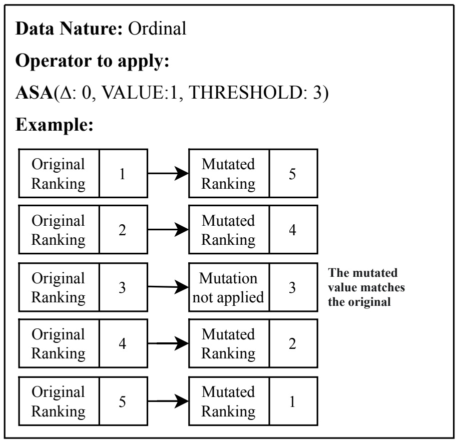

Fig. 13. Example of the *DaMAT* methodology applied to ordinal data.

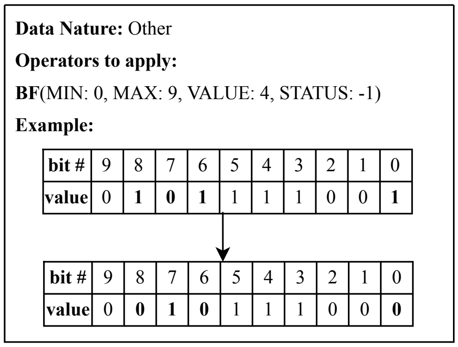
Data Nature: Other
Operators to apply:
BF(MIN: 0, MAX: 9, VALUE: 4, STATUS: -1)
Example:
bit # 9 8 7 6 5 4 3 2 1 0
value 0 1 0 1 1 1 1 0 0 1
bit # 9 8 7 6 5 4 3 2 1 0
value 0 0 1 0 1 1 1 0 0 0

Fig. 14. Example of the *DaMAT* methodology applied to generic data in binary form.

3.4 Automated Generation of Mutation API (Step 2) and Probe Insertion (Step 3)

DaMAT automatically generates a *mutation API* to perform mutations at runtime. The API implements a set of functions (called *mutate_FM_<name>*) that mutate a data buffer according to the given fault model. These functions select the data item to mutate and the mutation procedure to apply based on the mutant under test (see Section 3.5).

The *DaMAT* mutation API works with C/C++ code;

TABLE 3
Portion of the fault model specification for ESAIL

#	Fault Model	Position	Span	Type	Op	MIN	MAX	T	DELTA	STATE	VALUE
1	IfHK	0	2	DOUBLE	VAT	-	-	33.53	0.01	-	-
2	IfHK	0	2	DOUBLE	FVAT	-	-	33.53	0.01	-	-
3	IfHK	0	2	DOUBLE	VBT	-	-	24	1	-	-
4	IfHK	0	2	DOUBLE	FVBT	-	-	24	1	-	-
5	IfHK	10	2	DOUBLE	VAT	-	-	6	1	-	-
6	IfHK	12	2	DOUBLE	VOR	-20	50	-	1	-	-
7	IfHK	14	2	DOUBLE	VOR	-20	50	-	1	-	-
8	IfStatus	0	1	BIN	BF	3	3	-	-	0	1
9	IfStatus	0	1	BIN	BF	3	3	-	-	1	1
10	IfStatus	0	1	BIN	BF	4	4	-	-	0	1
11	IfStatus	0	1	BIN	BF	4	4	-	-	1	1
12	IfStatus	0	1	BIN	BF	5	5	-	-	0	1
13	IfStatus	0	1	BIN	BF	5	5	-	-	1	1
14	IfStatus	1	1	BIN	BF	0	0	-	-	0	1
15	IfStatus	1	1	BIN	BF	0	0	-	-	1	1
16	IfStatus	1	1	BIN	BF	1	1	-	-	0	1
17	IfStatus	1	1	BIN	BF	1	1	-	-	1	1
18	IfStatus	1	1	BIN	BF	2	2	-	-	0	1
19	IfStatus	1	1	BIN	BF	2	2	-	-	1	1
20	IfStatus	1	1	BIN	BF	3	3	-	-	0	1
21	IfStatus	1	1	BIN	BF	3	3	-	-	1	1
22	IfStatus	1	1	BIN	BF	4	4	-	-	0	1
23	IfStatus	1	1	BIN	BF	4	4	-	-	1	1

Note: a ”-” is used for parameters not required to configure a mutation operator.

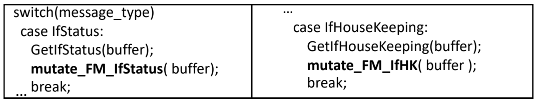

```
switch(message_type)
 case IfStatus:
   GetIfStatus(buffer);
   mutate_FM_IfStatus( buffer);
…  break;

   …
    case IfHouseKeeping:
      GetIfHouseKeeping(buffer);
      mutate_FM_IfHK( buffer );
      break;
```

Fig. 15. Example of *DaMAT* mutation probes (in bold).

however it may be extended to deal with other programming languages. Since it is not possible to automatically determine which data buffer to mutate, *DaMAT* requires engineers to modify the source code of the CPS under test by introducing a mutation probe which consists of an invocation of the *DaMAT* function that mutates the data buffer according to a specific fault model.

Mutation probes should be inserted into the functions that handle components' communication within either the SUT or the simulator. In our experiments, we inserted them into functions that either receive or produce data (see Section 4.1). We expect engineers to insert probes into functions that are easier to modify (e.g., requiring fewer probes, with more lenient time requirements).

We insert probes into the source code of the software under test (or the simulator used to drive testing) instead of the transmission layer (e.g., OS network functions) because data mutation is driven by the data semantics, which varies according to message type (e.g., the thresholds for VBT). Since the transmission layer does not distinguish low-level messages based on their content, we insert probes into code locations (i.e., SUT functions) where the message type is known. Probes inserted into the transmission layer would need to process the message content to match it to message types (e.g., through IDs); such a solution would introduce runtime overhead and further modeling effort.

Note that the effort required to insert probes is limited; indeed, the exchange of data between components is usually managed in a single location (e.g, the function that serializes the data buffer on the network) and thus it is usually sufficient to introduce one function call for each message type to mutate.

Figure 15 shows how the implementation of ESAIL has been modified to add the mutation probes. The SVF function was modified to handle the message requests sent to the ADCS by inserting one mutation probe for each message type to mutate, e.g., IfStatus and IfHouseKeeping in Figure 15. Function *mutate_FM_IfStatus* is part of the generated mutation API; it loads the fault model *IfStatus* into memory (our API relies on a tree data structure) and then invokes the function *mutate*. The function *mutate* performs data-driven mutation according to the provided fault model; the implementation of *mutate* is part of the *DaMAT* toolset.

The behavior of function *mutate* depends on the value of a unique identifier (i.e., the *MutantID*) associated at compile time to the mutant; the *Mutant ID* univocally identifies the performed mutation operation (each mutant executes one mutation operation, see Section 3.5). At a high level, *mutate* performs four activities. First, it checks if the mutation should be performed (i.e., if the data buffer is targeted by the mutation operation identified with the *Mutant ID*). Second, it casts the data item instance targeted by the mutant to a support variable of the type specified in the fault model. Third, it mutates the data stored in the support variable; for each mutation operator, we have implemented a distinct set of instructions for each data representation type. Fourth, before terminating, the function *mutate* writes the mutated data back to the data buffer.

