

# Equinox: Decentralized Scheduling for Hardware-aware Orbital Intelligence

Ansel Kaplan Erol, Divya Mahajan



**Problem:** Low-Earth Orbit satellites face an "inter-temporal coupling" where processing tasks during sunlit phases can exhaust battery reserves needed for high-value events during eclipse. Existing schedulers either prioritize task value while ignoring hardware limits or enforce resource safety without accounting for the varying scientific importance of data.


**Our Approach:** We introduce **Equinox**, a decentralized runtime that compresses complex hardware states (battery, thermal, and congestion) into a single, state-dependent "marginal cost of execution." This barrier-based signal enables autonomous, value-aligned load shedding and serves as a coordination primitive for inter-satellite offloading without requiring global state.

**Findings:** In a 143-satellite simulation, **Equinox** achieved 20% higher scientific goodput with 2.2x greater battery reserves than priority-based baselines by strategically shedding low-value tasks. We show preserving energy headroom is critical for maximizing long-term orbital throughput.



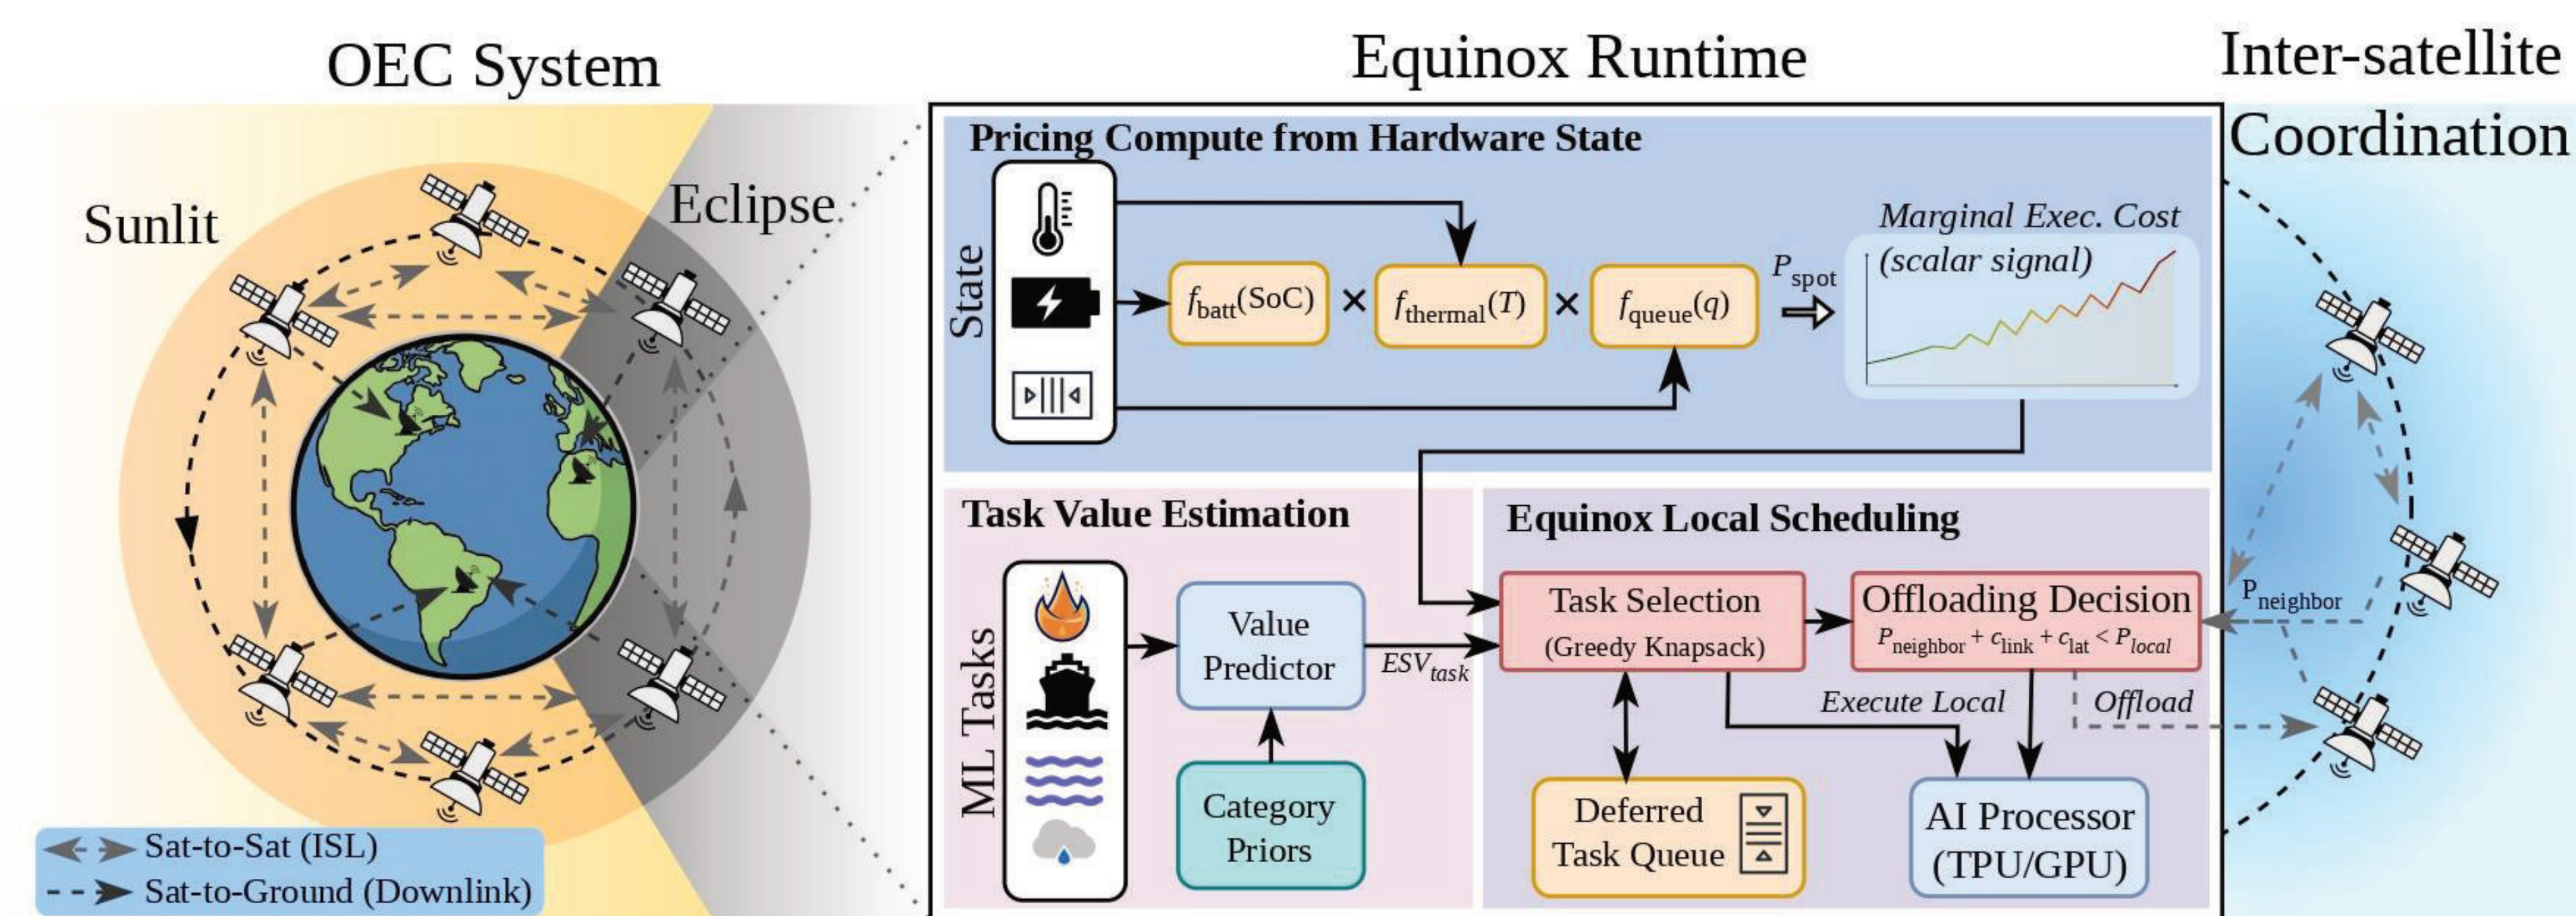
Preprint/Code

## Inter-temporal Coupling for Orbital Scheduling



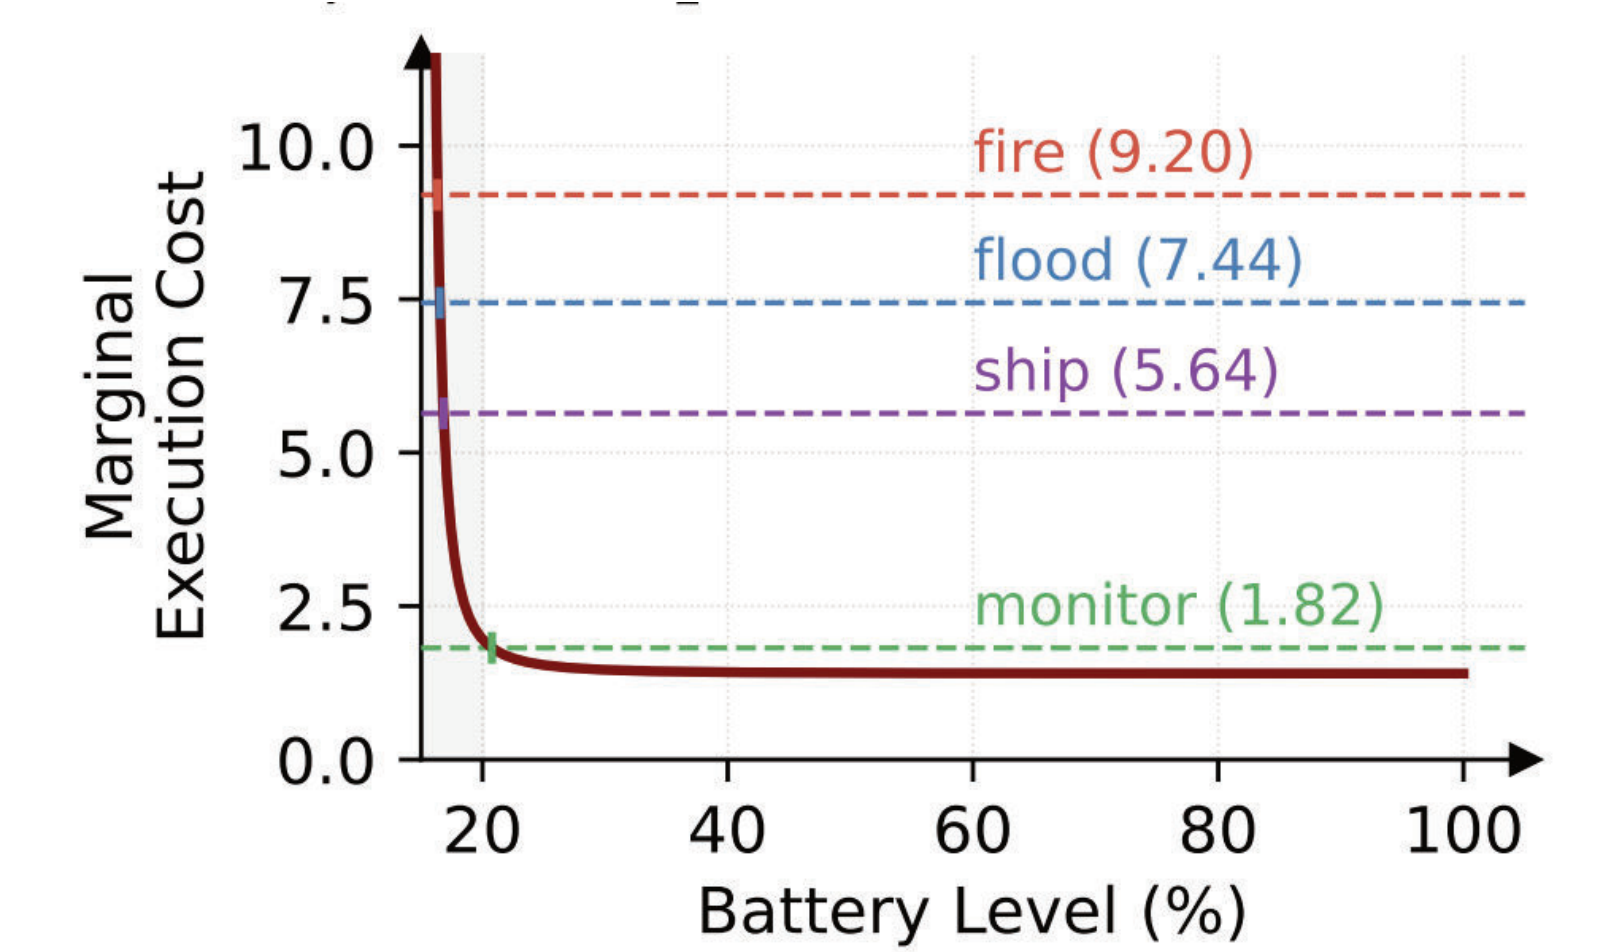
Scheduling for OEC satellites must maintain processing power for the 30% of orbit spent in eclipse, or risk missing critical events!

## Equinox Overview



Overview of Equinox's distributed resource management, illustrating per-satellite marginal execution cost, local scheduling, and ISL arbitration components.

