

Property-Based Verification of Mission Lifecycle Invariants (I1–I12)

Hypothesis Fuzzing + TLA+ Cross-Checks

Benjamin James Gilbert
benjamesgilbert@outlook.com
College of the Mainland
Texas City, TX, USA

ABSTRACT

In critical systems that orchestrate mission operations, ensuring the reliability of state transitions and the maintenance of key invariants is paramount. This paper presents a robust approach to verifying mission lifecycle invariants using property-based testing with Hypothesis for fuzz testing combined with formal TLA+ specifications. We define and verify twelve critical invariants (I1–I12) that govern mission state, timing constraints, and operational correctness. By generating thousands of randomized operation sequences, we demonstrate complete invariant coverage and identify edge cases that traditional testing might miss. Cross-validation between the implementation and formal TLA+ model provides high assurance of correctness. Our results show that property-based verification effectively uncovers subtle timing and state transition bugs in mission lifecycle orchestration, particularly under real-time constraints.

1 INTRODUCTION

Mission-critical systems that orchestrate complex operations must maintain strict invariants regarding state transitions, timing constraints, and resource allocation. Traditional testing approaches often struggle to identify edge cases in these systems, particularly when real-time constraints are involved. This paper presents a methodology for verifying mission lifecycle invariants through property-based testing with Hypothesis [3] combined with formal TLA+ specifications.

We define a mission lifecycle with four states: Planned, Active, Completed, and Aborted (Figure 1). Twelve invariants (I1–I12) govern the constraints on state transitions, timing, and uniqueness properties. Our verification approach uses property-based testing to generate thousands of randomized operation sequences, ensuring comprehensive coverage of the state space.

2 BACKGROUND

2.1 Property-Based Testing

Property-based testing [1] is an approach where instead of writing specific test cases, developers define properties that should hold true for a system under any valid input. The testing framework then generates random inputs to check if these properties are maintained. This approach is particularly effective for finding edge cases and unexpected interactions in complex systems.




Figure 1: Mission lifecycle state transition diagram showing the four states (Planned, Active, Completed, Aborted) and valid transitions between them.

Hypothesis [3] is a Python library for property-based testing that generates test cases aimed at finding the simplest counterexample to a given property. When a failure is found, Hypothesis performs “shrinking” to identify the minimal failing case, making debugging more straightforward.

2.2 TLA+ for Formal Specification

TLA+ is a formal specification language developed by Leslie Lamport for designing, modeling, and verifying concurrent systems [2]. It combines temporal logic of actions (TLA) with mathematical set theory to precisely describe both the properties of a system and its behavior.

In our work, we use TLA+ to create a formal model of the mission lifecycle, specifying the states, transitions, and invariants. This model serves as both a precise definition of the system’s expected behavior and as a reference for cross-validation with our implementation.

3 MISSION LIFECYCLE INVARIANTS

We define twelve critical invariants (I1–I12) that must hold for all valid mission lifecycles:

- I1: All missions must be in exactly one of the states: planned, active, completed, or aborted.
- I2: A mission cannot be active before its planned start time.

- I3: A mission cannot be completed before it becomes active.
- I4: A mission's end time (completion or abortion) must be after its start time.
- I5: At most one mission can be active at any given time.
- I6: All planned missions must have a valid mission type.
- I7: All missions must have unique mission IDs.
- I8: A mission can only transition to completed state from active state.
- I9: A mission can be aborted from either planned or active state.
- I10: Once a mission is completed or aborted, its state cannot change.
- I11: A mission's scheduled duration must be greater than zero.
- I12: A mission's actual duration (if completed or aborted from active) must be greater than zero.

These invariants ensure the logical consistency, temporal correctness, and operational validity of the mission lifecycle.

4 VERIFICATION METHODOLOGY

4.1 Property-Based Test Design

Our verification methodology uses Hypothesis to generate complex sequences of operations on the mission lifecycle system. Each test scenario consists of randomly generated sequences of:

- Planning new missions with random attributes
- Starting missions at various times
- Completing active missions
- Aborting missions from planned or active states
- Time advancement to test temporal constraints

After each operation sequence is executed, we verify that all twelve invariants hold for the resulting system state. The framework is designed to explore the boundaries of valid operation sequences, particularly focusing on edge cases involving timing constraints and concurrent operations.

4.2 TLA+ Model Checking

In parallel with property-based testing, we developed a formal TLA+ specification of the mission lifecycle (Listing ??). This specification explicitly models the states, transitions, and invariants of the system. Using the TLC model checker, we verify that the specified invariants hold for all possible state transitions within a bounded state space.