3.5 Automated Generation of Mutants (Step 4)

Consistent with code-driven mutation analysis, *DaMAT* generates one mutant for each mutation procedure of the mutation operators configured in the fault model. Each mutant performs exactly one *data mutation operation* (i.e., a data mutation procedure configured for a specific data item). For example, the specification in row 6 of Table 3 makes *DaMAT* generate two mutants: each mutant modifies the value of the data item starting at position 12 but one mutant replaces the current value with the value 51 (i.e., $50 + 1$) while the other replaces the current value with the value -21 (i.e., $-20 - 1$).

The mutant generation is invisible to the end-user who does not need to modify the source code further; indeed, we rely on a C macro to specify, at compile time, which mutation operation must be performed by every mutant. Mutants are generated by compiling the SUT multiple times, once for each mutation operation. At runtime, the mutate function executes only the mutation operation selected for the mutant under test.

3.6 Mutants Execution (Step 5)

As for code-driven mutation analysis, the test suite under analysis is executed iteratively with every data-driven mutant. At runtime, all the data items targeted by a mutant are mutated whenever the mutation preconditions hold (e.g., the STATE of the BF operator); we leave the mutation of a sampled subset of data item instances to future work [50], [51]. To speed up the mutation analysis process, the test suite under analysis is first executed with a special mutant that, instead of mutating data items, keeps trace of the fault models loaded by each test case; in other words, it traces what are the data types covered by each test case. The collected information enables the execution, for every mutant, of the subset of test cases that cover the message type targeted by the mutant, thus speeding up mutation analysis.

The UML sequence diagram in Figure 16 exemplifies the execution of a test case for ESAIL with a mutant generated by the operator configured in the first row of Table 3, targeting the *IfHK* message. During the test case execution, there are three interactions between the on-board software and the ADCS. The first interaction is a request for *IfHK*, which makes the ADCS generate, in a buffer, an *IfHk* message with a voltage value below 33.53 (DataItem 1 in Figure 3). During the execution, the buffer is mutated by the *DaMAT* probe through a call to the *Mutation API*. The second interaction is a request for an *IfStatus* message; in this case the message is not mutated because, in the current mutant, the API is not configured for targeting *IfStatus* messages. The third interaction is a request for *IfHK*, which is the mutant target. However, in this case, the data is not mutated because the voltage value is above 33.53 (i.e., the threshold for the VAT operator).

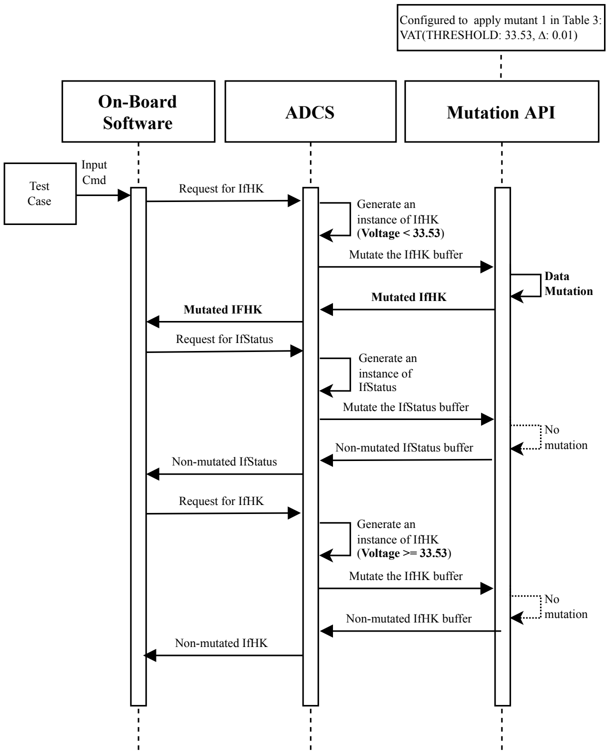

Fig. 16. UML sequence diagram for the execution of an ESAIL test case.

3.7 Mutation Analysis Results (Step 6)

Inspired by work on abstract mutation analysis [52], we have defined three metrics to evaluate test suites with data-driven mutation analysis: fault model coverage, mutation operation coverage, and covered mutation score. These metrics measure the frequency of the following scenarios: (case 1) the message type targeted by a mutant is never exercised, (case 2) the message type is covered by the test suite but it is not possible to perform some of the mutation operations (e.g., because the test suite does not exercise out-of-range cases), (case 3) the mutation is performed but the test suite does not fail.

Fault model coverage (FMC) is the percentage of fault models covered by the test suite. Since we define a fault model for every message type exchanged by two components, it provides information about the extent to which the message types actually exchanged by the SUT are exercised and verified by the test suites. Since different component functionalities often require different message types, low fault model coverage may indicate that only a small portion of the integrated functionalities have been tested.

Mutation operation coverage (MOC) is the percentage of data items that have been mutated at least once, considering only those that belong to the data buffers covered by the test suite. It provides information about the input partitions covered for each data item; for example, the FVOR operator leads to two mutation operations, which are applied only if the observed value is outside range. Otherwise the two mutation operations will not be covered, thus enabling the engineer to identify such shortcoming in the test suite.

The *covered mutation score (CMS)* is the percentage of mutants killed by the test suite (i.e., leading to at least one test case failure) among the mutants that target a fault model and for which at least one mutation operation was successfully exercised (i.e., covered). It provides information about the quality of test oracles; indeed, a mutant that performs a mutation operation and is not killed (i.e., is *live*) indicates that the test suite cannot detect the effect of the mutation (e.g., the presence of warnings in logs). Also, a low CMS may indicate missing test input sequences. Indeed, live mutants may be due to either software faults (e.g., the SUT does not provide the correct output for the mutated data item instance) or the software not being in the required state (e.g., input partitions for data items are covered when the software is paused); in such cases, with appropriate input sequences, the test suite would have discovered the fault or brought the SUT into the required state. Both poor oracles and lack of inputs indicate flaws in the test case definition process (e.g., the stateful nature of the software was ignored).

Different from the ones produced by code-driven mutation analysis, these three metrics enable engineers to distinguish among possible test suite shortcomings. A test suite shortcoming means that the test suite is not exercising a specific input or message or that oracles are not being (fully) checked. A test suite shortcoming prevents the detection of observable failures caused by data—exchanged between SUT components—not being equivalent to the data assumed by test cases. The categories of test suite shortcomings include (1) untested message types, (2) uncovered input partitions, (3) poor oracle quality, and (4) lack of test inputs. *Fault model coverage* spots untested message types. *Mutation operation coverage* reports uncovered partitions and is computed after excluding mutants belonging to untested message types. The *covered mutation score* is computed after excluding mutants that implement a mutation operation that was not applied; for this reason, a low CMS indicates poor oracle quality and lack of test inputs bringing the system into a specific state.

4 Empirical Evaluation

We address the following research questions:

RQ1. What are the types of test suite shortcomings identified by DaMAT*?* We aim to assess the effectiveness of *DaMAT* in identifying various test suite shortcomings, as described in Section 3.7. In other words, we want to know if mutation analysis based on *DaMAT* can provide clear guidance in terms of what to improve in a test suite.

RQ2. What is the impact of equivalent and redundant data-driven mutants on the mutation analysis process? In general, mutation analysis may lead to the generation of equivalent and redundant mutants. In the specific context of *DaMAT*, we analyze the extent of their impact on mutation scores.

RQ3. Is data-driven mutation feasible? To assess its feasibility in practice, we evaluate the cost of setting up data-driven mutation analysis (i.e., defining fault models and instrumenting the CPS with probes), the duration of the mutation analysis process, and the runtime overhead introduced during test case execution.

