Crystal (ball): I Look at Physics and Predict Control Flow!

Sriharsha Etigowni, Shamina Hossain-McKenzie[+], Maryam Kazerooni[+], Bogdan Pinte[*], Kate Davis[*], Saman Zonouz

Department of Electrical and Computer Engineering
Rutgers University, [*]University of Texas A&M, [+]University of Illinois

Introduction

Motivating Scenario: Recent major attacks against unmanned aerial vehicles (UAV) and their controller software necessitate domain-specific cyber-physical security protection. Detecting unsafe states during or after the event aids mitigation measures, but predicting unsafe states provides more beneficial and significant impact for recovery.

JCR Architecture

High Level Architecture: Drones typically follow a dual processor architecture. The *flight control unit processor* takes care of the real-time sense-process-actuate executions. The *main processor* interacts with other peripheral devices. Crystal runs on the main processor to monitor each control logic update and its execution on the flight control unit processor.

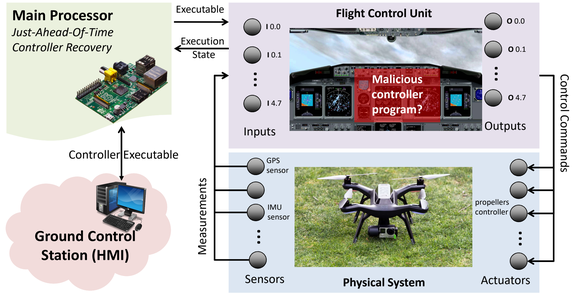

Figure: JCR's High-Level Architecture

Threat Model

We assume that the underlying software stack (e.g., operating system or firmware) and hardware are trusted, while the control logic, guidance and navigation algorithms on the drone's flight control processor can be malicious.

Safety Requirement

The safety requirements do not have to correspond to specific actuator outputs and could be defined for global system parameters. eg.,

$$\frac{Thrust\_on\_adjacent\_Motors}{dt} < |\gamma|, \quad \forall altitude > 0$$

Drone Physics Modelling

1. **Normal operational mode physical modelling** : The sensor and actuator values are estimated ahead of time by Extended Kalman filter.
2. **Failure mode data drive modelling** : The sensor and actuator values are estimated ahead of time by Neural network.
3. **Full flight operation mode modelling** : The sensor and actuator values are estimation ahead of time by a hybrid approach.

Cyber-physical Security Modelling

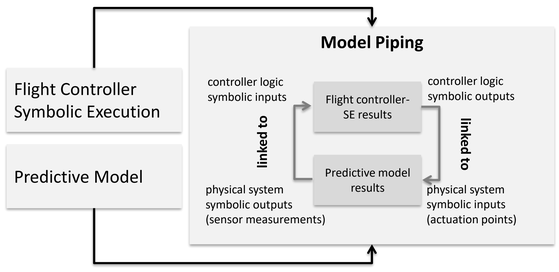

Figure: Symbolic hybrid model. The flight controller and physical dynamics are interconnected through sensing and actuation channel.

1. **Drone requirements →Controller output constraints** : Hybrid symbolic model is used to calculate the control unit's output constraints given the drone's global safety requirements.
2. **Flight controller output constraints →Controller's primary input constraints** : Control unit's code is analyzed to determine the primary constraints on its inputs that guarantee the above-mentioned output constraints if the code executes.
3. **Controller output constraints →Secondary input constraints** : Crystal emulates the flight dynamics symbolically using the above-mentioned control unit's output constraints and calculates a secondary set of constraints on the light control unit's inputs.
4. **Formal proof of the drone safety** : Formal theorem provers are used to prove the drone safety.

JAT Verification and Recovery

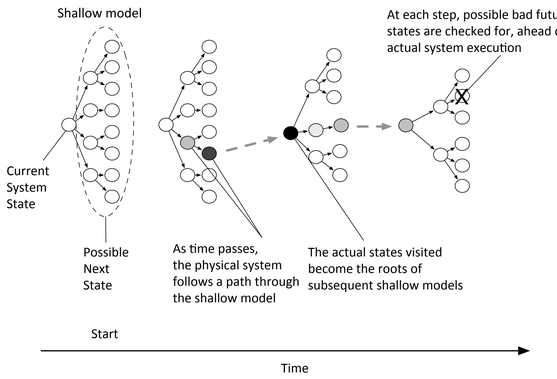

Figure: Discarding unreachable states. Ahead of time model checking using symbolic execution for pruning unreachable states to determine unsafe states before execution on UAV

1. **Just-ahead-of-time analysis** : Possible symbolic hybrid states of the drone are explored and the corresponding state-based finite state automaton is created for verification.
2. **Human assisted drone safety recovery** : If JAT encounters a potentially unsafe future state, Crystal asks for the operator's recommendation.
3. **Optimization for practical feasibility** : Runtime model pruning, parallel JAT and physics aware flight dynamics prediction are used to optimize.

Evaluations

Table: Average mean absolute error (MAE) for extended Kalman filter (EKF) and neural network (NN) model during minimal and heavy transitions

Sensor Data	NN MAE	EKF MAE
Roll during minimal transition	1.3242	0.1136
Roll during heavy transition	0.1713	0.8268
Yaw during minimal transition	1.9644	1.6359
Yaw during heavy transition	3.5643	18.5647

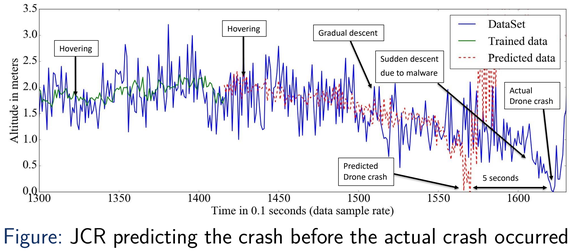

Figure: JCR predicting the crash before the actual crash occurred

Our experimental results show that JCR can proactively detect unsafe states, and recover the system with a negligible performance overhead.

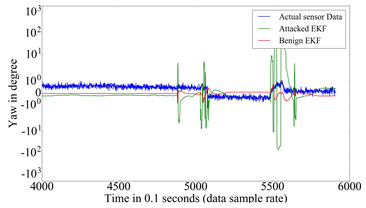

Figure: Predicting the attack on attitude and heading reference system (AHRS)

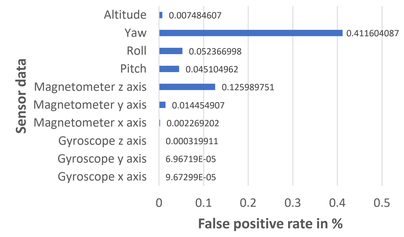

Figure: False positive rate due to sensor prediction

References

[1] Michael Robinson.
Knocking my neighbor's kid's cruddy drone offline.
In *Defcon*, 2015.

[2] Stephen McLaughlin, Saman Zonouz, Devin Pohly, and Patrick McDaniel.
A trusted safety verifier for process controller code.
In *Proc. ISOC Network and Distributed Systems Security Symposium (NDSS)*, 2014.

[3] Edward J Schwartz, Thanassis Avgerinos, and David Brumley.
All you ever wanted to know about dynamic taint analysis and forward symbolic execution (but might have been afraid to ask).

Acknowledgements

We would like to thank our sponsor:
National Science Foundation (NSF)