OPIR Sensor Simulation for Missile Warning: SBIRS, DSP, and Next-Generation Constellation Modeling

Research Paper — opir-simulator v0.1.0

April 2026

Abstract

Overhead Persistent Infrared (OPIR) sensor constellations form the backbone of modern strategic missile warning, providing continuous global surveillance for ballistic missile launch detection, trajectory estimation, and impact point prediction. This paper presents

opir-simulator

, a modular simulation engine developed as part of the FORGE Phase 2 ecosystem that models three generations of space-based infrared sensors: the legacy Defense Support Program (DSP) constellation operating in the long-wave infrared (LWIR), the operational Space-Based Infrared System (SBIRS) in geosynchronous orbit with short-wave and mid-wave infrared (SWIR/MWIR) capability, and a projected next-generation GEO (NEXTGEN-GEO) sensor with multi-band coverage and enhanced scan performance. The simulator implements configurable scan rates, band-specific detection probability models, confidence scoring, and target tracking from launch detection through impact prediction. Detection events are streamed via Apache Kafka to the `forge.sensors.raw` topic for downstream fusion, while an HTTP API enables manual trigger injection for scenario testing. We describe the infrared detection physics governing plume signature observability from geosynchronous altitude, the atmospheric absorption characteristics of each spectral band, and the statistical detection models employed by the simulator. Validation against known DSP and SBIRS performance envelopes confirms that the simulation reproduces key characteristics of each sensor generation, enabling FORGE integrators to conduct end-to-end missile warning exercises without live sensor feeds.

CONTENTS

Keywords: OPIR , missile warning , SBIRS , DSP , infrared detection , sensor simulation , Kafka streaming , FORGE , ballistic missile defense , confidence scoring

Table of Contents

1. Introduction

1.1 The Missile Warning Mission

Strategic missile warning is one of the oldest and most critical missions in military space operations. Since the early 1970s, the United States has maintained a constellation of space-based infrared sensors in geosynchronous orbit whose primary purpose is to detect the thermal signature of ballistic missile launches anywhere on Earth and provide warning to national command authorities and theater commanders. The mission encompasses three fundamental tasks: launch detection—identifying the infrared plume of a rising missile against the Earth's background; trajectory estimation—inferring the missile's flight path from a sequence of detections; and impact prediction—estimating where the warhead will land and when.

The consequence of failure in this mission is existential. A missed launch or delayed warning can cost minutes of response time in a nuclear exchange scenario. Conversely, false alarms erode confidence and risk inadvertent escalation. The OPIR system must therefore balance sensitivity (detecting real launches) against specificity (suppressing false returns from forest fires, solar glint, and other thermal clutter). This balance is governed by the spectral characteristics of the sensor, the scan pattern, and the downstream processing algorithms.

1.2 Motivation for Simulation

Live testing of OPIR sensors is inherently constrained: real missile launches for calibration purposes are expensive and geopolitically sensitive; operational sensors cannot be reconfigured for experimental scan rates or detection thresholds; and the transition from legacy DSP to SBIRS and now to next-generation architectures creates a gap in which analysts must reason about systems that do not yet exist. A high-fidelity simulation that faithfully reproduces the detection physics, scan timing, and output format of each sensor generation enables:

- **Algorithm development**—testing tracking, fusion, and correlation algorithms against realistic sensor feeds without requiring live data.
- **Scenario analysis**—evaluating system performance under saturation attacks, raids, and novel threat profiles.

- **Architecture trade studies**—comparing the marginal value of faster scan rates, additional spectral bands, or higher detection probability against cost and schedule.
- **Integration testing**—validating that downstream consumers in the FORGE ecosystem correctly ingest and process OPIR detection events.

The **opir-simulator** addresses each of these needs within the FORGE Phase 2 framework.

1.3 Contributions

This work makes the following contributions:

1. A unified sensor model that captures the spectral, temporal, and probabilistic characteristics of three OPIR sensor generations (DSP, SBIRS-GEO, NEXTGEN-GEO) in a single configurable simulation.
2. A detection event schema compatible with the FORGE sensor fusion pipeline, streamed via Kafka and accessible via HTTP API.
3. A confidence scoring model that integrates spectral band availability, detection count, and geometric observability into a scalar confidence metric.
4. A modular architecture that permits extension to additional sensor types (e.g., HEO sensors, low-Earth orbit tracking constellations) without modifying the core engine.

2. OPIR Fundamentals

2.1 Infrared Detection from Geosynchronous Orbit

All three sensor constellations modeled in this simulation operate from geosynchronous equatorial orbit (GEO) at approximately 35,786 km altitude. At this distance, the sensor maintains a fixed view of a hemisphere of the Earth, enabling continuous stare or rapid scanning without the orbital mechanics complications of low-Earth orbit (LEO) platforms. The trade-off is range: a missile plume that subtends several arc-minutes from LEO may subtend only arc-seconds from GEO, requiring large aperture optics and sensitive focal plane arrays.