4.1 Subjects of the study

To assess our research questions, we considered CPS components used in cubesat constellations and in *ESAIL*, which is a micro-satellite [19] launched into space on September 2020 [53]. More precisely, we consider *LIBP*, which is a client-server component to manage configuration parameters in cubesats. Also, we examine three *ESAIL* software subsystems (1) the Attitude Determination And Control System (*ESAIL-ADCS*), the Global Positioning System (*ESAIL-GPS*), and the Payload Data Handling Unit (*ESAIL-PDHU*). These are representative examples of CPS control and utility software, as well as sensor and actuator drivers.

We rely on *DaMAT* to evaluate the *LIBP* integration test suite by mutating the data exchanged between the client and server components of *LIBP*. Similarly, *DaMAT* is used to evaluate how well the *ESAIL* test suite covers interoperability problems affecting the integration between the control software of *ESAIL* (hereafter, CSW) and the *ESAIL-ADCS*, *ESAIL-PDHU*, and *ESAIL-GPS* components. We thus mutate the data exchanged between *ESAIL* CSW and these three components. Since each of these sub-systems have a different purpose (i.e., their data is processed by distinct CSW functions and affect distinct *ESAIL* features) we treat them as distinct case study subjects although they are tested using the same test suite. We focus on the *ESAIL* test suite that makes use of an SVF to simulate the *ESAIL-ADCS*, *ESAIL-PDHU*, and *ESAIL-GPS* components. The main reason is that these three components can only be executed on the target hardware and thus most of the scenarios involving them are tested in a simulated environment first. We do not mutate messages or data items that are tested only with HIL.

In the case of *LIBP*, we inject mutation probes into the *LIBP* server to mutate both received and generated messages. For *ESAIL*, we insert mutation probes into the SVF that mutate the messages it generates; we avoid mutating the messages received by the SVF because such mutations may lead to input data it does not support. ESAIL features 74 kLoC and its SVF 65 kLoC. The ESAIL test suite includes 384 test cases, takes approximately 10 hours to execute, and relies on three simulated *ESAIL-SVF* sub-systems (i.e., *ESAIL-ADCS*, *ESAIL-GPS*, and *ESAIL-PDHU*). Instead, *LIBP* contains 3 kLoC and is tested through an integration test suite which is composed by 170 test cases. The *LIBP* integration test suite takes approximately 1 minute to execute. We have executed all the test suites with every mutant without imposing time limits, except for test timeouts: a test case is considered to be failing if its execution takes more than three times the time required with the original SUT. By considering both a quick integration test suite and an extensive system test suite, we aim to cover the diversity of scenarios in which our approach can be applied.

4.2 Experimental Setup

With the support of our industry partners, we relied on the systems' specification documents to define the fault models for each subject.

Table 4 provides information about the fault models. The fault models (FMs) for the *ESAIL-ADCS* include multiple configurations (*Configured operators*) of eight mutation operators: BF, VAT, VBT, VOR, IV, FVOR, FVBT, and FVAT. The *ESAIL-PDHU* fault models include four operators: BF, IV, VAT and FVAT. Even though the *ESAIL-GPS* fault model concerns only one data type, it makes use of six operators: ASA, HV, IV, SS, VAT, and FVAT. For *LIBP*, we relied on the operators BF, HV, IV, SS, VAT, and FVAT. They have led to 172 mutation operations for *ESAIL-ADCS*, 23 for *ESAIL-GPS*, 29 for *ESAIL-PDHU*, and 44 for *LIBP*; the number of configured operators and mutation operations match except when we rely on VOR and FVOR.

TABLE 4
Fault models and mutation operators.

Subject	Fault Models	Configured Operators	Mutation Operations
ESAIL-ADCS	10	142	172
ESAIL-GPS	1	23	23
ESAIL-PDHU	3	29	29
LIBP	6	44	44

Except for INV, all the mutation operators provided by *DaMAT* have been used in at least one fault model, which shows their usefulness. The INV operator replaces the observed value with another one in the valid domain; to demonstrate its usefulness, we used it to mutate the destination port and address of an opensource network library produced by GomSpace Luxembourg and tested it in our tutorial on *DaMAT* [22], which we did not include in our experiments.

We performed our experiments using an HPC cluster with Intel Xeon E5-2680 v4 (2.4 GHz) nodes.

4.3 RQ1 - Approach effectiveness

We analyzed the extent to which *DaMAT* helps identify limitations in test suites. For each subject, we inspected uncovered fault models, uncovered mutation operations, and live mutants. We then analyzed how they could potentially be explained by the types of shortcomings introduced in Section 3.7: untested message types (UMT), uncovered input partitions (UIP), poor oracle quality (POQ), and lack of test inputs (LTI). To achieve the above, we proceeded as follows. For each uncovered fault model, we discussed with developers if the functionality triggering the exchange of the targeted message was tested by the test suite. For uncovered mutation operations, we discussed with engineers if they match an uncovered input partition. For live mutants, we determined if they could be killed by improving test oracles (see how equivalent mutants are detected for RQ2).

To address RQ1, based on the above analysis, we discuss below how our metrics (i.e., *fault model coverage - FMC*, *mutation operation coverage - MOC*, and *covered mutation score - CMS*) relate to the predefined shortcoming categories (e.g., a low CMS may indicate missing test oracles). Further, to understand how variations in test effectiveness could be explained, we investigate how our metrics relate to the number of functionalities under test (i.e., the number of fault models - FM), the number of mutation operations (MO), and the number of covered mutation operations (CMO), respectively. To get an idea of observable trends, we compute the Spearman's correlation coefficients between them, hereafter denoted ρ_{FM}, ρ_{MO}, ρ_{CMO}.

Results

Table 5 reports the mutation analysis results according to the metrics introduced in Section 3.7. In Table 6, we report how uncovered fault models, uncovered mutation operations,

TABLE 5
Mutation Analysis Results.

Subject	# FMs	FMC	#MOs-CFM	#CMOs	MOC	Killed	Live	CMS
ESAIL-ADCS	10	90.00%	135	100	74.00%	45	55	45.00%
ESAIL-GPS	1	100.00%	23	22	95.65%	21	1	95.45%
ESAIL-PDHU	3	100.00%	29	24	82.76%	24	0	100.00%
LIBP	6	100.00%	44	41	93.20%	37	4	90.24%

CMO=Covered Mutation Operation, MOs-CFM=Mutation Operations in covered FMs.

TABLE 6
Shortcomings of CPSs test suites.

Short-	*ESAIL-ADCS*			*ESAIL-GPS*			*ESAIL-PDHU*			*LIBP*		
coming	UF	UM	LM	UF	UM	LM	UF	UM	LM	UF	UM	LM
UMT	1	-	-	-	-	-	-	-	-	-	-	-
UIP	-	35	-	-	1	-	-	5	-	-	7	-
POQ	-	-	55	-	-	1	-	-	-	-	-	4
LTI	-	-	-	-	-	-	-	-	-	-	-	-
Total	1	35	55	-	1	1	-	5	-	-	7	4

UF=Uncovered Fault model, UM=Uncovered Mutation operation, LM=Live mutant.

and live mutants are distributed with respect to the different shortcomings we noticed on each subject.

Concerning *fault model coverage*, *ESAIL-ADCS* reached a coverage of 90.00%, while *ESAIL-GPS*, *ESAIL-PDHU*, and *LIBP* all achieved 100%. As expected, the much higher number of messages to test for *ESAIL-ADCS* leads to incomplete testing.

