OASIS: Optimized Adaptive System for Intelligent SLAM

SDSU San Diego State University

Alles Rebel
Computational Science Research Center

Nikil Dutt
University of California, Irvine

Bryan Donyanavard
San Diego State University

UCI

INTRODUCTION

Motivation: Autonomous cyber-physical systems need **robust** machine perception to **build accurate environmental and self models in real time**.

Perception pipeline: Feature extraction → depth estimation → sensor fusion → online spatial map construction.

SLAM Challenge: **Perceptual errors accumulate**, causing drift between estimated and true pose; unchecked drift risks collisions or system failure.

Key metrics:

- **Robustness**: consistently low mean ATE.
- **Predictability**: bounded max ATE (worst-case error).

Baseline system: **ORB-SLAM3** [1], a state-of-the-art visual inertial SLAM designed for desktop-class hardware.

Our contribution (**OASIS**): Controller that adapts perception online to maintain accuracy under resource limits and dynamic scenes.

Embedded deployment: Integrate **OASIS with ORB-SLAM3** on NVIDIA **Jetson Orin NX** (representative end device).

Realtime focus: Respect compute/latency constraints of the embedded hardware and demonstrate improved **robustness (↓ mean ATE)** and **predictability (↓ max ATE)** on embedded hardware.

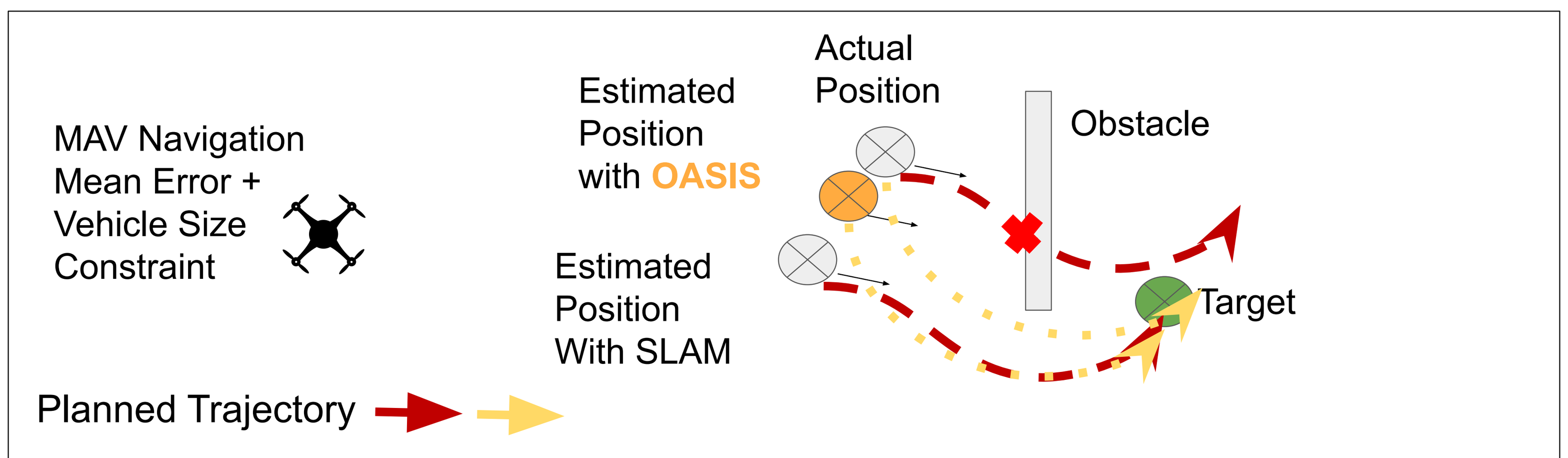

Figure 1. Planner outputs in a MAV navigation task using SLAM-only vs. OASIS-enhanced pose estimates. With larger localization error, the SLAM-only plan (red) intersects the obstacle; using OASIS (yellow) reduces error so the plan clears the obstacle and reaches the target. Both plans are applied from the same actual pose; discs indicate mean pose error plus a vehicle-size safety margin.