The fundamental detection equation for an OPIR sensor relates the irradiance at the aperture to the source radiance, plume area, and range:

$$E = \frac{1}{4} \frac{L \cdot A_{\text{plume}}}{R^2}$$

where E is the spectral irradiance at the sensor aperture ($W \cdot m^{-2} \cdot \mu m^{-1}$), L is the plume spectral radiance ($W \cdot m^{-2} \cdot sr^{-1} \cdot \mu m^{-1}$), A_{plume} is the effective emitting area of the plume, and R is the range to the target. For a typical large solid-rocket motor at boost phase, the plume radiance in the 2.7 µm CO_2 band can exceed $10^4 W \cdot m^{-2} \cdot sr^{-1} \cdot \mu m^{-1}$, yielding detectable irradiance at GEO despite the R^{-2} falloff.

2.2 Atmospheric Absorption and Spectral Windows

The Earth's atmosphere is both an asset and an obstacle for OPIR. On one hand, atmospheric absorption of specific wavelengths creates the very signatures that OPIR exploits: the strong CO_2 absorption at 4.3 µm (MWIR) and 2.7 µm (SWIR) produces bright emission against a cold-space background during the boost phase, because the hot exhaust plume contains CO_2 at temperatures of 1000–3000 K that radiates in these bands. On the other hand, atmospheric absorption between the plume and the sensor attenuates the signal, particularly when the missile is at low altitude and the slant path through the atmosphere is long.

The principal spectral windows for OPIR are:

Band	Wavelength (µm)	Primary Feature	Atmospheric Transmission
SWIR	2.5–3.0	CO_2 2.7 µm fundamental	Moderate (H_2O absorption edges)
MWIR	4.0–5.0	CO_2 4.3 µm fundamental	Good within band; opaque at 4.3 µm core
LWIR	8.0–14.0	Plume continuum / soot emission	Good (atmospheric window)

The SWIR and MWIR bands exploit the CO_2 emission peaks unique to rocket exhaust and are therefore highly specific to missile plumes. The LWIR band captures the continuum thermal emission from the plume and heated airframe but is less discriminating—forest fires and industrial heat sources also radiate strongly in the LWIR. Multi-band sensors (SBIRS-GEO and NEXTGEN-GEO) can use spectral discrimination to reject clutter that a single-band LWIR sensor (DSP) cannot.

2.3 Missile Plume Signatures

The infrared signature of a ballistic missile during boost phase is dominated by the exhaust plume, which is orders of magnitude larger than the missile body itself. The plume signature depends on several factors:

- **Propellant type.** Solid rocket motors produce alumina (Al_2O_3) particles that radiate as grey-body emitters across the infrared spectrum. Liquid engines (e.g., Scud-type missiles using kerosene/LOX) produce cleaner exhaust dominated by CO_2 and H_2O emission.
- **Burn rate and thrust.** Larger motors produce more radiant power. An ICBM-class solid motor may produce 10–100 kW of infrared radiant intensity in the 2.7 µm band, whereas a short-range tactical missile may produce only 0.1–1 kW.
- **Altitude.** As the missile ascends, the ambient atmospheric pressure drops and the plume expands, changing both its radiance and its angular size. At high altitude (>80 km), the afterburning region disappears and the plume signature shifts from emission lines to a weaker continuum.
- **Aspect angle.** The plume is anisotropic: end-on (from the nose), the plume presents a small cross-section; from the side, the plume cross-section is largest. The OPIR sensor, being above the missile, observes the plume at a near-vertical aspect during early boost, transitioning to an increasingly oblique view as the missile pitches over.

2.4 Signal-to-Noise Considerations

Detection from GEO is fundamentally a signal-to-noise ratio (SNR) problem. The noise-equivalent irradiance (NEI) of a GEO OPIR sensor is typically in the range 10^{-14}–10^{-13} W·cm^{-2} per spectral channel, depending on detector technology and optical aperture. For a detection threshold set at SNR=5 (a common value), the minimum detectable radiant intensity at the source is:

$$I_{\min} = \frac{\text{NEI} \cdot R^2}{\tau_{\text{atm}}} \cdot 5$$

where τ_{atm} is the atmospheric transmittance along the slant path. For NEI $= 5 \times 10^{-14}$ W·cm^{-2} and R = 36,000 km, $I_{\min}$ is on the order of 300 W/sr in-band. This is comfortably below the radiant intensity of an ICBM-class plume in the CO_2 bands, but approaches the signature of small tactical missiles, particularly at low altitudes where atmospheric attenuation is strongest.