ESAIL-ADCS reached 74% *mutation operation coverage*. *ESAIL-GPS*, *ESAIL-PDHU*, and *LIBP* achieved even higher coverage with 95.65%, 82.76%, and 93.20%, respectively. Since ρ_{MO} = -0.8, results suggest that lower mutation operation coverage is more likely when systems are more complex (i.e., there are many mutation operations, whose numbers depend on the number of input partitions).

Regarding *covered mutation scores*, we report 45.00% for *ESAIL-ADCS*, 95.45% for *ESAIL-GPS*, and 100.00% for *ESAIL-PDHU*. These results indicate a varying performance of the *ESAIL-SVF* test suite across sub-systems. *LIBP* obtained a CMS of 90.24%. Given that ρ_{CMO} = -0.8, we conclude that the mutation score tends to be lower for complex systems with a large number of covered mutation operations [4]; indeed, a large number of mutation operations derive from data items and input partitions in the data model that are difficult to exercise when test cases are manually defined.

Table 6 provides the shortcomings identified for all our subjects. Our analysis confirms that (1) uncovered fault models (i.e., low *FMC*) indicate lack of coverage for certain message types (*UMT*) and, in turn, the lack of coverage of a specific functionality (i.e., setting the pulse-width modulation in *ESAIL-ADCS*); (2) uncovered mutation operations (i.e., low *MOC*) highlight the lack of testing of certain input partitions (*UIP*); (3) live mutants (i.e., low *MS*) suggest poor oracle quality (*POQ*). In our case study systems the presence of live mutants was not explained by the lack of test inputs in the original test suite. Moreover, we have not uncovered latent faults, which is unsurprising given that all these systems went through all testing stages, including HIL, and are on orbit. If applied at earlier stages of development, *DaMAT* could have supported the early improvement of test suites, thus likely preventing the discovery of faults at late development stages (e.g., HIL testing). Our future work includes evaluating the benefits of *DaMAT* at earlier stages in industrial projects.

4.4 RQ2 - Equivalent and redundant mutants

As they potentially have significant impact on the applicability of any mutation analysis approach, we assess the impact of equivalent and redundant mutants generated by *DaMAT*.

We determined if a live mutant is nonequivalent by verifying, with the support of our industry partners, if there existed a test case that, after performing the mutation operation, would generate one observable output (e.g., log entry, state variable, or data sent in response to test inputs) that differs from the one generated by the original program. Otherwise a mutant was considered equivalent.

According to related work, two mutants should be considered redundant if they produce the same observable output for every possible input [54]. Since, with large CPSs, it is not possible to automatically determine if such a condition holds (e.g., differential symbolic execution may not scale and is hardly applicable when components communicate through a network), we rely on manual inspection. To make such an analysis feasible, we first need to select a subset of mutant pairs that appear to be redundant (e.g., mutants that produce the same output for every executed test case). However, the size of the CPSs under analysis prevents the collection of all the observable outputs produced by the system. We thus select as seemingly redundant all the pairs of killed mutants that (1) are exercised by the same test cases and (2) present the same failing assertions for every test case. We then manually inspect the test cases to determine if an additional assertion or a different test input might lead to different results for the two mutants, thus determining if the two mutants are actually redundant. Similar to related work, we exclude live mutants from this analysis [55].

Results.

All live mutants (i.e., 55 mutants for *ESAIL-ADCS*, 1 mutant for *ESAIL-GPS*, and 45 mutants for *LIBP*) generate outputs that differ from the original CPS and, therefore, we did not detect any equivalent mutant. Though it needed to be confirmed, such result was expected since our methodology (Section 3.3), if correctly applied, suggests, for every data item, a set of mutation operators that, by construction, should not lead to mutated data that is equivalent to the original data. Live mutants can be killed by introducing oracles that (1) verify additional entries in the log files (39 instances for *ESAIL-ADCS*, 1 instance for *ESAIL-GPS*), (2) verify additional observable state variables (14 instances for *ESAIL-ADCS*, 45 instances for *LIBP*), and (3) verify not only the presence of error messages but also their content (2 instances for *ESAIL-ADCS*).

We did not find redundant mutants either, which was expected since (1) mutations concerning different data items, by definition, are expected to lead to different outputs, (2) the set of operators applied to a same data item, if selected

[4] Please note that the covered mutation score is computed considering only the mutation operations that are exercised (i.e., covered).

according to our methodology, cannot lead to mutated data that is redundant. When analysing pairs of *seemingly redundant mutants*, we identified five reasons why these mutants were actually not redundant: (1) the test case does not distinguish failures across data items (e.g., temperature values collected by different sensors), (2) the test case does not distinguish errors across different messages (e.g., in *ESAIL-ADCS*, the IfHK message reporting a broken sensor or the message sent by a sensor reporting malfunction), (3) the test case does not distinguish between errors in nominal and non-nominal data (e.g., it does not distinguish between VOR and FVOR), (4) the test case does not distinguish between upper and lower bounds (e.g., the mutants for VOR lead to the same assertion failures), and (5) the test case does not distinguish between different error codes (i.e, it simply verifies that an error code is generated). Addressing such shortcomings make test cases more useful for root cause analysis.

4.5 RQ3 - Feasibility

The feasibility of data-driven mutation analysis depends on the required manual effort, which includes defining fault model specifications and injecting probes into the SUT source code. Also, the overhead introduced at runtime by the execution of the mutation operations may introduce delays in real-time systems and consequently cause failures. Finally, feasibility also depends on the overall duration of the mutation analysis process.

To discuss manual effort, we measured, for each subject, (1) the number of rows in the fault model specifications, as they match the number of operators manually identified and configured by an engineer, and (2) the number of lines of code (LoC) added to the source code of our subjects. Since the number of added lines of code depends on the number of fault models per case study, we report the ratio of LoC per fault model. The lines of code added to the source code of our subjects include invocations to function *mutate* (see Section 3.4) and additional utility code such as exit handlers used to clear the fault models loaded into memory.

To address the overhead, we measured the execution time taken by every passing test case when executed with the original software and with any of the mutants generated by the approach. We exclude failing test cases because they may bias the results (e.g., failing assertions may terminate a test case earlier, while test timeout failures are detected when a test case execution takes too long). To account for performance variations due to the varying load of our HPC, we executed every test case three times. For every test case, we then computed the overhead of every mutant as the difference between the average execution time obtained with the mutants and that with the original software. Since different subjects are characterized by different types of messages being exchanged, we discuss the distribution of such overhead among our subjects.

Last, to discuss the overall duration of the mutation analysis process, we report the average time taken to execute the test cases selected by the approach for every mutant, across three runs.

TABLE 7
Manual effort and execution time.

Subject	Configured	Configured	LoC	Execution time	
	Operators	Operators / FM	/ FM	Original	*DaMAT*
ESAIL-ADCS	142	14.20	6.10	8.34 [h]	947.17 [h]
ESAIL-GPS	23	23.00	2.72		42.02 [h]
ESAIL-PDHU	29	9.66	4.33		60.36 [h]
LIBP	44	13.33	7.64	0.015 [h]	0.12 [h]

Results

Manual effort. The left part of Table 7 reports the measures related to manual effort. The number of operators configured per subject varies from 23 to 142, with an average between 9.66 and 23 operators per fault model. In our experiments, on average, it took between five and ten minutes to configure an operator (i.e., to read the paragraph of the software specifications related to a specific data item and to write the configuration values for an operator in the *DaMAT* fault model). Given that the definition of test cases for safety-critical CPS components, such our case study subjects, takes days to complete, our industry partners found the required effort acceptable. The same considerations hold for the number of LoC per fault model, whose average across subjects varied between 2.72 and 7.64[5]. In total, the configuration of *DaMAT* (i.e., configure operators and insert mutation probes) for *ESAIL-ADCS*, *ESAIL-GPS*, *ESAIL-PDHU*, and *LIBP* took around 20, 3, 4, and 6 working hours, respectively.