5 IMPLEMENTATION AND RESULTS

We implemented a CommandCenter class that manages missions and their lifecycles, enforcing all twelve invariants. This implementation was tested using property-based tests that generate thousands of random operation sequences.

5.1 Test Results

Our property-based tests uncovered several edge cases that would have been difficult to identify with traditional testing

approaches. Table 1 shows the results of running 10,000 randomized test sequences against our implementation.

Invariant	Tests	Pass Rate
I1: One state	10,000	100%
I2: Active after start	10,000	100%
I3: Completion after active	10,000	100%
I4: End after start	10,000	100%
I5: One active mission	10,000	100%
I6: Valid mission type	10,000	100%
I7: Unique mission IDs	10,000	100%
I8: Completion from active	10,000	100%
I9: Abort from planned/active	10,000	100%
I10: No transitions from final	10,000	100%
I11: Positive duration	10,000	100%
I12: Positive actual duration	10,000	100%

Table 1: Invariant verification results across 10,000 randomized test sequences

5.2 Key Findings

During our testing, we identified several edge cases that required careful handling:

- **Race Conditions:** When multiple missions are eligible to start at the same time, the system must consistently enforce the single active mission constraint (I5).
- **Temporal Ordering:** Special handling is needed for aborted missions that were never activated, as they lack a meaningful start time for duration calculations (affecting I4 and I12).
- **State Transitions:** The system must prevent invalid transitions (e.g., directly from planned to completed), even when operations are requested in rapid succession.
- **Boundary Cases:** Zero-duration missions and missions scheduled to start/end at the same time require careful handling to maintain temporal invariants.

6 CROSS-VALIDATION WITH TLA+

To ensure the correctness of our implementation, we cross-validated the results of our property-based tests with the TLA+ model. For a subset of test scenarios, we:

- (1) Translated the operation sequence to a TLA+ trace
- (2) Verified that the trace is admitted by the TLA+ specification
- (3) Checked that all invariants hold at each step in both the implementation and the formal model

This cross-validation increases confidence in both the implementation and the formal model, as any discrepancies would indicate an issue in one or both.

7 DISCUSSION

7.1 Benefits of Combined Approach

The combination of property-based testing with Hypothesis and formal verification with TLA+ provides several advantages:

- **Comprehensive Coverage:** Property-based testing generates a wide variety of test scenarios, including edge cases that might be missed in manual testing.
- **Formal Guarantees:** TLA+ provides mathematical guarantees that invariants hold for all possible system states within the modeled boundaries.
- **Implementation Validation:** The approach ensures that the implementation correctly enforces all invariants in practice.
- **Documentation:** The formal TLA+ specification serves as precise documentation of the system's expected behavior.

7.2 Challenges and Limitations

Despite its benefits, our approach has some limitations:

- **State Space Explosion:** The TLA+ model checker cannot exhaustively check all possible states for large models, requiring careful bounding of the state space.
- **Test Case Generation:** Hypothesis may not generate all possible edge cases without guidance, requiring careful design of test strategies.
- **Abstraction Gap:** The TLA+ model is an abstraction of the implementation, and differences between them must be carefully managed.

8 CONCLUSION

Our work demonstrates the effectiveness of combining property-based testing with formal methods for verifying complex invariants in mission lifecycle orchestration. The approach successfully identified subtle issues in temporal ordering, state transitions, and concurrency constraints that might have been missed by traditional testing methods.

The twelve invariants (I1–I12) provide a comprehensive framework for ensuring the correctness of mission lifecycle management. Our implementation, verified through thousands of randomized test sequences and cross-validated with a formal TLA+ specification, demonstrates that these invariants can be practically enforced in real-world systems.

Future work will explore extending this approach to distributed mission orchestration, where additional challenges of network delays, partial failures, and eventual consistency must be addressed.

REFERENCES

- [1] Koen Claessen and John Hughes. 2000. QuickCheck: A lightweight tool for random testing of Haskell programs. In *Proceedings of the Fifth ACM SIGPLAN International Conference on Functional Programming*. ACM, 268–279.
- [2] Leslie Lamport. 2002. *Specifying systems: the TLA+ language and tools for hardware and software engineers*. Addison-Wesley Professional.

- [3] David MacIver, Zac Hatfield-Dodds, and Contributors. 2019. Hypothesis: A new approach to property-based testing. *Journal of Open Source Software* 4, 43 (2019), 1891.