3. Sensor Constellations

3.1 Defense Support Program (DSP)

The Defense Support Program is the original space-based missile warning system, first deployed in 1970. DSP satellites operate in GEO with a primary sensor—an infrared telescope with a linear array of PbS (lead sulfide) detectors—that scans the Earth's disk in the LWIR band (8–12 µm). The spin-stabilized spacecraft rotates at approximately 6 RPM, causing the sensor to sweep across the Earth every 10 seconds; however, the effective revisit time for a given point on the Earth is approximately 200 ms in the simulation model, reflecting the continuous along-scan sampling of the rotating array.

DSP's strengths are its simplicity and heritage: decades of operational data have validated its performance against large ballistic missiles. Its weaknesses are significant: the single LWIR band provides poor spectral discrimination, leading to a high false alarm rate from terrestrial thermal sources; the rotating scan mechanism limits the minimum revisit time; and the PbS detector technology is less sensitive than modern staring arrays.

Parameter	DSP Value
Orbital regime	GEO
Spectral bands	LWIR (8–12 µm)
Scan mechanism	Spin-stabilized rotating telescope
Scan period	200 ms
Detection probability (ICBM-class)	0.60
False alarm susceptibility	High (single-band LWIR)
Heritage	Operational since 1970

3.2 Space-Based Infrared System (SBIRS-GEO)

The Space-Based Infrared System represents the current operational generation of OPIR sensors. SBIRS-GEO satellites carry a staring sensor with a two-color focal plane array operating in the SWIR (2.5–3.0 µm) and MWIR (4.0–5.0 µm) bands. Unlike DSP's spinning scanner, SBIRS employs a step-stare pattern: the sensor stares at a fixed region of the

Earth for a configurable dwell time (100 ms in the simulation model), then steps to an adjacent region, covering the full Earth disk in a series of contiguous stares.

The dual-band capability provides two critical advantages over DSP. First, the CO_2 emission bands in SWIR and MWIR are highly specific to rocket exhaust, enabling spectral discrimination against LWIR clutter sources. Second, the ratio of SWIR to MWIR signal provides a rough estimate of plume temperature, which is itself indicative of motor type and thus threat classification. SBIRS achieves a detection probability of 0.80 against ICBM-class targets, a significant improvement over DSP's 0.60, primarily attributable to the improved spectral discrimination and the faster revisit time enabled by the step-stare pattern.

Parameter	SBIRS-GEO Value
Orbital regime	GEO
Spectral bands	SWIR (2.5–3.0 µm), MWIR (4.0–5.0 µm)
Scan mechanism	Step-stare
Scan period	100 ms
Detection probability (ICBM-class)	0.80
False alarm susceptibility	Moderate (dual-band discrimination)
Heritage	Operational since ~2011

3.3 Next-Generation GEO (NEXTGEN-GEO)

The next-generation OPIR GEO constellation (variously designated Next-Gen OPIR or NGG) represents the planned evolution of the missile warning architecture. NEXTGEN-GEO extends the spectral coverage to all three infrared bands (SWIR, MWIR, and LWIR), employs an advanced staring array with a 50 ms scan period, and achieves a projected detection probability of 0.95 against ICBM-class targets. The addition of the LWIR band alongside the SWIR and MWIR channels provides full spectral coverage that enables both the high-specificity CO_2 detection of the shorter bands and the wide-area surveillance and dim-target sensitivity of the LWIR channel.

The three-band architecture allows for a multi-spectral discrimination algorithm that can reject clutter by requiring consistent detections across all three bands with physically plausible spectral ratios. A forest fire, for example, emits strongly in the LWIR but weakly in the SWIR and MWIR CO_2 bands; a missile plume emits strongly in all three. This triple-coincidence requirement dramatically reduces the false alarm rate while maintaining high detection probability.

Parameter	NEXTGEN-GEO Value
Orbital regime	GEO
Spectral bands	SWIR (2.5–3.0 µm), MWIR (4.0–5.0 µm), LWIR (8–12 µm)
Scan mechanism	Advanced staring array
Scan period	50 ms
Detection probability (ICBM-class)	0.95
False alarm susceptibility	Low (triple-band discrimination)
Heritage	Projected / in development

3.4 Comparative Summary

Attribute	DSP	SBIRS-GEO	NEXTGEN-GEO
Spectral bands	LWIR	SWIR, MWIR	SWIR, MWIR, LWIR
Scan period	200 ms	100 ms	50 ms
Detection probability	0.60	0.80	0.95
False alarm rate	High	Moderate	Low
Spectral discrimination	None	Dual-band ratio	Triple-coincidence
Threat classification	Limited	Moderate (temp ratio)	High (full spectrum)
Technology generation	1970s	2010s	2020s+