METHODS

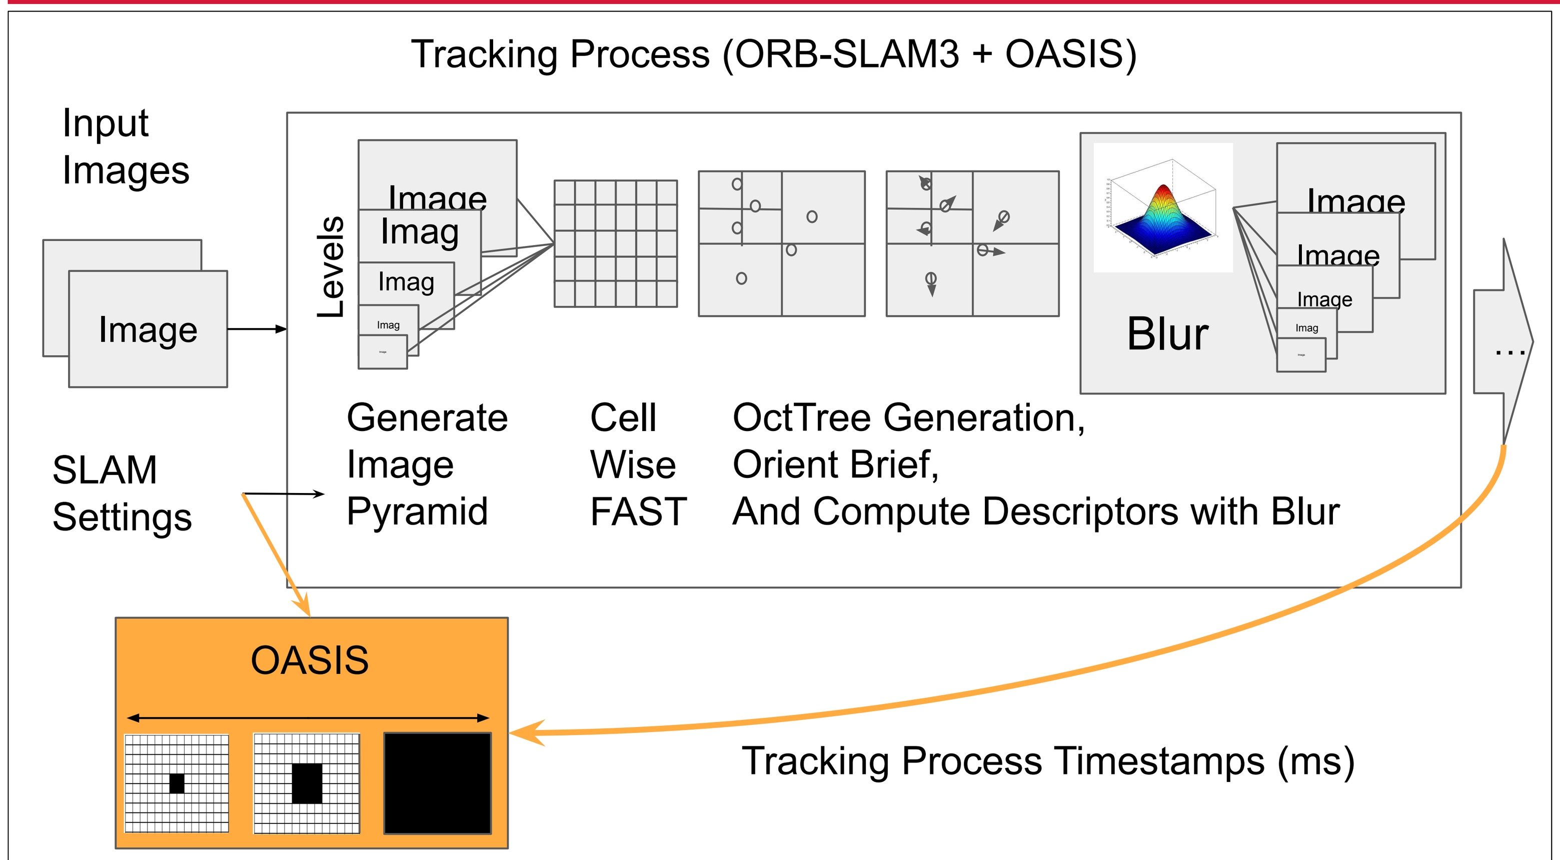

Figure 2. OASIS Block diagram showing SLAM configuration and timestatmp Inputs and binary Outputs onto the tracking process pipeline of ORB-SLAM3 [1]. Yellow indicates additions to the regular ORB-SLAM3 pipeline.

