



FROM SENSOR TO DECISION

OBSERVATION DATA PROCESSING AND THE CCR DUAL-AI COMMAND LOOP

Technical brief · fire.ecopter.ai · sar.ecopter.ai · CONFIDENTIAL

How the surveyor platform's sensor streams become a shared fire model, how the two CCR planners reason over that model to offer the Commander first a global plan and then tactical options, how the referee resolves their disagreements, and a worked example on a two-front brush fire. CCR — Conflicting Computer Resolver — © Helen-B Advanced Engineering, Colorado.

01

The loop in one picture

Everything in this paper sits on one closed loop that runs continuously for the life of the incident. The surveyor platform observes; the processing tier turns observations into a single, time-stamped fire state model; the two CCR planners read that model and each propose what to do; the referee resolves the conflict and hands the Commander a decision; the aircraft execute; the surveyor observes the result. Loop period is seconds for tactics and minutes for strategy.

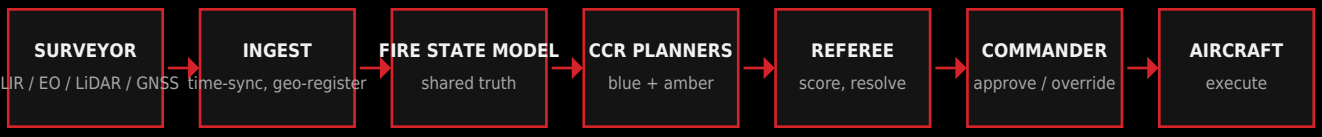


Figure 1 — the observe / model / plan / decide / act loop.

Two design rules govern the whole system. First, the model is the only thing the planners see: no planner reads a raw camera frame, so every decision is traceable to a model element with a timestamp and a confidence. Second, the Commander never receives raw AI output: they receive the referee's resolved recommendation, the losing option, and the reason one won, so the human sees the disagreement rather than a single confident answer.

02

What the observation platform produces

The surveyor is the fixed-wing or multirotor aircraft orbiting above the fire traffic area. Its payload is deliberately simple: every stream below is available off the shelf and the value is in the fusion, not in exotic hardware.

Stream	Sensor	Rate / format	What it is used for
Thermal	Radiometric LWIR (e.g. 640×512, 30 Hz)	Frames + per-pixel temperature	Fire front, hot spots, spot fires, drop effect, people/animals (SAR)
Visible	EO zoom camera, gimbal-stabilised	Frames, 30 Hz	Smoke column, fuel type, structures, roads, crews, water sources
Geometry	LiDAR (optional, pre-season or on-station)	Point cloud, LAS	Terrain, canopy, wires, towers — the obstacle layer of the altitude stack
Own position	RTK GNSS + IMU, gimbal encoders	100 Hz pose	Geo-registration of every pixel; without it nothing else works
Weather	Onboard pitot/static + RASS or nearby station	1 Hz wind, T, RH	Rate-of-spread, drop drift, density altitude
Air picture	ADS-B in, FTA radio digitised	events	Who else is in the fire traffic area; deconfliction

The thermal stream is the heart of it. Because the camera is radiometric, every pixel carries a temperature, so the pipeline can separate active flame (above ~500 °C), smouldering fuel, residual heat on already-burnt ground, and warm bodies, rather than guessing from brightness.

03

Processing tier: from pixels to a fire state model

Processing runs on the ground command node (and in reduced form on the surveyor itself so the aircraft can keep a coarse picture if the link drops). It is a chain of five stages; each stage writes to the model and nothing skips a stage.

Stage 1 — Time-sync and pose

Every frame, point and measurement is stamped with GNSS time and the sensor's full pose (position, attitude, gimbal angles). This is the single most important engineering discipline in the system: a 50 ms timing error at 30 m/s is 1.5 m on the ground.

Stage 2 — Geo-registration

Each thermal and visible pixel is projected through the camera model and the pose onto the terrain surface (DTM from LiDAR or the best available public elevation model). Output: a georeferenced temperature raster at 0.5–2 m ground sample distance, re-generated every few seconds and accumulated into a mosaic. Accuracy budget: 1–3 m with RTK, which is adequate for planning; the centimetre layer for the actual drop is a separate system (see the precision-guidance paper).