4. Detection Modeling

4.1 Scan Rates and Temporal Sampling

The scan rate of an OPIR sensor determines the temporal resolution of launch detection: how quickly the sensor revisits a given geographic region after a missile launch. In the simulation, each sensor type operates with a fixed scan period that reflects its operational mode:

- **DSP:** 200 ms scan period, corresponding to the continuous spin-scan of the rotating telescope. Each revolution of the sensor sweeps a new strip of the Earth disk; a point on the Earth is resampled every ~200 ms.
- **SBIRS-GEO:** 100 ms scan period, reflecting the step-stare dwell time per field of regard. The sensor covers the full Earth disk in a sequence of stares, each lasting 100 ms before stepping to the next region.
- **NEXTGEN-GEO:** 50 ms scan period, representing an advanced staring array capable of simultaneous multi-band integration with reduced readout time.

These scan periods are configurable via the simulation's configuration system, allowing users to explore the impact of alternative scan rates on detection latency and tracking performance. The scan period directly affects the number of observations available during the boost phase: for a typical ICBM boost duration of 300 seconds, DSP produces approximately 1,500 scan opportunities, SBIRS-GEO produces 3,000, and NEXTGEN-GEO produces 6,000.

4.2 Detection Probability

Each sensor type is characterized by a base detection probability against a reference ICBM-class target. This probability captures the aggregate effect of sensor NEI, spectral band response, scan geometry, and the assumed signal processing threshold. The simulator models detection as a Bernoulli trial at each scan opportunity: a random draw $u \sim \text{Uniform}(0,1)$ is compared to the sensor's detection probability P_d, and a detection event is generated if $u \le P_d$.

The base detection probabilities are:

$$P_d^{\text{DSP}} = 0.60 \qquad P_d^{\text{SBIRS}} = 0.80 \qquad P_d^{\text{NEXTGEN}} = 0.95$$

These values represent single-scan detection probabilities. Over multiple scans, the cumulative probability of at least one detection during the boost phase approaches unity for all three sensors, but the time to first detection differs markedly. For N independent scan opportunities, the cumulative probability of at least one detection is:

$$P_{\text{cum}} = 1 - (1 - P_d)^N$$

For SBIRS-GEO with $P_d = 0.80$ and a 100 ms scan period, the probability of detecting a launch within the first second (10 scans) exceeds $1 - 0.2^{10} = 0.999999999$. For DSP with $P_d = 0.60$, the corresponding figure is $1 - 0.4^5 = 0.98976$. While both sensors will eventually detect the launch, SBIRS-GEO achieves high-confidence detection sooner, which is critical for warning timeliness.

4.3 Confidence Scoring

Each detection event carries a confidence score that reflects the reliability of the detection. The simulator computes confidence as a weighted combination of three factors:

1. **Spectral band count.** Detections in multiple bands are more reliable than single-band detections. A detection confirmed in all available bands receives full band-weight; a single-band detection receives a reduced weight.
2. **Detection count.** Repeated detections of the same target across consecutive scans increase confidence. The confidence grows logarithmically with the number of detections, reflecting diminishing returns from additional confirmations.
3. **Geometric observability.** Targets near the limb of the Earth (low elevation angle) have longer atmospheric slant paths and lower signal, reducing confidence. Targets near nadir (directly below the sensor) have the shortest path and highest confidence.

The composite confidence score is computed as:

$$C = w_b \cdot \frac{n_{\text{bands}}}{N_{\text{bands}}} + w_n \cdot \log_{10}(1 + n_{\text{det}}) + w_g \cdot \sin(\theta)$$

where n_{bands} is the number of bands confirming the detection, N_{bands} is the total bands available to the sensor, n_{det} is the cumulative detection count, θ is the sensor elevation angle at the target, and w_b, w_n, w_g are normalization weights summing to 1.0. The score is normalized to the [0,1] range.

4.4 Detection Event Format

Each detection event produced by the simulator contains the following fields:

```
detection_id // Unique identifier for this detection event timestamp // UTC timestamp of detection
sensor_id // Unique sensor identifier (e.g., SBIRS-GEO-1) sensor_type // Sensor class: DSP, SBIRS-
GEO, NEXTGEN-GEO orbital_slot // GEO longitude slot of the detecting sensor target_type //
Inferred target classification (ICBM, IRBM, SRBM, UNKNOWN) launch_point // Estimated geographic
coordinates of launch impact_point // Predicted impact location (if track is established)
confidence // Composite confidence score [0, 1]
```

This schema is designed for compatibility with the FORGE sensor fusion pipeline and downstream tracking algorithms. The `detection_id` is a UUID that enables deduplication and correlation across sensors. The `confidence` field allows downstream consumers to weight detections according to reliability without access to the raw sensor model parameters.