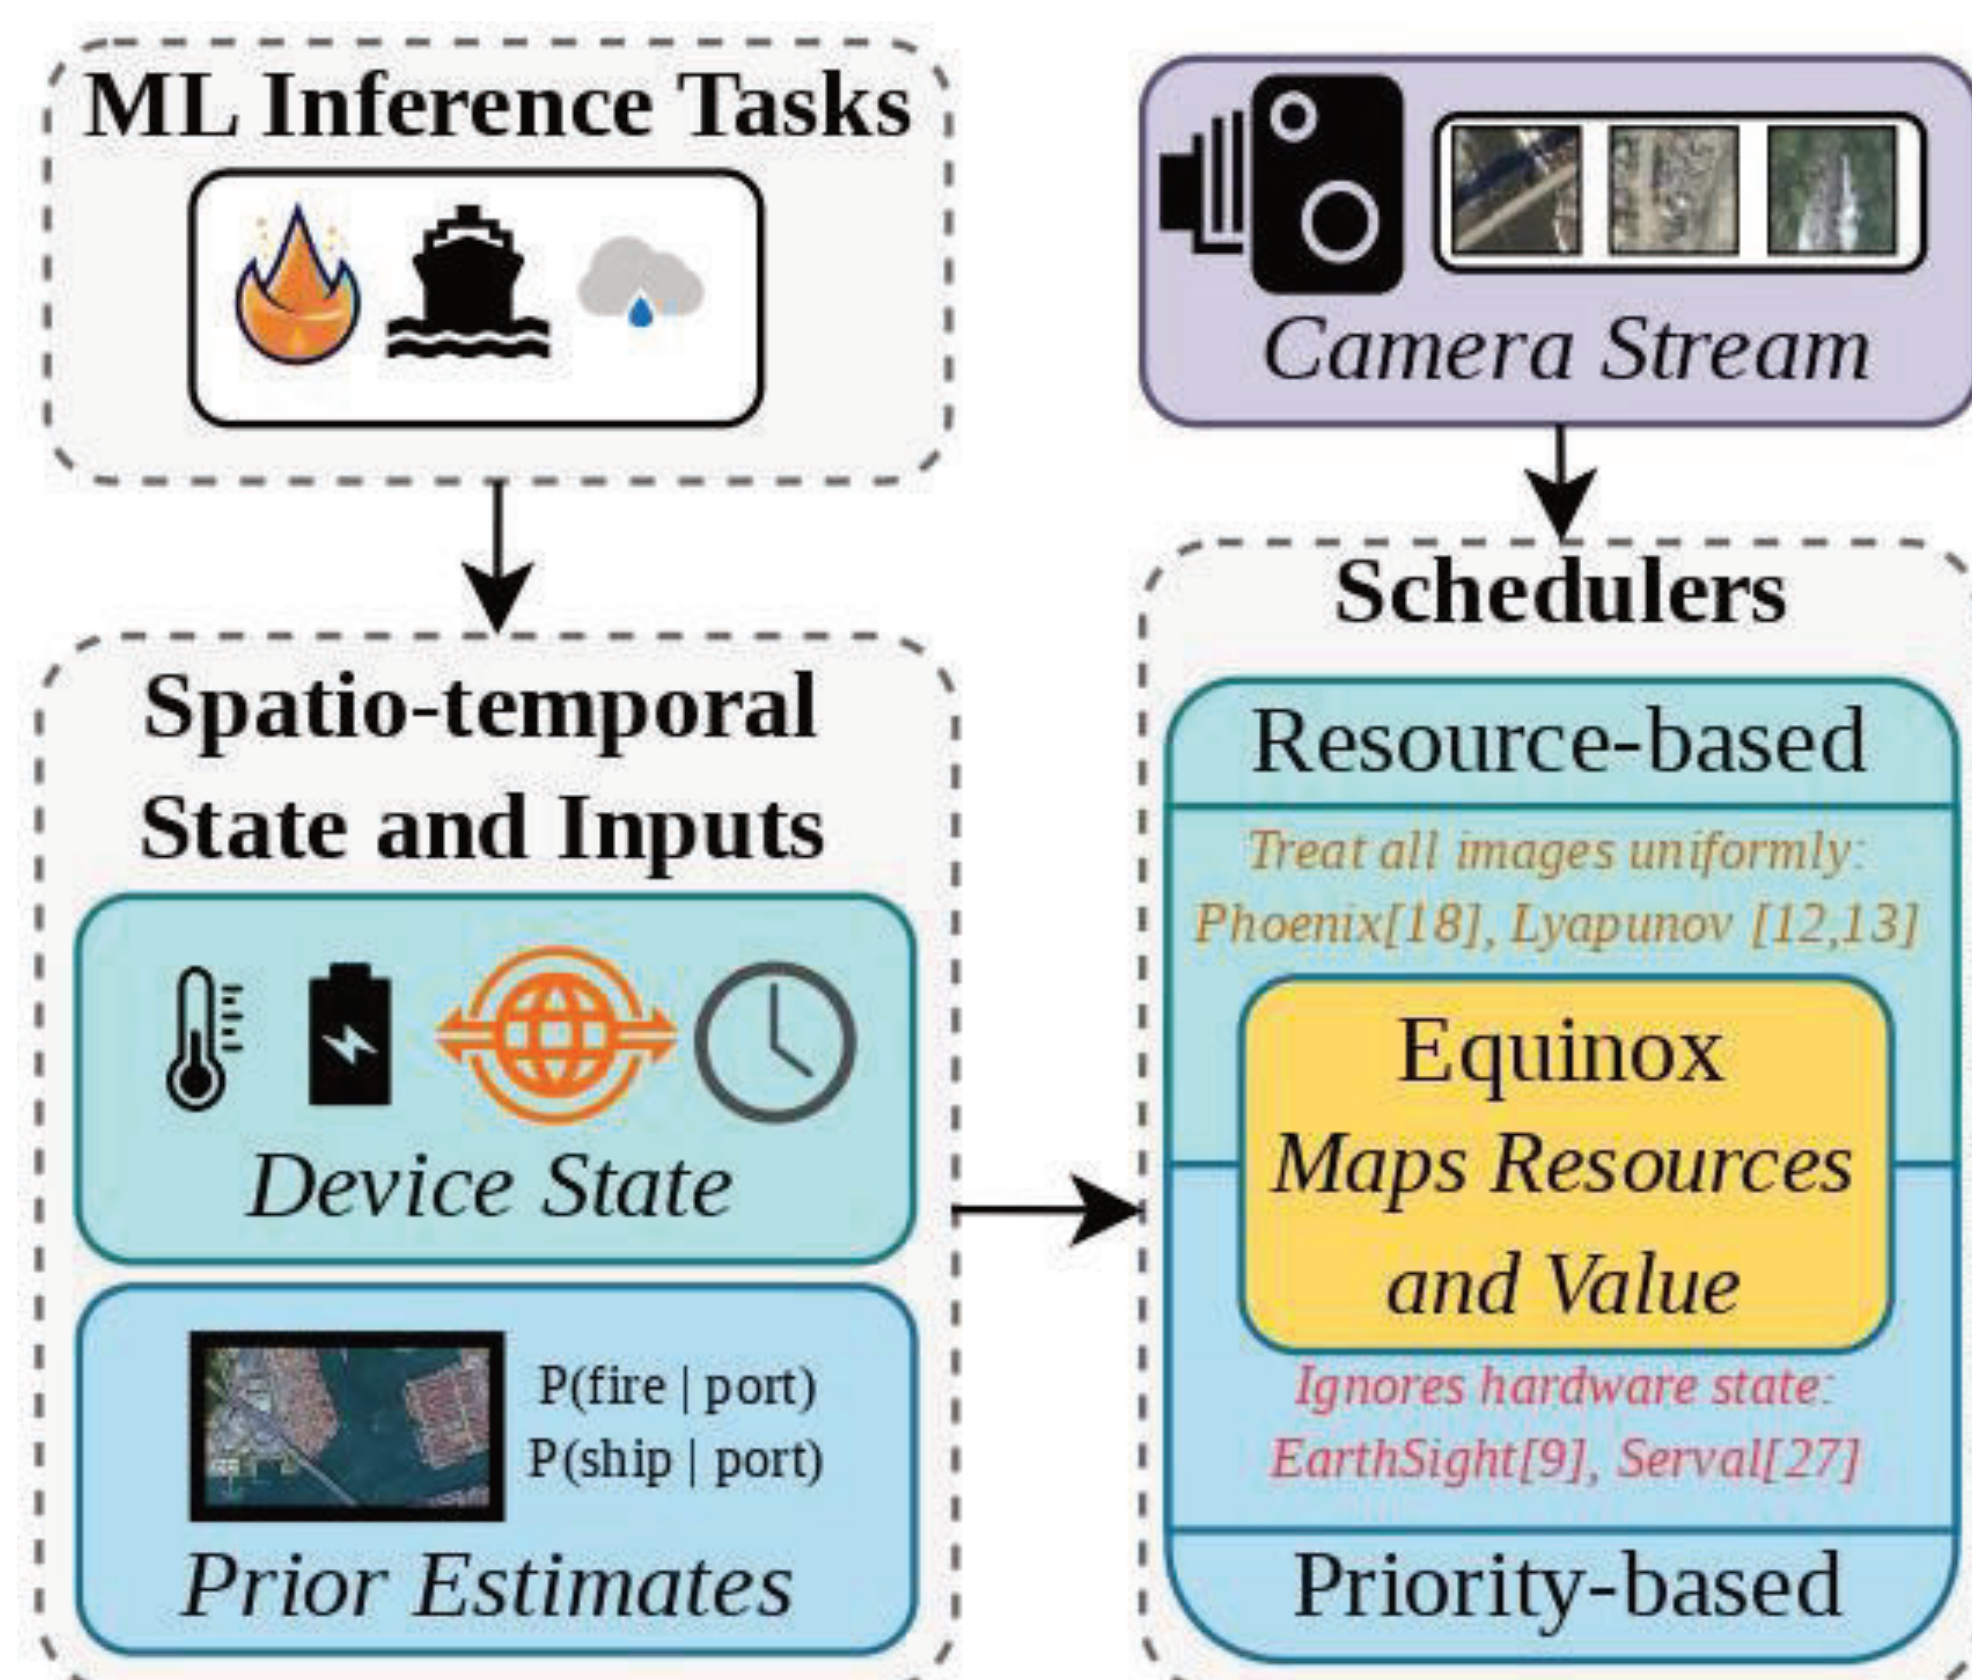
## Staggered Dropout



As battery declines, low-value work is shed naturally as a task's expected value cannot overcome the marginal execution cost.

## Need for Co-optimization

Device state and prior estimates must be considered together for effective scheduling.



However, *orbital constraints* require the scheduler have minimal compute and bandwidth overhead!

## System Metrics

### (1) Scientific Value and Goodput

Each detection category (flood, fire, etc.) is assigned a value for its successful detection. *Scientific goodput* measures how much total value all satellites collect per hour, i.e., weighted throughput.

