CHESS: Flight-Calibrated, AI-Generated Threat Scenarios for Air Defense

Abdeldjalil DEKKICHE – Guillaume MINARD
Opus Aerospace, 91220 Brétigny-sur-Orge, France - www.opus-aerospace.com

1. Identification of Use Cases

Air defense command and control (C2) chains (detection, threat characterization, engagement decision) are predominantly trained and evaluated against constructive simulation fed by synthetic threat models: trajectories modeled from estimated aerodynamics, radar cross-section (RCS) and infrared (IR) signatures seldom flight-validated, uncertainty rarely quantified, behavior mostly scripted. Operators and algorithms are therefore assessed against threats of unverified realism, with an unquantified model error. This gap widens as threats become faster and more maneuverable.

This paper presents CHESS, a scenario generation platform whose threat models are calibrated against real, fully instrumented suborbital flights. CHESS generates complex multi-vehicle raids, doctrinally plausible and physically achievable, for three uses: an evaluation bench for C2 chains and embedded algorithms; a training simulator for air defense operators; and a planning tool for ground-based air defense (GBAD). Scenarios can be replayed in minutes, without engaging any real asset.

Direct users are air defense units and training centers, C2 program teams, and evaluation authorities. CHESS provides a sovereign, physically grounded threat reference, reducing dependence on foreign or unvalidated threat models. It is inherently dual-use: the vehicle, its propulsion and its flight qualification derive from a civil suborbital launcher program.

The carrier vehicle is MESANGE, a suborbital launcher with an existing six-degree-of-freedom (6-DOF) simulation toolchain. Engine qualification on the test bench was completed in July 2026. The first ballistic flight (MES V1) is planned for Q1 2027, followed by a VX version used as a maneuvering target, then as a target carrier. A first version of the platform already runs in simulation, from radar set-up to raid debrief (Figure 2).

CHESS is built as an iterative loop (Figure 1). An initiation flight seeds the calibration base. Each cycle then reconciles predictions with measured trajectories, signatures and dispersion, quantifies the residual model error, and defines the next flight from envelope gaps and users' operational needs. The tool thus becomes usable early and its error bounds are expected to tighten with each cycle, without waiting for a full flight campaign. AI acts on the ground, generating and shortlisting raids within the calibrated envelope, and, in a later phase, in flight, selecting branches of a pre-loaded mission decision tree in bounded time and within the safety envelope.

At this stage, this paper and the associated presentation aim to test the concept against the forces' operational requirements, gather feedback on the use cases, and identify candidate experimentation frameworks, partners and flight-test opportunities.



Figure 1 — CHESS iterative loop: fly, measure, calibrate, refine. An initiation flight (⓪) seeds the models; each cycle then chains instrumented flight (①), ground processing (②), scenario study (③) and parameter refinement (④), whose outputs are loaded onboard for the next flight. Users (⑤) exploit the scenarios and feed back operational needs.

2. Proposed Technical Approach

(a) Flight-based threat model. Each instrumented suborbital flight yields a calibration dataset: 6-DOF trajectory reconstructed from ground and onboard sensors, RCS and IR signatures measured on the emulated (jet-like) threat profile, and dispersion statistics used to reconcile predictions with the observed flight. The tool, a constructive simulator, thus becomes calibrated, not just predictive: its threat models are anchored in measurements, and their known error bars make a C2 evaluation verdict defensible (Figure 1, steps ①–②). Flight data calibrate the tool; they do not directly train the onboard AI. With every flight, the models are refined and the residual error bound is expected to tighten; the remaining envelope gaps (step ④) define what the next flight must measure.



Figure 2 — CHESS Platform interface views: radar detection simulation around single threat object MESANGE

(b) Scenario generation and mission optimization. Raid generation (step ③) is a planning and optimization problem over trajectories, timings, approach axes and saturation patterns, under doctrinal, terrain and vehicle constraints, so that every scenario is doctrinally plausible and physically achievable within the flight-calibrated envelope. Generation proceeds in two stages. From the first flight, an optimization core (a constrained multi-objective search) treats the calibrated envelope as a hard feasibility constraint. A learned generative layer follows once the calibration base and a corpus of operationally validated scenarios exist. It is trained offline on validated scenarios only, neither directly on flight data nor in flight, and its output is filtered by the optimizer, so physical admissibility is enforced, not learned. Four criteria reduce candidates to a few reference scenarios per experimentation window: coverage of distinct C2 decision modes, discriminating power between configurations, plausibility against the calibrated envelope, and cost/readiness (simulation-only vs. flight-instrumented). These scenarios feed the C2 evaluation bench, operator training and GBAD planning, as well as step ④, where the next mission decision tree is compiled.

(c) AI across ground and flight: one continuous decision loop. On the ground, free of onboard computing constraints, step ④ compiles the shortlisted scenarios into a mission decision tree: a reference line plus a bounded set of contingency branches, like a chess opening book anticipating the plausible responses of the C2 chain. Once deployed in flight (a later phase), the adaptive onboard AI does not re-plan. Its four-step loop (sensing, processing, branch selection, trajectory adaptation) selects and locally adjusts the branch matching what it observes, so the vehicle stays representative of the emulated threat despite atmospheric, propulsion and alignment dispersion. Measurements are processed onboard, reducing reliance on telemetry for in-flight decisions in a degraded or contested electromagnetic environment. Across a multi-vehicle raid (King: emulated threat; Queen: main decoy; Pawn: decoy), a shared scenario clock and pre-computed, time-stamped waypoints preserve relative timing (saturation, axis convergence) without continuous inter-vehicle links. Computational frugality is a hard constraint: fixed-topology models, static memory allocation and verified worst-case execution time (WCET) bounds on the existing flight computer, with no added flight-critical hardware. The onboard AI is adaptive but always bounded: its models are trained offline on Monte Carlo runs of the calibrated simulator over each mission decision tree, and frozen before flight; it decides in flight but does not learn in flight. Every adaptation stays inside the flight safety envelope declared in the safety case, and the flight termination system remains unchanged, independent of the onboard AI, and retains full authority. Post-flight replay closes the loop (steps ②–④).

Maturity. Concept, moving to early prototype. Achieved to date: engine qualification on the test bench; an operational 6-DOF simulation toolchain inherited from the launcher program; a first version of the scenario platform running in simulation. Not yet achieved: exploitation of flight data; the calibration base remains empty by construction until the initiation flight (MES V1), the learned components in (b) and (c) are specified but not yet trained, and the onboard AI has not yet flown.

Open challenges. The signature measurement quality actually achievable in flight conditions; the representativeness gap between a suborbital carrier and the threat classes to be emulated, and how that gap is quantified rather than asserted; validation of a learned scenario generator against expert judgment of operational plausibility; and authorization of in-flight trajectory adaptation by the range safety authority.