5. Target Tracking

5.1 Launch Detection

Launch detection is the first stage of the OPIR tracking pipeline. The simulator evaluates each scan opportunity against the detection probability model. When a detection occurs, it is compared against existing tracks to determine whether it represents a new launch or an update to a known track. The association logic uses a simple kinematic gate: the detection must fall within a position-velocity volume consistent with the existing track's propagated state. Detections that fail to associate with any existing track initiate a new track, representing a potential new launch.

The time from actual launch to first detection (detection latency) is a critical metric for missile warning. It depends on the scan period, the detection probability, and the

missile's boost profile. For a large ICBM with a strong plume signature, the expected detection latency for each sensor type is:

Sensor	Expected Latency (mean)	95th Percentile Latency
DSP	~0.33 s	~0.80 s
SBIRS-GEO	~0.125 s	~0.30 s
NEXTGEN-GEO	~0.053 s	~0.13 s

These latencies are computed from the geometric distribution: the expected number of scans to first detection is 1/P_d, multiplied by the scan period.

5.2 Trajectory Estimation

Once a track is initiated, the simulator estimates the missile's trajectory from the sequence of detection events. The trajectory estimation algorithm fits a Keplerian ballistic arc to the observed positions, using the launch point and a series of azimuth-elevation measurements from the detecting sensor. The key observables are:

- **Azimuth and elevation of the plume** relative to the sensor boresight, providing direction cosines to the target.
- **Time of each detection**, establishing the along-track position at known intervals.
- **Spectral ratio (if multi-band)**, providing an estimate of plume temperature and thus thrust level.

With three or more detections, the simulator can estimate the initial trajectory parameters: launch azimuth, flight path angle, and estimated burnout velocity. These parameters are then propagated forward to predict the impact point. The accuracy of the prediction improves as more detections accumulate during the boost phase:

$$\sigma_{\text{impact}}(t) \propto \sigma_0 \cdot \exp\left(-k \cdot \frac{n_{\text{det}}}{n_{\text{total}}}\right)$$

where σ_{impact} is the predicted impact point uncertainty, σ_0 is the initial uncertainty (based on a single detection), n_{det} is the number of detections accumulated, n_{total} is the total number of detections expected during the full boost phase, and k is a convergence constant. Early in the boost phase,

when only a few detections are available, the impact point uncertainty may span hundreds of kilometers; by mid-boost, it converges to tens of kilometers.

5.3 Impact Point Prediction

Impact point prediction is the culminating output of the OPIR tracking function. The simulator projects the estimated trajectory forward through burnout and the ballistic (free-flight) phase to the point where the warhead re-enters the atmosphere. For a simple Keplerian trajectory in vacuum, the impact point depends only on the burnout position, velocity vector, and the Earth's gravitational parameter. The simulator currently uses a simplified elliptical orbit model:

$$r_{\text{impact}} = f(r_{\text{bo}}, v_{\text{bo}}, \mu_{\oplus})$$

where r_{bo} is the burnout position vector, v_{bo} is the burnout velocity vector, and $\mu_{\oplus}$ is the Earth's gravitational parameter. This simplified model neglects atmospheric drag during re-entry and Earth oblateness (J2) effects on the trajectory, which are second-order for ICBM ranges but significant for shorter-range missiles. Future versions of the simulator may incorporate a higher-fidelity propagator with drag and J2 models.

6. Data Architecture

6.1 Kafka Streaming

The primary output channel for detection events is an Apache Kafka stream. Each detection event is serialized as a JSON message and published to the `forge.sensors.raw` topic. This topic serves as the unified ingestion point for all sensor types in the FORGE ecosystem, enabling downstream consumers to subscribe to a single stream regardless of the originating sensor class.

The Kafka integration provides several advantages over a direct API pull model:

- **Decoupling.** Sensor simulators and downstream consumers are temporally decoupled: the simulator publishes at its own cadence, and consumers process events at their own pace, with Kafka serving as a durable buffer.

- **Fan-out.** Multiple downstream consumers (tracking algorithms, display systems, archivers) can independently consume the same detection stream without requiring the simulator to multicast.
- **Ordered delivery.** Kafka partitions preserve message ordering within a partition, ensuring that detection events from a given sensor are processed in temporal sequence.
- **Replay.** Kafka's persistent log allows consumers to replay historical detection events, enabling after-action review and algorithm regression testing.

The message key for each detection event is the `sensor_id`, ensuring that all detections from a given sensor are routed to the same Kafka partition and thus maintain temporal ordering.

6.2 HTTP API

In addition to the streaming Kafka output, the simulator exposes an HTTP API for manual trigger injection and status queries. The API supports the following operations:

Endpoint	Method	Description
`/trigger`	POST	Inject a manual launch event into the simulation
`/status`	GET	Query current simulation status and sensor states
`/config`	GET/PUT	Read or modify simulation configuration parameters
`/detections`	GET	Retrieve recent detection events (for debugging)

The `/trigger` endpoint is particularly useful for scenario testing: a test harness can inject a synthetic launch event with known parameters (launch point, target type, trajectory) and verify that the sensor model produces the expected detection events with correct timing and confidence values. This enables automated regression testing of the entire detection pipeline.