Overhead. Excluding outliers (i.e., values above 90^{th} percentile), the maximum execution overhead for *ESAIL-ADCS*, *ESAIL-GPS*, *ESAIL-PDHU*, and *LIBP* is 1.47%, 3.16%, 1.7%, and 1.3%, respectively. We did not observe any failure due to violated timing constraints (i.e., constraints defined within SUT test suites that verify both functional and time requirements); this indicates that *DaMAT* does not introduce significant delay. Therefore, overall, the small overhead incurred by the subjects is acceptable and does not prevent the application of *DaMAT* to real-time CPSs.

The right part of Table 7 shows the *DaMAT* analysis time. Although it is much larger than the execution times of test suites for the original SUTs, it is practically feasible. Indeed, in the worst case (i.e., *ESAIL-ADCS*), mutation analysis can be performed in 9 hours with 100 parallel computation nodes; in safety-critical contexts, where development entails large costs, buying computation time on the Cloud is affordable. Code-driven mutation analysis for systems with similar characteristics lasts considerably more [56], [57]. For ESAIL and *LIBP*, for example, mutants sampling[6] makes code-driven mutation analysis feasible with parallel computation nodes, leading to 1800 hours (ESAIL) and 3 hours (*LIBP*) of execution time, which is still significantly higher than the time required by *DaMAT*. Without mutants sampling, mutation analysis of the whole ESAIL software may take up to 589,000 hours. We will evaluate the applicability of mutants sampling to *DaMAT* in future work; how-

[5] The exit handler includes one call for each fault model and thus subjects with less fault models show a lower average. For *LIBP*, the larger number of lines of code is due to the need for recomputing a message checksum after the invocation of function *mutate*.

[6] Mutants sampling consists of selecting a subset of mutants to compute the mutation score [58]. In our previous work [57], we have demonstrated that a fixed-width sequential confidence interval approach guarantees the accurate estimation of the mutation score.

ever, although useful to derive an accurate mutation score, mutants sampling prevents the complete identification of test suite shortcomings. For this reason, we aim to study how mutants and test cases can be prioritized to generate accurate mutation analysis results while executing a subset of the mutants with a subset of the test cases.

4.6 Data Availability

The source code of our toolset, usage examples, and the data collected in our empirical evaluation are available [20]. Our case study software cannot be provided since it is proprietary.

4.7 Discussion

Our results demonstrate that data-driven mutation analysis, as implemented in the *DaMAT* approach, can effectively identify test suite shortcomings. Based on our results, *DaMAT* does not lead to false positives, requires limited manual effort, and does not significantly affect the real-time properties of the software. Also, the overall execution time of the mutation analysis process is much lower than that of code-driven mutation analysis, which is the most similar approach proposed in the literature.

Our industry partners confirmed the correctness and usefulness of our results, as well as the practical feasibility of the approach. Further, the effectiveness of the approach is particularly encouraging given that our case study subjects are safety-critical CPS software systems (i.e., satellite software components) that were—before we started our experiments—already well tested based on several test suite inspection and validation activities, as per ECSS standards [59], [60].

To discuss the ease of adoption of *DaMAT* in industry, we also need to consider the cost for the manual configuration of the approach (i.e., the definition of fault models and the insertion of probes into the SUT source code). Overall, setting up *DaMAT* requires an amount of effort that is comparable to configuring other software testing and mutation analysis tools. For example, to generate unit test cases with KLEE [61], which is a well-known test generation tool based on symbolic execution, it is necessary to implement test templates distinguishing the outputs and the inputs to be treated symbolically for each function under test, which may entail substantial effort for large software systems. Also, to perform code-driven mutation analysis with MASS [57], which is our recent code-driven mutation analysis tool for CPS software, it is necessary to update the configuration and execution scripts for the SUT (e.g., to collect code coverage data, compile with multiple compiler optimization options, determine test timeouts, and track test failures), an activity that requires a few hours to complete. Further, other mutation analysis tools require manual effort to be executed; for example, MART [62] and Mull [63] require the software to be compiled with LLVM [64], which implies modifying compilation scripts and often leads to compilation errors that are hard to solve for large CPS software with many dependencies on third party libraries [57]. Finally, note that our implementation of *DaMAT* enables engineers to introduce probes into the SUT source code using dedicated annotation labels (i.e., specific comment keywords); such annotations do not alter the execution of production code but can be used to automatically perform mutation analysis within a CI/CD pipeline without requiring manual intervention.

Compared to traditional code-driven mutation analysis, data-driven mutation analysis may still require more manual effort and end-user expertise (e.g., knowledge of the SUT) because the former is usually performed automatically by software toolsets without the need for a manually specified fault model. However, data-driven mutation analysis provides more readily interpretable results than code-driven mutation analysis, thanks to our three mutation analysis metrics. For example, the lack of coverage for a FVAT mutation operation applied to a specific data item indicates that the test suite does not exercise a value above the FVAT threshold parameter for that data item (e.g., it does not exercise non-nominal cases); in the case of code-driven mutation analysis, it is often difficult to determine which missing test input prevents the test suite from killing a mutant (e.g, first, engineers need to determine if the mutant is nonequivalent to the original program, which is in itself complicated). Finally, based on our experience, data-driven mutation analysis leads to a lower number of mutants than traditional code-driven mutation analysis. For example, in the case of *LIBP*, *DaMAT* generates 44 mutants while code-driven mutation analysis, based on our previous work [57], leads to 3931 mutants (346, if mutants sampling is applied). The low number of mutants generated by *DaMAT* leads to a low number of mutants to be inspected (e.g., because not detected by the test suite), which greatly reduces the cost of the analysis. For example, in the case of *LIBP*, with *DaMAT*, we inspected only 4 mutants, which sharply contrasts with code-driven mutation analysis that led, in our previous studies, to examine 1179 mutants (or 50 with mutants sampling [57]).

Based on the above, we can conclude that the manual effort required by *DaMAT* is comparable or lower than that required by code-driven mutation analysis, which is already adopted in certain contexts [65], [66], and other well-known software testing tools. Consequently, we believe that data-driven mutation analysis can be realistically adopted in the CPS industry, given that it targets problems that are particularly relevant for CPS software.

Finally, we report the results of an independent evaluation of code-driven and data-driven mutation analysis performed by our industry partners in the context of the ESA project supporting this research. According to our industry partners, the percentage of critical test suite shortcomings detected by code-driven mutations is 38.24%; for data-driven mutation, critical shortcomings add up to 57%. Examples of critical shortcomings include message types not being exercised, input partitions not being covered, or test suites not detecting the presence of unexpected non-nominal data values. These results further highlight the potential of data-driven mutation analysis.

4.8 Threats to validity

External validity. We have selected industrial CPSs of diverse size, tested with different types of test suites. They are developed according to space safety standards and are thus representative of CPS software adhering to safety regulations. Also, ESAIL is larger than any other industrial system considered in the mutation analysis literature to date [56], [63], [67], [68].