Stage 3 — Segmentation and tracking

The temperature raster is classified into active front, interior burning, burnt-out, hot spot (isolated, small), and unburnt. Connected regions are tracked frame-to-frame so each fire or spot fire keeps a stable identity (F1, F2, S1 …) and a history: when it was first seen, how fast its perimeter is moving, in which direction. Tracking is what lets the system notice a spot fire that was lit later and add it to the plan without a human spotting it.

Stage 4 — Fire behaviour estimate

For each tracked front the system estimates rate of spread and flame length from the perimeter history, wind, slope and the fuel class read from the visible imagery and any fuel map. A simple Rothermel-type model is enough: the goal is a 10–30-minute projection of where each front will be, with an uncertainty band, not a research-grade simulation.

Stage 5 — Values at risk and constraints

Structures, roads, crews on the ground, water sources, wires and towers, no-drop zones and the fire traffic area boundaries are overlaid from the visible stream, the LiDAR obstacle layer, and incident data. The result is the fire state model: a single geospatial object store that the planners query.

The model, in one line each

```
fires[]: id, polygon(t), front_segments[], ros_m_per_min, flame_len_m, projection(t+10, t+30),
confidence
hotspots[]: id, point, temp, first_seen, growing?
assets_at_risk[]: id, type, polygon, time_to_impact_min
aircraft[]: id, type, position, fuel/water state, capability, last_drop
water_sources[]: id, point, type (dip tank / pool / lake), approach_ok?
obstacles[]: wires, towers, terrain, no-fly volumes (from LiDAR + charts)
weather: wind vector(z), T, RH, density altitude, forecast(t)
```

04

CCR — two planners, one referee, one Commander

CCR — Conflicting Computer Resolver — is deliberately built on disagreement. Two independent planners are given the same fire state model and the same objective, but are built to think differently. Each produces a complete, ranked set of options. A referee compares them against a fixed scorecard, resolves the conflict, and presents the Commander with the chosen option, the rejected option, and the margin. The Commander approves, picks the other one, or overrides.

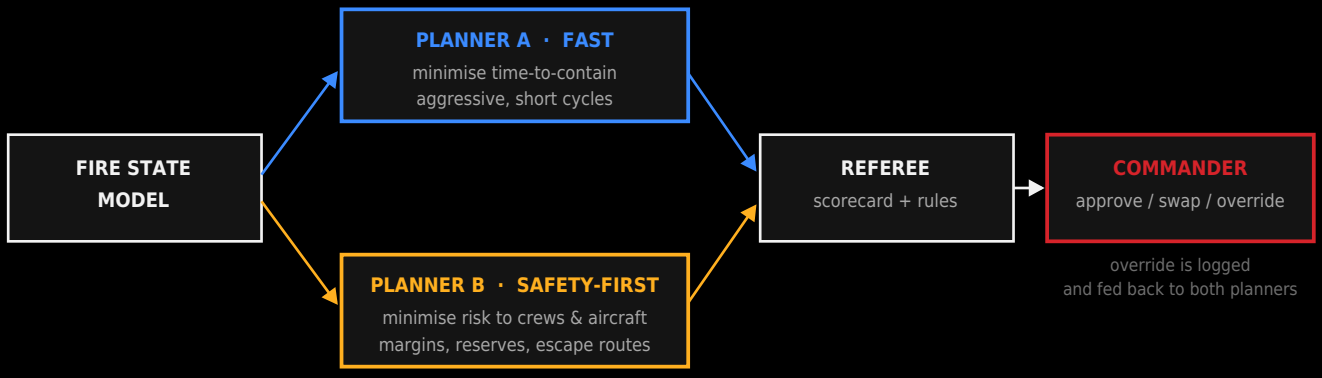


Figure 2 — CCR architecture. Blue and amber are the two planners' identities throughout the operator interface.

Why two planners rather than one good one

A single optimiser always returns a single confident answer, and a Commander has no way to tell a well-founded answer from an artefact of the model. Two planners with different objectives expose the shape of the trade-off: when they agree, the Commander can act quickly; when they disagree, the disagreement itself is the information, and it is presented in those terms (“A gets there 4 minutes sooner; B keeps 12 minutes of fuel margin instead of 6”).