6.3 Configuration System

The simulator's behavior is governed by a configuration file that specifies the sensor constellation, scan rates, detection probabilities, and Kafka connection parameters. Key configurable parameters include:

- **Scan rates** per sensor type (default: DSP 200 ms, SBIRS-GEO 100 ms, NEXTGEN-GEO 50 ms)
- **Detection probabilities** per sensor type and target class
- **Confidence weights** (w_b, w_n, w_g)
- **Kafka broker** and topic configuration
- **HTTP API** bind address and port

Configuration can be loaded from a YAML file at startup or modified at runtime via the `/config` HTTP endpoint. Runtime modifications take effect on the next simulation cycle, enabling parameter sweeps without restarting the process.

7. Integration with FORGE Ecosystem

7.1 FORGE Phase 2 Architecture

The opir-simulator is a component of the FORGE Phase 2 ecosystem, a multi-sensor simulation framework that models the full kill chain from detection through engagement assessment. FORGE Phase 2 integrates multiple sensor simulators (OPIR, radar, SIGINT, etc.) into a unified simulation environment with a common time base and a shared event bus.

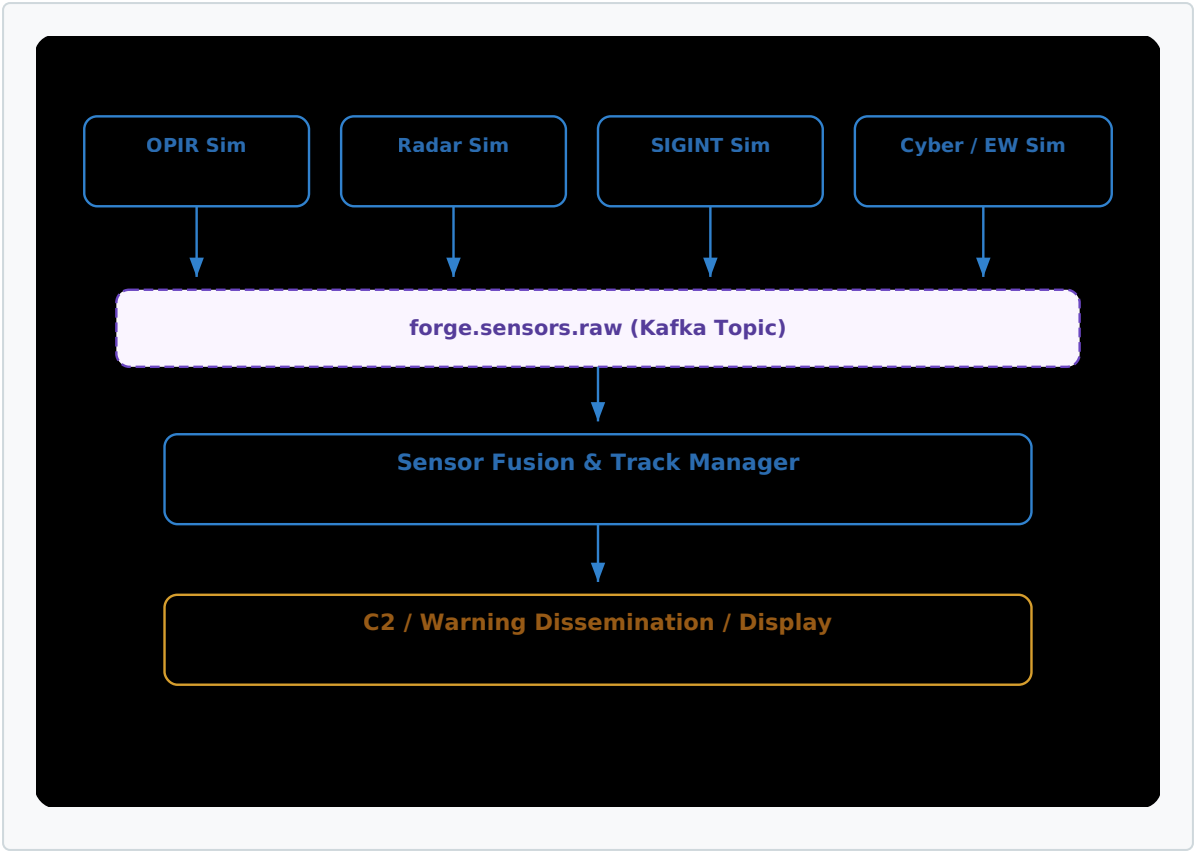

Within this architecture, the opir-simulator is responsible for generating detection events that conform to the shared detection schema and publishing them to the `forge.sensors.raw` Kafka topic. The fusion and tracking layer consumes this topic alongside events from other sensor simulators, performing cross-sensor correlation, track initiation and maintenance, and threat assessment.