Internal validity. To minimize implementation errors, we have extensively tested our toolset; we provide both the test cases and the *DaMAT* source code. Also, the *DaMAT* toolset has been developed in the context of an ESA project which required following ECSS practices for category D software validation and documentation.

Construct validity. The indicators selected for cost estimation (configured operators and LoC) are directly linked to the activities of the end-user and are thus appropriate. We rely on LOC per Fault Model since LOC is an objective though imperfect indicator of effort. Ideally, though far from being straightforward, effort discussions should also be based on data collected from practitioners. However, since the reported LOC/FM values are very small, there is not much room for interpretation in our context. We therefore leave empirical studies with human subjects to future work. Future studies should evaluate not only the effort required to configure *DaMAT* operators but also the effort required to determine what data to mutate; however, based on our experience, determining what to mutate is an easy task in safety-critical CPS because the software specification documents provide a clear description of what are the communicating components, the type of messages being exchanged, and the message structure. Further, future studies should also report on the likelihood of equivalent mutants caused by human errors; for example, operators can be misconfigured (e.g., the misconfigured operator may generate values outside the representable range that will simply not be transmitted over the communication channel) or probes can be inserted in the wrong location (e.g., inserting the mutation probe in a block of code that is executed after the data to be mutated has been processed by the SUT). Such mistakes may lead to false alarms that are noticed only when investigating uncovered fault models, uncovered mutation operators, and live mutants; after fixing such mistakes, the engineers need to re-execute the test suite with the mutants generated by the subset of operators that have been reconfigured. Though user studies are required to confirm this statement, our industry partners found the effort entailed by the approach to be justified by its benefits (e.g., they identified message types not being tested).

5 Conclusion

Assessing the quality of test suites is necessary in the case of large, safety-critical systems, such as most CPSs, where software failures may lead to severe consequences including loss of human lives or environmental damages. This problem is made even more acute due to the fact that such test suites are usually manually constructed.

In this paper, we have introduced data-driven mutation analysis, a new paradigm and associated techniques to assess the effectiveness of a test suite in detecting interoperability faults. Interoperability faults are a major source of failures in CPSs but they are not targeted by existing mutation analysis approaches. Data-driven mutation analysis works by modifying (i.e., mutating) the data exchanged by software components during test execution. Mutations are performed through a set of mutation operators that are configured according to a fault model provided by software engineers, following a proposed methodology. Test suite assessment is based on three metrics: (1) fault model coverage, (2) mutation operation coverage, (3) and covered mutation score. These metrics enable engineers to determine specific weaknesses in test suites (e.g., oracles with low covered mutation score).

We have proposed a partially automated technique, *DaMAT*, that is applicable to a vast range of systems since it mutates the data exchanged through data buffers, by relying on a tabular fault model that can be inexpensively loaded into memory. Based on interactions with engineers, we have defined a set of mutation operators that enable the simulation of common data faults causing interoperability problems in CPSs. We have provided a methodology that supports the definition of a fault model based on the nature, representation type, and dependencies of the data to mutate. In future work, additional mutation operators can be defined to address specific aspects of other data structures, e.g., changes to structures of trees or lists. Furthermore, we will assess mutant sampling strategies to reduce execution time without dampening effectiveness.

We empirically evaluated *DaMAT* by applying it to test suites of commercial CPS components in the space domain that are currently deployed on orbit. Our results show that *DaMAT* can identify a large and diverse set of test suite shortcomings, entails limited modelling and execution costs, and is not affected by redundant and equivalent mutants.

Acknowledgments

This work has been funded by the European Space Agency (ITT-1-9873/FAQAS), the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement No 694277), and NSERC Discovery and Canada Research Chair programs. Authors would like to thank the ESA ESTEC officers, the GomSpace Luxembourg team, LuxSpace software engineers, and Yago Isasi Parache for their valuable support.

References

[1] S. K. Khaitan and J. D. McCalley, "Design techniques and applications of cyberphysical systems: A survey," *IEEE Systems Journal*, vol. 9, no. 2, pp. 350–365, 2015.

[2] H. Roehm, J. Oehlerking, M. Woehrle, and M. Althoff, "Model conformance for cyber-physical systems: A survey," *ACM Trans. Cyber-Phys. Syst.*, vol. 3, no. 3, Aug. 2019. [Online]. Available: https://doi.org/10.1145/3306157

[3] Y. Jia and M. Harman, "An analysis and survey of the development of mutation testing," *IEEE transactions on software engineering*, vol. 37, no. 5, pp. 649–678, 2010.

[4] M. Papadakis, M. Kintis, J. Zhang, Y. Jia, Y. Le Traon, and M. Harman, "Mutation testing advances: an analysis and survey," in *Advances in Computers*. Elsevier, 2019, vol. 112, pp. 275–378.

[5] M. Papadakis, D. Shin, S. Yoo, and D.-H. Bae, "Are mutation scores correlated with real fault detection? A large scale empirical study on the relationship between mutants and real faults," in *2018 IEEE/ACM 40th International Conference on Software Engineering (ICSE)*. IEEE, 2018, pp. 537–548.

[6] O. Givehchi, K. Landsdorf, P. Simoens, and A. W. Colombo, "Interoperability for industrial cyber-physical systems: An approach for legacy systems," *IEEE Transactions on Industrial Informatics*, vol. 13, no. 6, pp. 3370–3378, 2017.

[7] V. Jirkovský, M. Obitko, and V. Mařík, "Understanding data heterogeneity in the context of cyber-physical systems integration," *IEEE Transactions on Industrial Informatics*, vol. 13, no. 2, pp. 660–667, 2017.

[8] S. Abbaspour Asadollah, R. Inam, and H. Hansson, "A survey on testing for cyber physical system," in *Testing Software and Systems*, K. El-Fakih, G. Barlas, and N. Yevtushenko, Eds. Cham: Springer International Publishing, 2015, pp. 194–207.

[9] P. Ammann and J. Offutt, *Introduction to Software Testing*, 1st ed. USA: Cambridge University Press, 2008.

[10] NASA, "Mars Climate Orbiter, Spacecraft lost," https://solarsystem.nasa.gov/missions/mars-climate-orbiter/in-depth/, 1998.

[11] M. E. Delamaro, J. Offutt, and P. Ammann, "Designing deletion mutation operators," in *2014 IEEE Seventh International Conference on Software Testing, Verification and Validation*. IEEE, 2014, pp. 11–20.

[12] M. E. Delamaro, L. Deng, V. H. S. Durelli, N. Li, and J. Offutt, "Experimental evaluation of SDL and one-op mutation for C," in *2014 IEEE Seventh International Conference on Software Testing, Verification and Validation*. IEEE, 2014, pp. 203–212.

[13] N. He, P. Rümmer, and D. Kroening, "Test-case generation for embedded Simulink via formal concept analysis," *Proceedings - Design Automation Conference*, pp. 224–229, 2011.