Planner A — fast

Objective: minimise time to contain, weighted by values at risk. Favours short dip-to-drop cycles, direct routing, tight sequencing of multiple aircraft on the same front, and accepting thinner fuel reserves. It is the planner that wins on a small fire with a structure in its path.

Planner B — safety-first

Objective: minimise risk to crews, aircraft and the mission, subject to containment within a target window. It budgets fuel and water reserves, enforces escape routes and separation from the ground crews, prefers upwind and uphill approaches, and will trade minutes for margin. It is the planner that wins in wind, at night, with wires, or when more than two aircraft share a fire traffic area.

Both planners run the same search (a receding-horizon planner over the fire state model with the projection as the cost surface), but with different cost functions and different hard constraints. Running them on separate hardware with separate code paths is what makes a shared bug or a shared blind spot unlikely.

The referee

The referee is not a third AI with its own opinion. It is a deterministic scorecard with published weights, so an after-action review can reproduce every decision. For each pair of competing options it computes:

- Containment: projected time until every front is held, from the model's own projection.
- Protection: time-to-impact on the most exposed value at risk, before and after the option.
- Crew risk: minimum separation from ground crews, approach over or across the fire, visibility, wind.
- Aircraft risk: fuel and water margin at the worst point, obstacle clearance, density altitude, number of aircraft in the pattern.
- Reversibility: how much of the option can be abandoned mid-way without leaving an aircraft committed.

Hard rules sit above the scorecard and cannot be outscored: no option that violates an aircraft's envelope, puts a drop on a crew, enters a no-fly volume, or leaves fuel below the published reserve is ever presented. If both planners' options fail a hard rule, the referee reports that and asks for human input rather than inventing a plan.

The Commander

The Commander (the air operations chief on a fire, the SAR chief on a search) sees three things on the tablet: the resolved option, the rejected option, and the referee's margin with the two or three scorecard lines that decided it. They can approve (one tap), swap to the rejected option, or override with their own instruction. Every override is logged with context and fed back to both planners as a constraint for the rest of the incident, so the same disagreement is not re-raised every cycle.

05

Two planning horizons: global first, then tactical

CCR runs the same two planners at two horizons. The global horizon is the incident: which fires, in which order, which aircraft, which water source, how long. The tactical horizon is the next few minutes for one aircraft: the drop line, the approach, the altitude, the abort point. Global options are produced once at the start and re-issued when the model changes materially (a new spot fire, a wind shift, an aircraft going unserviceable); tactical options are produced every cycle for every aircraft.

	Global options (strategy)	Tactical options (execution)
Horizon	30–120 min	2–10 min
Re-plan trigger	New fire / spot fire, wind shift, aircraft state change, Commander override	Every cycle (5–15 s) while an aircraft is tasked
Unit of decision	Fire priority order, aircraft allocation, water source, sequencing, holding points	Drop line start/end, heading, release altitude, speed, dip site, abort point
What the Commander sees	A plan on the map with ETA to containment and risk summary, A vs B	A drop card per aircraft: line, timing, go/no-go, A vs B differences highlighted
Typical disagreement	Which fire first; one aircraft per fire vs both on the worst one	Approach into wind over the burnt area vs crosswind with better escape

Separating the horizons keeps the Commander's attention where it belongs. Global decisions are few and consequential and deserve a considered approval. Tactical decisions are many and fast; once a global plan is approved, tactical options that stay inside it are executed on a pre-approved envelope — the Commander is alerted only when the two planners disagree beyond a threshold, or when an option would leave the approved envelope.

06

Worked example: a two-front brush fire

The scenario is small on purpose, so the reasoning can be followed line by line. It is the kind of fire the March 2027 drone demonstration will reproduce at candle scale.

Situation at T+0

- Surveyor on station at 2,500 ft AGL. Wind 12 kt from the south-west, gusting 18. Mid-afternoon, RH 14 %.
- Model has tracked F1: main fire, 2.1 ha, front moving north-east at 9 m/min, flame length 2–3 m in brush. Projection puts the front at a ranch house (asset R1) in 22 minutes.

- S1: a spot fire 300 m downwind of F1, 0.05 ha, growing, first seen 90 s ago.
- Two suppressors available: H1 (Bell 205, 320 gal bucket, 11 min to dip site D1) and H2 (same, currently at D1 filling). Ground crew G1 (6 people) holding a road on F1's south flank.