### (2) Battery Reserves

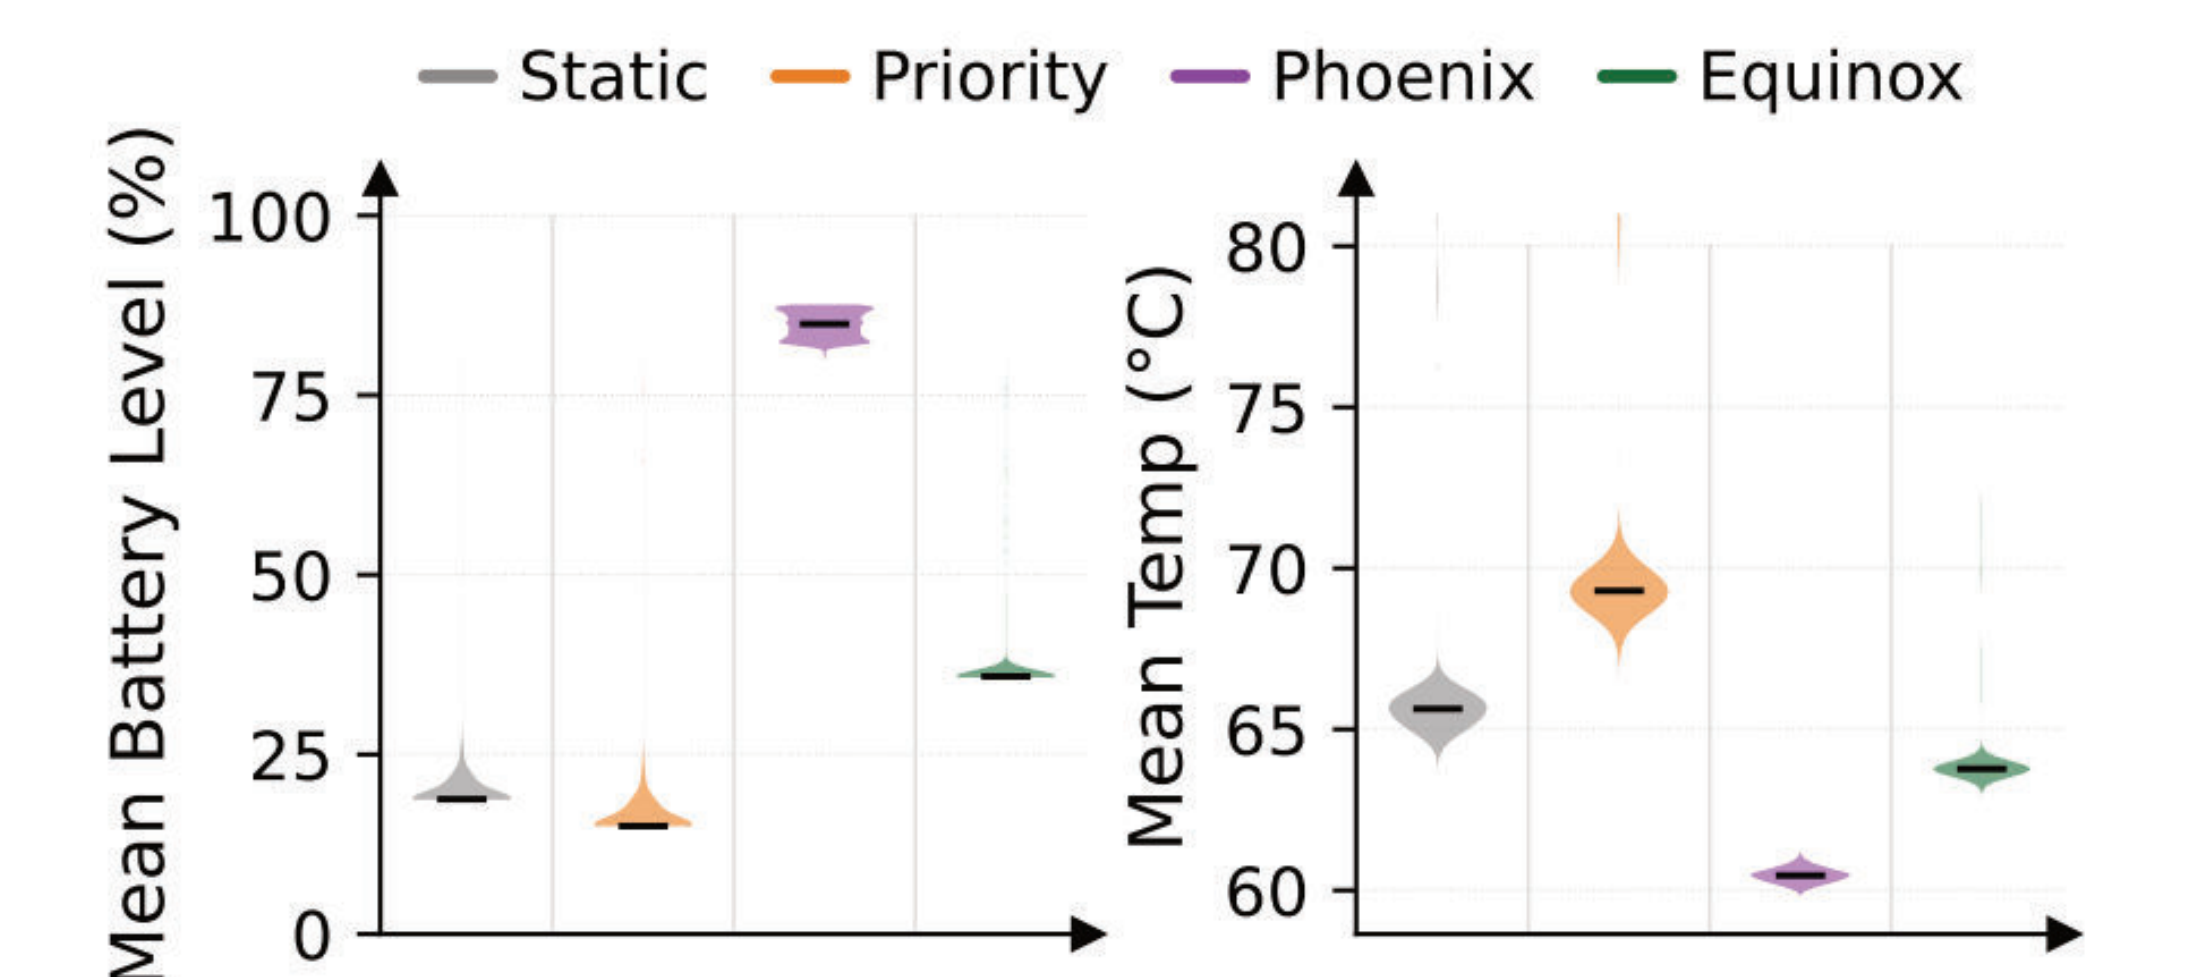
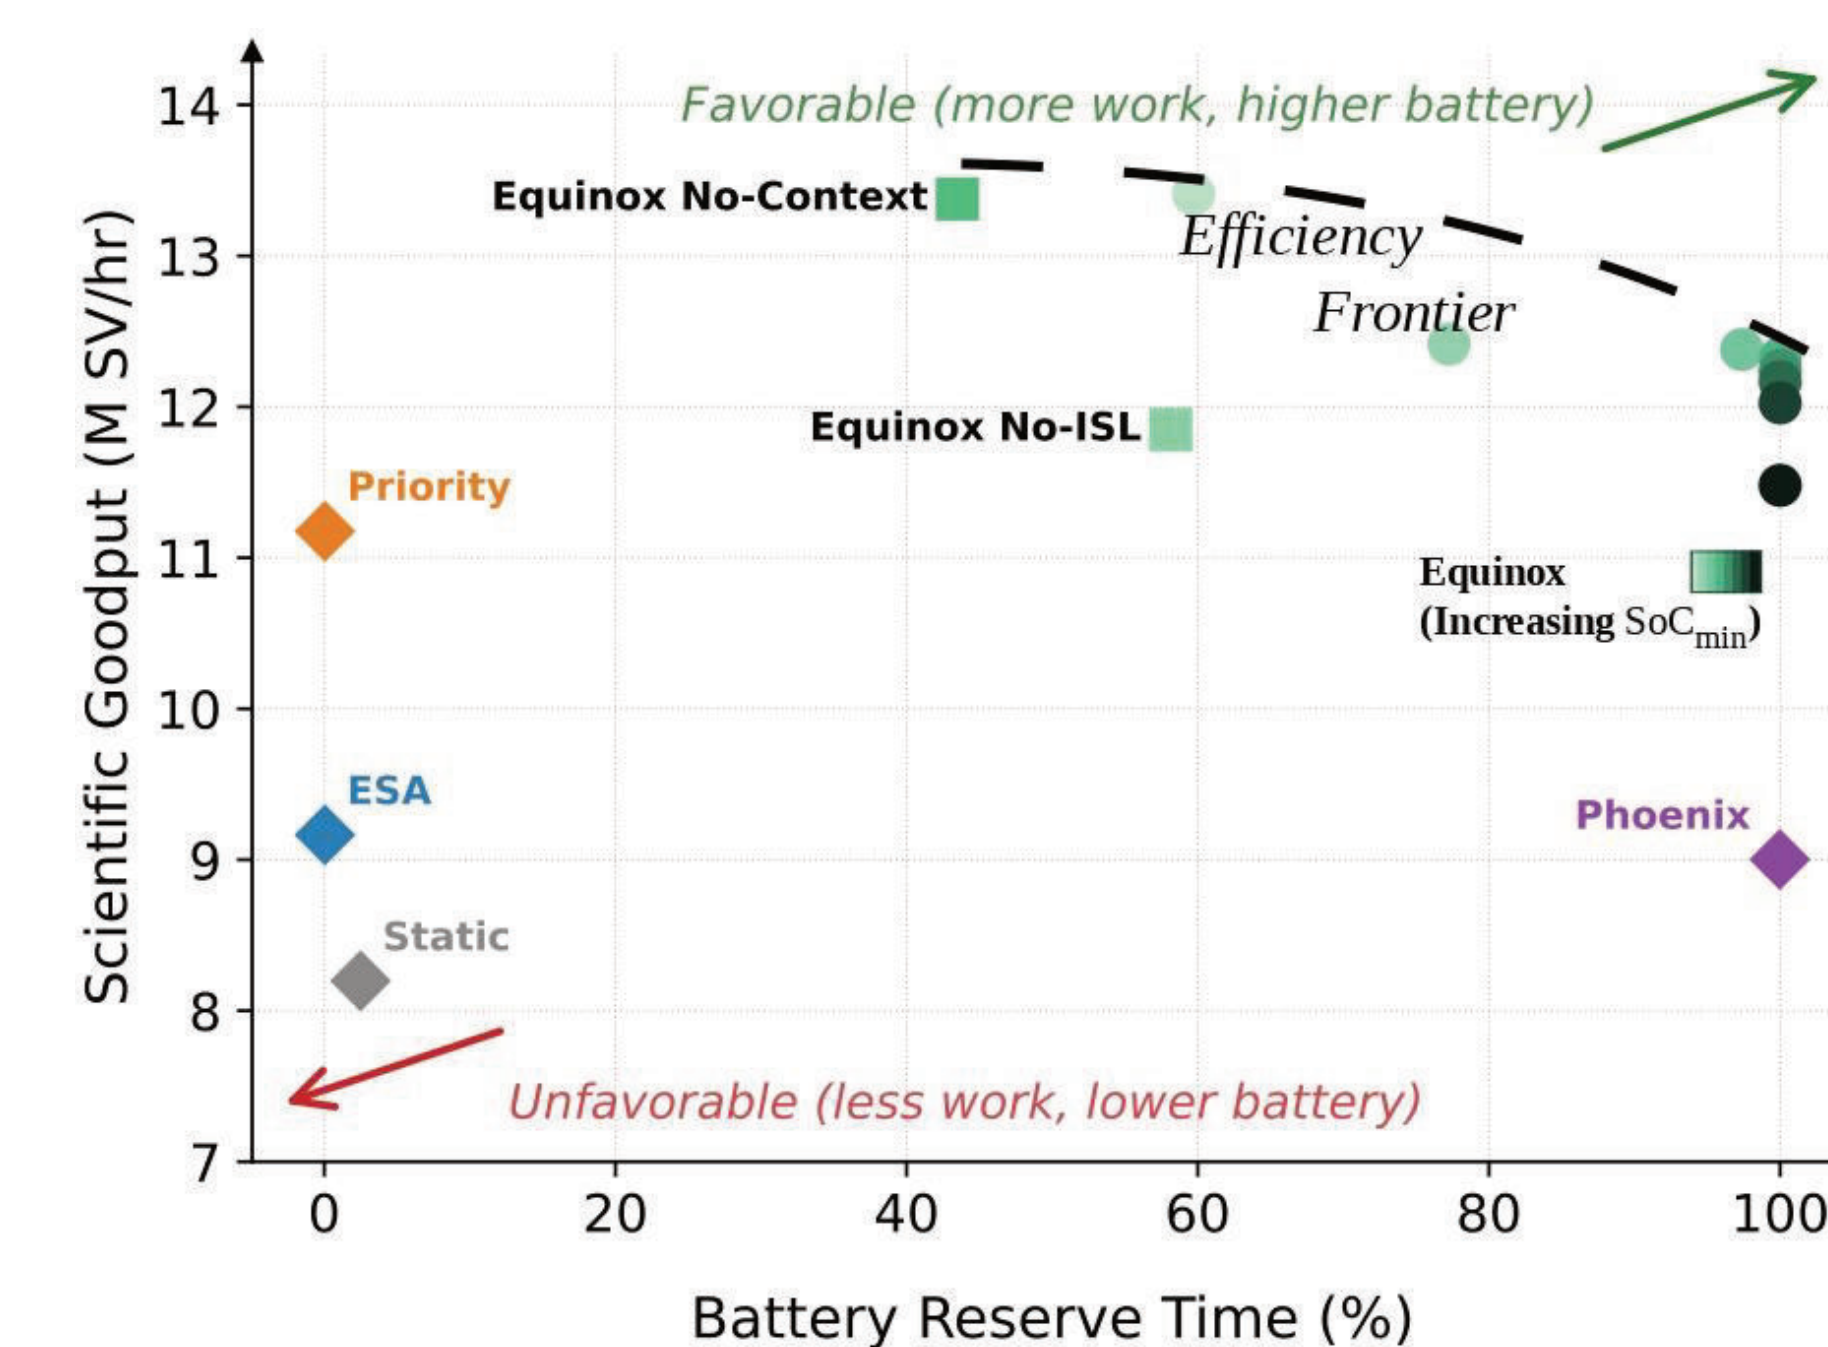
How prepared is each satellite to process a new wave of tasks? We measure *battery reserve time*, the amount of time a satellite spends with over 50% battery charge.