[14] B. Aichernig, H. Brandl, J. Elisabeth, W. Krenn, R. Schlick, and S. Tiran, "Model-based Mutation Testing for UML," in *ICST 2015*, 2015. [Online]. Available: www.momut.org

[15] X. Devroey, G. Perrouin, M. Papadakis, A. Legay, P.-Y. Schobbens, and P. Heymans, "Featured model-based mutation analysis," in *Proceedings of the 38th International Conference on Software Engineering*, ser. ICSE '16. New York, NY, USA: Association for Computing Machinery, 2016, p. 655–666. [Online]. Available: https://doi.org/10.1145/2884781.2884821

[16] F. Belli, C. J. Budnik, A. Hollmann, T. Tuglular, and W. E. Wong, "Model-based mutation testing—approach and case studies," *Science of Computer Programming*, vol. 120, pp. 25–48, 2016. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0167642316000137

[17] ESA, "European Space Agency," 2020. [Online]. Available: https://www.esa.int

[18] LuxSpace, "ESAIL - Maritime Microsatellite," 2020. [Online]. Available: https://directory.eoportal.org/web/eoportal/satellite-missions/e/esail

[19] F. Davoli, C. Kourogiorgas, M. Marchese, A. Panagopoulos, and F. Patrone, "Small satellites and cubesats: Survey of structures, architectures, and protocols," *International Journal of Satellite Communications and Networking*, vol. 37, no. 4, pp. 343–359, 2019. [Online]. Available: https://onlinelibrary.wiley.com/doi/abs/10.1002/sat.1277

[20] Authors of this paper, "Replicability package," https://doi.org/10.6084/m9.figshare.21276093, 2022.

[21] ——, "DAMAT source code," https://github.com/SNTSVV/DAMAT, 2022.

[22] Enrico Viganò, "DAMAT Tutorial," https://github.com/SNTSVV/DAMAT_Tutorial/, 2022.

[23] OASIS, "OASIS Web Services Business Process Execution Language (WSBPEL) TC," https://www.oasis-open.org/committees/tc_home.php?wg_abbrev=wsbpel, 2007.

[24] R. Natella, D. Cotroneo, and H. S. Madeira, "Assessing dependability with software fault injection: A survey," *ACM Computing Surveys (CSUR)*, vol. 48, no. 3, p. 44, 2016.

[25] A. J. Offutt, A. Lee, G. Rothermel, R. H. Untch, and C. Zapf, "An experimental determination of sufficient mutant operators," *ACM Transactions on Software Engineering and Methodology (TOSEM)*, vol. 5, no. 2, pp. 99–118, 1996.

[26] G. Rothermel, R. H. Untch, and C. Zapf, "An experimental determination of sufficient mutant operators a. jefferson offutt ammei lee george mason university," *ACM Transactions on software engineering methodology*, vol. 5, no. 2, pp. 99–118, 1996.

[27] J. H. Andrews, L. C. Briand, and Y. Labiche, "Is mutation an appropriate tool for testing experiments?" in *Proceedings of the 27th international conference on Software engineering*. ACM, 2005, pp. 402–411.

[28] M. Kintis, M. Papadakis, Y. Jia, N. Malevris, Y. Le Traon, and M. Harman, "Detecting trivial mutant equivalences via compiler optimisations," *IEEE Transactions on Software Engineering*, vol. 44, no. 4, pp. 308–333, 2017.

[29] M. Harman, Y. Jia, and W. B. Langdon, "A manifesto for higher order mutation testing," in *2010 Third International Conference on Software Testing, Verification, and Validation Workshops*. IEEE, 2010, pp. 80–89.

[30] Q. Zhu and A. Zaidman, "Mutation testing for physical computing," in *2018 IEEE International Conference on Software Quality, Reliability and Security (QRS)*. IEEE, 2018, pp. 289–300.

[31] M. E. Delamaro, J. Maidonado, and A. P. Mathur, "Interface mutation: An approach for integration testing," *IEEE transactions on software engineering*, vol. 27, no. 3, pp. 228–247, 2001.

[32] M. Grechanik and G. Devanla, "Mutation integration testing," in *2016 IEEE International Conference on Software Quality, Reliability and Security (QRS)*, 2016, pp. 353–364.

[33] Y. Jiang, S.-S. Hou, J.-H. Shan, L. Zhang, and B. Xie, "Contract-based mutation for testing components," in *21st IEEE International Conference on Software Maintenance (ICSM'05)*, 2005, pp. 483–492.

[34] P. R. Mateo, M. P. Usaola, and J. Offutt, "Mutation at system and functional levels," in *2010 Third International Conference on Software Testing, Verification, and Validation Workshops*. IEEE, 2010, pp. 110–119.

[35] T. K. Tsai, M.-C. Hsueh, H. Zhao, Z. Kalbarczyk, and R. K. Iyer, "Stress-based and path-based fault injection," *IEEE Transactions on Computers*, vol. 48, no. 11, pp. 1183–1201, 1999.

[36] J. H. Barton, E. W. Czeck, Z. Z. Segall, and D. P. Siewiorek, "Fault injection experiments using FIAT," *IEEE Transactions on Computers*, vol. 39, no. 4, pp. 575–582, 1990.

[37] S. Han, K. G. Shin, and H. A. Rosenberg, "Doctor: An integrated software fault injection environment for distributed real-time systems," in *Proceedings of 1995 IEEE International Computer Performance and Dependability Symposium*. IEEE, 1995, pp. 204–213.

[38] S. Dawson, F. Jahanian, T. Mitton, and T.-L. Tung, "Testing of fault-tolerant and real-time distributed systems via protocol fault injection," in *Proceedings of Annual Symposium on Fault Tolerant Computing*. IEEE, 1996, pp. 404–414.

[39] D. Di Nardo, F. Pastore, A. Arcuri, and L. Briand, "Evolutionary robustness testing of data processing systems using models and data mutation (t)," in *2015 30th IEEE/ACM International Conference on Automated Software Engineering (ASE)*. IEEE, 2015, pp. 126–137.

[40] D. Di Nardo, F. Pastore, and L. Briand, "Generating complex and faulty test data through model-based mutation analysis," in *2015 IEEE 8th International Conference on Software Testing, Verification and Validation (ICST)*. IEEE, 2015, pp. 1–10.

[41] R. Matinnejad, S. Nejati, L. C. Briand, and T. Bruckmann, "Test Generation and Test Prioritization for Simulink Models with Dynamic Behavior," *IEEE Transactions on Software Engineering*, vol. 45, no. 9, pp. 919–944, Sep. 2019.

[42] Peach Tech, "Peach Fuzzer." [Online]. Available: https://www.peach.tech

[43] A. K. Ghosh, M. Schmid, and V. Shah, "Testing the robustness of windows nt software," in *Proceedings Ninth International Symposium on Software Reliability Engineering (Cat. No. 98TB100257)*. IEEE, 1998, pp. 231–235.

[44] P. Godefroid, A. Kiezun, and M. Y. Levin, "Grammar-based whitebox fuzzing," in *Proceedings of the 29th ACM SIGPLAN Conference on Programming Language Design and Implementation*, ser. PLDI ?08. New York, NY, USA: Association for Computing Machinery, 2008, p. 206?215.

[45] P. Godefroid, M. Y. Levin, and D. Molnar, "Sage: whitebox fuzzing for security testing," *Communications of the ACM*, vol. 55, no. 3, pp. 40–44, 2012.

[46] E. Bounimova, P. Godefroid, and D. Molnar, "Billions and billions of constraints: Whitebox fuzz testing in production," in *Proceedings of the 2013 International Conference on Software Engineering*. IEEE Press, 2013, pp. 122–131.

[47] V.-T. Pham, M. Böhme, and A. Roychoudhury, "Model-based whitebox fuzzing for program binaries," in *2016 31st IEEE/ACM International Conference on Automated Software Engineering (ASE)*. IEEE, 2016, pp. 543–553.

[48] X. Bai, S. Lee, and Y. Chen, "Mutation-based simulation test data generation for testing complex real-time software," in *40th Annual Simulation Symposium (ANSS'07)*, 2007, pp. 73–80.