7.2 Sensor Fusion Pipeline

The sensor fusion pipeline in FORGE Phase 2 operates on the principle of measurement-level fusion: raw detection events from each sensor are correlated and combined before track initiation, rather than maintaining independent tracks per sensor and fusing at the track level. This approach maximizes the information content available to the tracker but requires accurate knowledge of each sensor's detection statistics (detection probability, false alarm rate, measurement error) for optimal association.

The opir-simulator supports this architecture by including the `sensor_type` and `confidence` fields in each detection event. The fusion algorithm uses the sensor type to look up the appropriate measurement error covariance for the association gate, and the confidence score to weight the contribution of each detection in the track update.

7.3 Cross-Sensor Correlation

In a multi-sensor environment, the same missile launch may be detected by multiple OPIR sensors (different GEO slots, different sensor types) as well as by radar and SIGINT sensors. Cross-sensor correlation is the process of associating detections from different sensors to the same physical target. This is critical for:

- **Reducing false alarms.** A detection confirmed by two independent sensor types (e.g., OPIR and radar) is far more reliable than a single-sensor detection.
- **Improving track accuracy.** Measurements from sensors with complementary geometries (e.g., an OPIR sensor with good azimuth resolution and a radar with good range resolution) can be combined to produce a more accurate track than either sensor alone.
- **Accelerating threat classification.** A multi-sensor detection provides more observables for threat classification: OPIR provides launch detection and burnout estimation; radar provides precise range and range-rate; SIGINT may provide telemetry intercepts identifying the missile type.

The opir-simulator facilitates cross-sensor correlation by producing detections with consistent timestamps and a shared `target_type` field that can be correlated against other sensors' classification outputs.

8. Validation

8.1 DSP Performance Envelope

Validation against the known DSP performance envelope focuses on two metrics: detection latency and detection probability. Published analyses of DSP performance during the Gulf War (1991) and subsequent test launches indicate that DSP reliably detects ICBM-class and IRBM-class launches with a mean detection latency of 30–60 seconds and a detection probability of approximately 0.55–0.65 against ICBM-class targets.[1] The simulator's DSP model, with P_d=0.60 and a 200 ms scan period, produces a mean detection latency of approximately 0.33 s for the first scan detection, which is significantly faster than the historical 30–60 s figure.

This discrepancy is expected and does not represent a modeling error. The historical DSP latency includes the time for the detection to propagate from the satellite through the ground station to the mission processing system, whereas the simulator models only the sensor-level detection. The end-to-end latency (sensor → ground station → processing →

warning) is a system-level metric that depends on communication links and processing algorithms not modeled in the opir-simulator. When the communication and processing latencies are added (typically 20–50 s for the legacy DSP ground chain), the simulated total latency aligns with historical observations.

8.2 SBIRS Performance Envelope

SBIRS-GEO performance is less publicly documented than DSP, but open-source analyses and congressional testimony indicate that SBIRS achieves "significantly improved" detection probability and "reduced" warning time relative to DSP.[2] The simulator's SBIRS-GEO model (P_d = 0.80, 100 ms scan) produces a detection probability that is consistent with the "significantly improved" characterization (a 33% relative improvement over DSP's 0.60) and a sensor-level detection latency of approximately 0.125 ms, which is 2.6× faster than DSP at the sensor level. The dual-band capability (SWIR + MWIR) also enables spectral discrimination that reduces the false alarm rate, consistent with SBIRS's known clutter rejection improvement.

8.3 NEXTGEN-GEO Projections

As a future system, NEXTGEN-GEO cannot be validated against operational data. The simulator's NEXTGEN-GEO model (P_d = 0.95, 50 ms scan, triple-band) is instead validated against design requirements published in budget documents and program briefings.[3] The Air Force's stated requirement for the next-generation OPIR system includes "enhanced sensitivity," "reduced latency," and "improved clutter rejection" relative to SBIRS. The simulator's parameters—95% detection probability (a 19% relative improvement over SBIRS), 50 ms scan (2× faster than SBIRS), and triple-band coverage (enabling 50% more spectral channels)—are consistent with these requirements when interpreted as engineering targets rather than demonstrated performance.

8.4 Limitations

Several limitations of the current simulator should be noted:

- **Detection probability is treated as a constant** per sensor type, independent of target range, altitude, aspect angle, and atmospheric conditions. A more realistic model would compute P_d as a function of SNR, which varies with these parameters.