### (3) Science Load Balance

Work distribution can be deceptive. What if satellites are processing many, meaningless images? We integrate value into load balance using the Gini coefficient.

## Evaluation and Results

Equinox moves the goodput vs. reserve efficiency frontier outward compared to prior work.



Equinox naturally balances and utilizes available thermal headroom and battery, without explicit control

System	Scientific Value			Operational Health			Fairness
	Scientific Goodput (M SV/hr)	Throughput (K Images/hr)	Event Coverage (%)	Mean Battery Level (%)	Battery Reserve Time (%)	Brownout Risk (%)	Science Load Balance (%)
Static	8.20	346	44.8	20.4	7.4	64.4	91.2
Phoenix [18]	9.00	380	49.1	84.8	99.9	0.0	90.2
ESA [13]	9.16	390	50.1	17.0	5.9	85.1	91.0
Priority	11.18	387	62.2	17.0	6.0	85.1	91.4
EQUINOX (no-ISL)	11.85	417	65.7	35.7	58.0	29.1	94.4
EQUINOX (no-context)	13.38	564	73.1	26.7	43.0	37.9	94.7
EQUINOX (noisy-context)	13.36	508	73.6	36.6	59.3	27.9	96.4
EQUINOX	13.41	509	73.8	36.7	59.5	27.7	96.4

Equinox increases scientific goodput by up to 20% with 2.2x battery reserves over priority baselines, while detecting up to 28% more events and achieving 2x less load imbalance