Connecting the World from the Skies | ITU Space Sustainability Forum

Organizers | CST هيئة الاتصالات والفضاء والتقنية Communications, Space & Technology Commission | ITU

Mostly Asleep

How Much of an Orbital Data Centre Needs to Be Awake?

Camilo J. Rojas Milla

Juan A. Fraire, Fabio Patrone, Nour Badini, Mario Marchese

University of Genoa, Italy

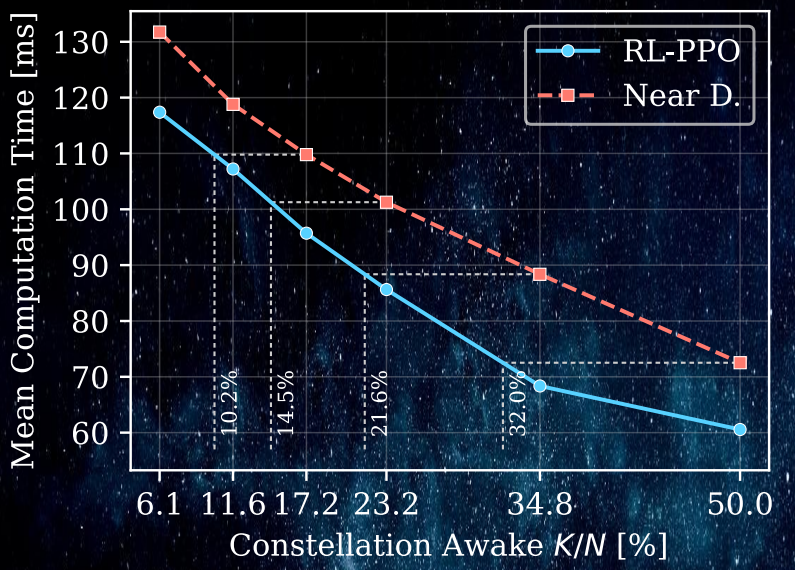
University of Santiago, Chile

Inria & INSA Lyon, CITI Laboratory, Lyon, France

Università di Genova | Inria | INSA Institut National des Sciences Appliquées Lyon | USACH Departamento de Ingeniería Informática

Orbital computing supports remote connectivity and high-performance computing. Selecting active payloads can deliver the same mean computation time with less capacity.

Equal mean computation time with fewer awake servers



198-satellite emulation; curves average loads of 45–270 tasks/s.
The 14.5% operating point is linearly interpolated on the RL-PPO curve.

≈38%

fewer awake servers

23.2% → 14.5%

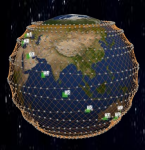
at 101.3 ms mean computation time

Near D.: greedy heuristic; activate the nearest K servers.
RL-PPO: learned activation policy.

Learning the awake set

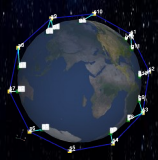
A shared actor–critic policy uses reinforcement learning and node state to select active servers.

We present two prior studies that examine key constellation resources, laying the groundwork for the learned policy.



1,080 satellites

Placement and network delay



10 satellites

Visibility and idle capacity



198 satellites

Learned server activation

From active capacity to power and heat

Fewer active payloads can reduce compute power and heat. The awake fraction informs energy and lifecycle models.

Energy savings require hardware measurements. Benefits depend on payload design, sleep power and wake-up costs.

Report awake fraction and latency as key sustainability metrics for orbital edge computing.

Contact:

camilo.rojas@edu.unige.it