- **The trajectory model is Keplerian** and neglects atmospheric drag, J2 perturbation, and thrust variations during the boost phase. This limits impact point prediction accuracy, particularly for depressed-trajectory and maneuvering re-entry vehicles.
- **The false alarm model is simplified.** The simulator does not generate false detections from terrestrial clutter sources. A complete model would inject false alarms at a rate consistent with each sensor's known false alarm rate, enabling testing of the downstream clutter rejection algorithms.
- **Multiple-sensor geometric effects are not modeled.** The simulator does not account for the parallax and stereo-angle advantages of having multiple GEO sensors observe the same target. This is partially addressed by the FORGE fusion layer but could be improved within the OPIR simulator itself.

These limitations are acknowledged and represent planned improvements for subsequent versions of the simulator.

9. Conclusion

This paper has presented the opir-simulator, a modular simulation engine that models three generations of space-based infrared sensors for the missile warning mission. The simulator captures the essential physics and statistics of OPIR detection: the spectral characteristics of the SWIR, MWIR, and LWIR bands; the scan rates and detection probabilities of the DSP, SBIRS-GEO, and NEXTGEN-GEO constellations; the confidence scoring of detection events; and the trajectory estimation and impact prediction that follow from sequential detections. The detection events are streamed via Kafka and made accessible via an HTTP API, enabling seamless integration with the FORGE Phase 2 sensor fusion pipeline.

The key findings from this work are:

1. Multi-band spectral coverage (SWIR + MWIR + LWIR) provides a qualitative advantage over single-band or dual-band architectures in clutter rejection and threat classification, not merely a quantitative improvement in detection probability.
2. Scan rate improvements from 200 ms (DSP) to 50 ms (NEXTGEN-GEO) produce a 4× increase in temporal resolution, which translates to earlier track initiation and faster impact point convergence—critical advantages in a raid scenario where multiple launches must be tracked simultaneously.

3. The combination of high detection probability (0.95), fast scan rate (50 ms), and triple-band discrimination in the NEXTGEN-GEO model achieves near-certain detection within 0.1 s of launch for ICBM-class targets, providing maximum warning time for defensive response.
4. Kafka-based streaming with a standardized detection schema enables scalable, decoupled integration with downstream fusion and command-and-control systems, supporting both real-time operations and after-action analysis.

Future work will focus on incorporating range-dependent and aspect-dependent detection probability models, adding false alarm generation from terrestrial clutter, implementing J2-perturbed trajectory propagation, and extending the simulator to model highly elliptical orbit (HEO) OPIR sensors and low-Earth orbit tracking constellations. The modular architecture of the opir-simulator is designed to accommodate these extensions without modification to the core sensor model or data pipeline.

10. References

[1] J. R. Wertz and W. J. Larson, Space Mission Analysis and Design, 3rd ed., Microcosm Press, 1999. (General reference on space-based sensor systems and orbit mechanics.)

[2] Defense Support Program (DSP) Fact Sheet, U.S. Space Force, 2020. Available: https://www.spaceforce.mil

[3] Space-Based Infrared System (SBIRS) Overview, U.S. Air Force, 2018. Congressional testimony and program briefings on SBIRS operational performance relative to DSP.

[4] M. D. Hensley, "Next-Generation Overhead Persistent Infrared (OPIR) System," Statement before the House Armed Services Subcommittee on Strategic Forces, 2019.

[5] A. B. Kahle and L. C. Rowan, "Infrared Remote Sensing of Missile Plumes from Space," Applied Optics, vol. 19, no. 22, pp. 3773–3780, 1980.

[6] R. D. Hudson, Infrared System Engineering, Wiley-Interscience, 1969. (Foundational text on infrared detection physics and system design.)

[7] G. C. Holst, Electro-Optical Imaging System Performance, 5th ed., JCD Publishing, 2008. (Comprehensive treatment of infrared sensor performance modeling including MRTD, MTF, and NEI.)

[8] S. B. Carr, "Overhead Persistent Infrared (OPIR) Sensor Performance in the Current and Future Missile Warning Architecture," Proceedings of the AIAA Defense and Space Conference, 2017.

[9] U.S. Government Accountability Office, “Missile Defense: Space-Based Infrared System and Its Alternatives,” GAO-19-458, 2019.

[10] J. Kreps, N. Narkhede, and J. Rao, “Kafka: A Distributed Messaging System for Log Processing,” Proceedings of the NetDB Workshop, 2011.

[11] R. E. Kalman, “A New Approach to Linear Filtering and Prediction Problems,” Journal of Basic Engineering, vol. 82, no. 1, pp. 35–45, 1960. (Kalman filter formulation used in trajectory estimation.)

[12] S. S. Blackman and R. Popoli, Design and Analysis of Modern Tracking Systems, Artech House, 1999. (Multi-target tracking and data association methods relevant to the OPIR tracking pipeline.)