[49] Y. Isasi, A. Pinardell, A. Marquez, C. Molon-Noblot, A. Wagner, M. Gales, and M. Brada, "The esail multipurpose simulator," in *Onlin Proceedings of Simulation and EGSE for Space Programmes (SESP2019)*, 2019,

https://atpi.eventsair.com/QuickEventWebsitePortal/sesp-2019/website/Agenda.

[50] L. Zhang, M. Gligoric, D. Marinov, and S. Khurshid, "Operator-based and random mutant selection: Better together," in *Proceedings of the 28th IEEE/ACM International Conference on Automated Software Engineering*. IEEE Press, 2013, pp. 92–102.

[51] R. Gopinath, A. Alipour, I. Ahmed, C. Jensen, and A. Groce, "How hard does mutation analysis have to be, anyway?" in *2015 IEEE 26th International Symposium on Software Reliability Engineering (ISSRE)*. IEEE, 2015, pp. 216–227.

[52] J. Offutt, P. Ammann, and L. Liu, "Mutation testing implements grammar-based testing," *2nd Workshop on Mutation Analysis (Mutation 2006 - ISSRE Workshops 2006), MUTATION'06*, p. 12, 2006.

[53] European Space Agency, "ESA Vega launch," 2020. [Online]. Available: https://www.esa.int/Enabling_Support/Space_Transportation/Vega/

[54] D. Shin, S. Yoo, and D. Bae, "A theoretical and empirical study of diversity-aware mutation adequacy criterion," *IEEE Transactions on Software Engineering*, vol. 44, no. 10, pp. 914–931, Oct 2018.

[55] M. Papadakis, C. Henard, M. Harman, Y. Jia, and Y. Le Traon, "Threats to the validity of mutation-based test assessment," in *Proceedings of the 25th International Symposium on Software Testing and Analysis*. ACM, 2016, pp. 354–365.

[56] R. Ramler, T. Wetzlmaier, and C. Klammer, "An empirical study on the application of mutation testing for a safety-critical industrial software system," *Proceedings of the ACM Symposium on Applied Computing*, vol. Part F128005, no. Section 4, pp. 1401–1408, 2017.

[57] O. E. Cornejo Olivares, F. Pastore, and L. Briand, "Mutation analysis for cyber-physical systems: Scalable solutions and results in the space domain," *IEEE Transactions on Software Engineering*, pp. 1–1, 2021.

[58] G. Guizzo, F. Sarro, and M. Harman, "Cost measures matter for mutation testing study validity," in *Proceedings of the 28th ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering*, ser. ESEC/FSE 2020. New York, NY, USA: Association for Computing Machinery, 2020, p. 1127–1139. [Online]. Available: https://doi-org.proxy.bnl.lu/10.1145/3368089.3409742

[59] European Cooperation for Space Standardization., "ECSS-E-ST-40C - Software general requirements." 2009. [Online]. Available: http://ecss.nl/standard/ecss-e-st-40c-software-general-requirements/

[60] ——, "ECSS-Q-ST-80C Rev.1 - Software product assurance." 2017. [Online]. Available: http://ecss.nl/standard/ecss-q-st-80c-rev-1-software-product-assurance-15-february-2017/

[61] C. Cadar, D. Dunbar, D. R. Engler *et al.*, "KLEE: Unassisted and Automatic Generation of High-Coverage Tests for Complex Systems Programs." in *OSDI*, vol. 8, 2008, pp. 209–224.

[62] T. T. Chekam, M. Papadakis, and Y. Le Traon, "Mart: A Mutant Generation Tool for LLVM," in *Proceedings of the 2019 27th ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering*, ser. ESEC/FSE 2019. New York, NY, USA: Association for Computing Machinery, 2019, p. 1080–1084. [Online]. Available: https://doi.org/10.1145/3338906.3341180

[63] A. Denisov and S. Pankevich, "Mull it over: mutation testing based on LLVM," in *2018 IEEE International Conference on Software Testing, Verification and Validation Workshops (ICSTW)*. IEEE, 2018, pp. 25–31.

[64] LLVM, " The LLVM Compiler Infrastructure ," https://llvm.org/, 2021.

[65] M. Beller, C.-P. Wong, J. Bader, A. Scott, M. Machalica, S. Chandra, and E. Meijer, *What It Would Take to Use Mutation Testing in Industry: A Study at Facebook*. IEEE Press, 2021, p. 268–277. [Online]. Available: https://doi.org/10.1109/ICSE-SEIP52600.2021.00036

[66] G. Petrović and M. Ivanković, "State of Mutation Testing at Google," in *Proceedings of the 40th International Conference on Software Engineering: Software Engineering in Practice*, ser. ICSE-SEIP '18. New York, NY, USA: Association for Computing Machinery, 2018, p. 163–171. [Online]. Available: https://doi.org/10.1145/3183519.3183521

[67] P. Delgado-Pérez, I. Habli, S. Gregory, R. Alexander, J. Clark, and I. Medina-Bulo, "Evaluation of mutation testing in a nuclear industry case study," *IEEE Transactions on Reliability*, vol. 67, no. 4, pp. 1406–1419, 2018.

[68] R. Baker and I. Habli, "An empirical evaluation of mutation testing for improving the test quality of safety-critical software," *IEEE Transactions on Software Engineering*, vol. 39, no. 6, pp. 787–805, 2013.

Enrico Viganò is a Research and Development Specialist at the Interdisciplinary Centre for Security, Reliability and Trust (SnT), University of Luxembourg.

His research interests lie in automated testing for space software. He is currently involved in projects with the European Space Agency and industry partners from the space domain.

Oscar Cornejo is Senior Software QA Engineer at SES Luxembourg. Previously, Oscar had been a Research Associate at the Interdisciplinary Centre for Security, Reliability and Trust (SnT), University of Luxembourg. He obtained his PhD degree in Computer Science from the University of Milano - Bicocca in 2019.

His interests are in software engineering, focusing on automated software testing and program analysis. He is currently working on R&D projects in the space domain.

Fabrizio Pastore is Chief Scientist II at the Interdisciplinary Centre for Security, Reliability and Trust (SnT), University of Luxembourg. He obtained his PhD in Computer Science in 2010 from the University of Milano - Bicocca.

His research interests concern automated software testing, including security testing and testing of AI-based systems; his work relies on the integrated analysis of different types of artefacts (e.g., requirements, models, source code, and execution traces). He is active in several industry partnerships and national, ESA, and EU-funded research projects.

Lionel C. Briand is professor of software engineering and has shared appointments between (1) School of Electrical Engineering and Computer Science, University of Ottawa, Canada and (2) The SnT centre for Security, Reliability, and Trust, University of Luxembourg. He is the head of the SVV department at the SnT Centre and a Canada Research Chair in Intelligent Software Dependability and Compliance (Tier 1).

He has conducted applied research in collaboration with industry for more than 25 years, including projects in the automotive, aerospace, manufacturing, financial, and energy domains. In 2016, he received an ERC Advanced grant, the most prestigious European individual research award. He was elevated to the grades of IEEE and ACM fellow, granted the ACM SIGSOFT Outstanding Research Award (2022), the IEEE Computer Society Harlan Mills award (2012), and the IEEE Reliability Society Engineer-of-the-year award (2013) for his work on software verification and testing. His research interests include: Testing and verification, search-based software engineering, model-driven development, requirements engineering, and empirical software engineering.