Global phase — the planners disagree

Planner A (fast): both helicopters on S1 immediately, two drops back-to-back, then both to F1's north-east head. Reasoning: S1 is small now and will not be in ten minutes; killing it while it is 0.05 ha buys back the whole north-east flank. ETA containment 48 min. Worst-case fuel margin H1: 9 min.

Planner B (safety-first): H2 (already loaded) to S1, H1 straight to F1's head between the front and R1, then both cycle on F1. Reasoning: R1 is the only life-safety exposure; a single aircraft on S1 is enough at 0.05 ha; splitting keeps one aircraft always between the fire and the house. ETA containment 56 min. Worst-case fuel margin H1: 16 min.

Referee: containment favours A by 8 min; protection favours B (time-to-impact on R1 rises from 22 to 41 min under B, to 31 under A); aircraft risk favours B (fuel margin); crew risk equal. Published weights put protection of an occupied structure above containment time. B wins, margin 0.14. The card shows the Commander both plans on the map with the three deciding lines. The Commander approves B in one tap, adding one override: “G1 is pulling back to the road junction, treat the south flank as unstaffed.” Both planners receive that constraint.

Tactical phase — H2's first drop on S1

Every 10 seconds, with H2 inbound, both planners issue a drop card for H2 inside the approved global plan:

	Planner A	Planner B	Referee
Drop line	Along S1's long axis, NE→SW, 60 m	Across S1's head, into wind, 45 m	—
Approach	Downwind, straight in from current heading, 40 s sooner	Into wind over burnt ground of F1's tail	B: downwind drop drifts 15–20 m at 12 kt gusting 18
Release height	40 ft AGL for a tight pattern	60 ft AGL, wider but no rotor-wash flare-up	B: gust factor
Abort point	Over S1	500 m out, with escape to the SE	B
Result			B selected, margin 0.22 — above the alert threshold, so the Commander is pinged and confirms

H2 flies B's card. Fifteen seconds after the drop the surveyor's thermal mosaic shows S1's head cooled below 150 °C with a 6 m hot remnant on the south edge. The model updates, both planners now agree on a second short drop on the remnant, the referee marks it as a consensus option, and it executes without a Commander prompt because it is inside the pre-approved envelope. The Commander's screen shows S1 turning from red to grey.

Re-plan at T+14

The surveyor's tracker opens a new object: S2, 0.02 ha, 450 m north of F1, in the lee of a ridge. This is a material model change, so the global planners re-run. This time they agree — H2, now closest and lightly loaded, takes S2; H1 continues the F1 head — and the referee presents a consensus update rather than a conflict. The Commander approves. The whole re-plan, from first thermal detection to an aircraft tasked, took under a minute and required one tap.

What the example shows

- The planners' disagreement at the global level was real and informative: speed versus structure protection. The Commander saw the trade-off in two lines, not a page of reasoning.
- At the tactical level the disagreement was about wind and drift, and the safety-first planner was simply right; the referee's scorecard captured that without a human having to notice the gust factor.
- Consensus options flow through without friction. The human is interrupted only when the machines disagree or the plan changes shape.
- Every step is reproducible from the logged model snapshot, the two option sets and the scorecard, which is what a post-incident review — or a certification authority — will ask for.

07

Implementation notes for the demonstration

- Model store: a geospatial time-series database on the command node; every object versioned; the surveyor keeps a reduced replica.
- Planners: same planner core, two cost functions, two hard-constraint sets, built and run as separate processes on separate machines. For the demo both can run on one ground laptop plus the Orin; the separation is logical first, physical later.
- Referee: a few hundred lines of deterministic scoring with a versioned weights file. The weights file is the policy document and is what an agency will want to review.
- Commander UI: the tablet. Map, two option overlays in blue and amber, three-line justification, approve / swap / override. Override text is a first-class input to the next planning cycle.
- Demo scaling: candles replace fire fronts, the thermal camera on the surveyor drone sees them as hot spots, the two X650-class suppressors are H1 and H2, and the dip site is a bucket. The pipeline, the two planners, the referee and the tablet are the real ones; only the cost surface is scaled down.

CCR — Conflicting Computer Resolver — is © Helen-B Advanced Engineering, Colorado. This document is confidential to ecopter.ai and its partners.