EXPERIMENTAL SETUP

Dataset utilized is the **ETH Zurich EuRoC MAV Dataset** [2]. A summary of the dataset:

- 11 **Indoor** Industrial Scenes
- Camera Calibration and Intrinsics
- 20 FPS, 752 x 480 **Stereo** Images
- 200 Hz **IMU Samples**
- Ground truth measurements (collected externally)

ORB-SLAM3 default configuration and parameters are used for all trials.

Specification	Details
Jetson Orin NX 16GB Developer Kit	
CPU	8-core Arm Cortex-A78AE
Memory	16 GB 128-bit LPDDR5
Power Limit	25W
Host Machine	
CPU	10-core Intel Core i7-6950X
Memory	64 GB 256-bit DDR4
Power Limit	140W

RESULTS

SLAM with OASIS has fewer spikes a overall smoother trace. Under realtime deadlines, OASIS (blue) reduces high-error bursts that dominate the baseline (red) due to frame drops, resulting in a predictable error profile.

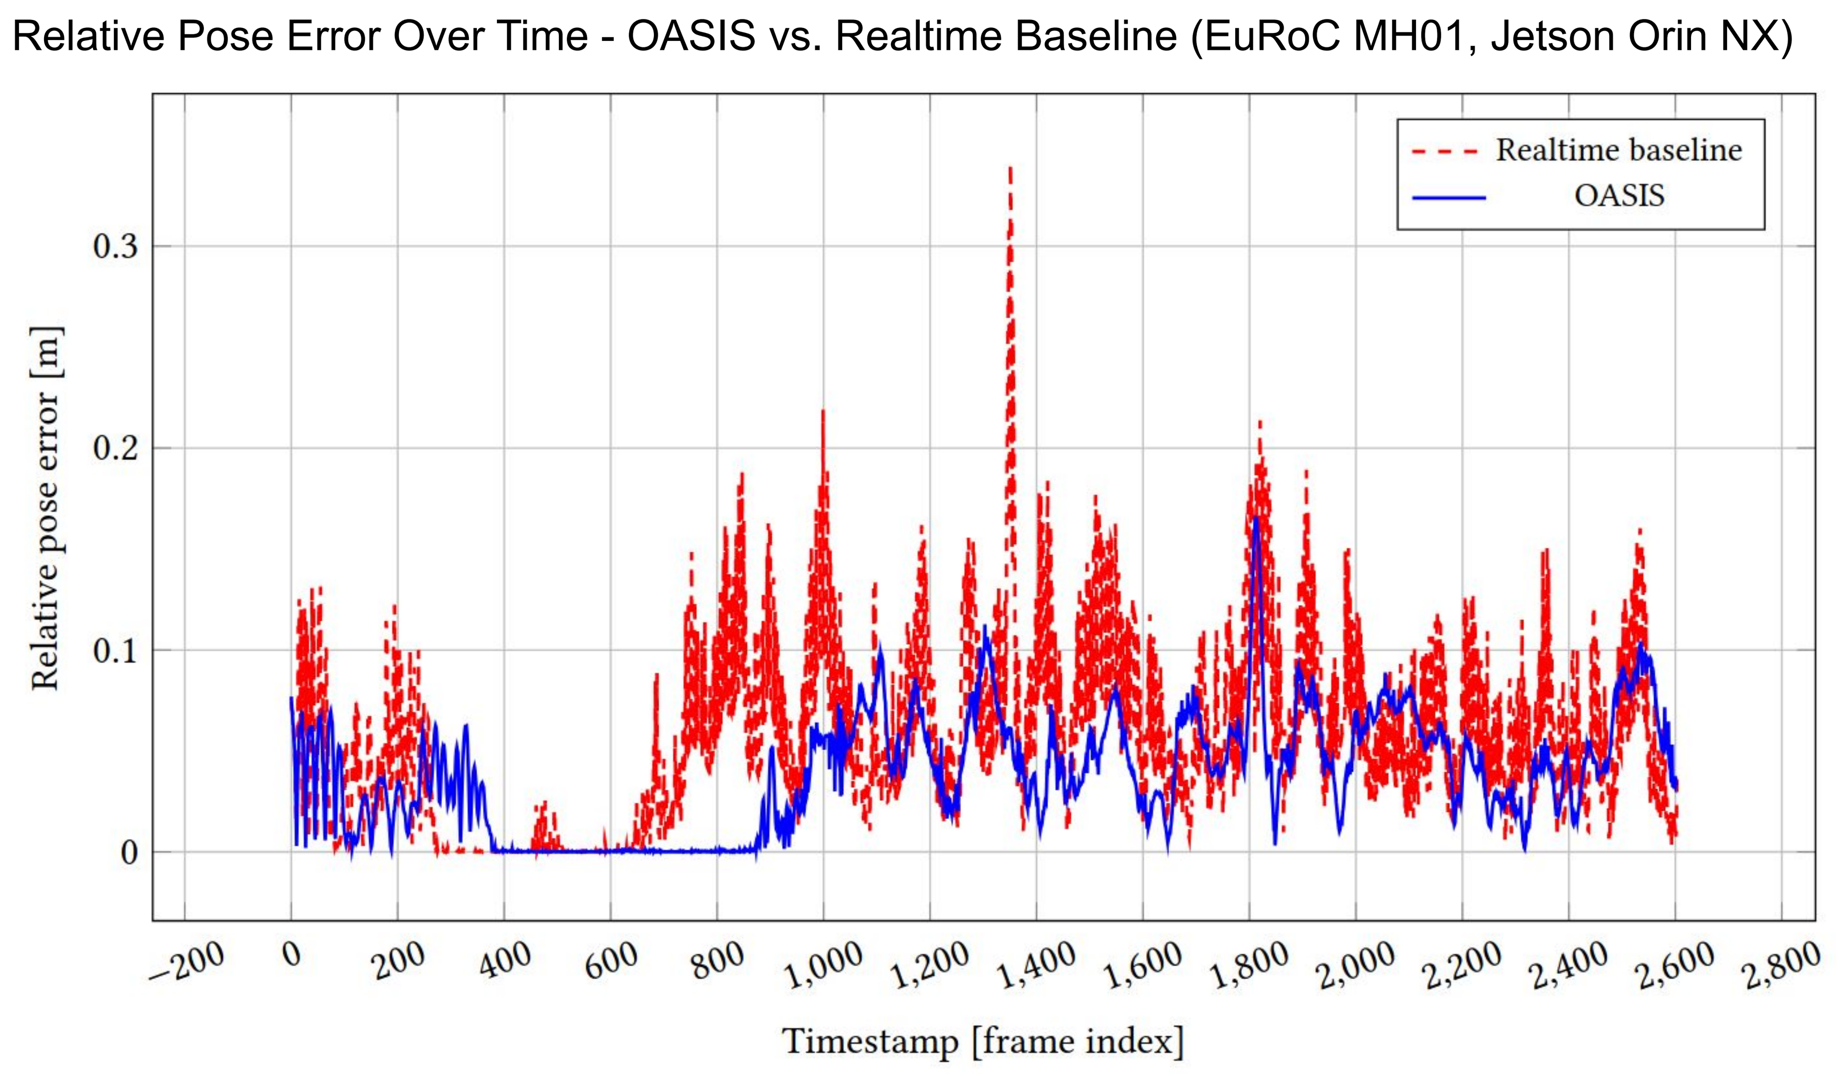

Figure 3. A Trial of EuRoC MH01 running on Nvidia Jetson Orin NX with ORB-SLAM3 (with deadlines), and another with OASIS enabled. Relative Pose Error in meters vs Timestamp of Frame. Red dotted trace indicates the ORB-SLAM3 realtime baseline. Blue solid indicates ORB-SLAM3 realtime baseline with OASIS enabled.

- **Realtime Baseline vs OASIS** — Jetson Orin NX
  - Mean ATE 0.0888 → 0.0547 m **(−38.4%)**
  - Max ATE 0.928 → 0.137 m **(−85.2%, ~6.75×)**
  - Dropped 11.5% → 0% (**baseline ~18.1 FPS, OASIS 20 FPS**)

- **Data-ready Baseline vs Realtime OASIS** — Intel Host (i7-6950X)
  - Mean ATE 0.0546 → 0.0538 m **(~−1.4%)**
  - Max ATE 0.1386 → 0.1312 m **(~−5.4%)**
  - Dropped **~0.0% vs 0.0%**

- **Realtime adaptive controllers vs OASIS** — Jetson Orin NX
  - PID-SLAM: mean ATE **−28.8%**, max ATE **−82.2%**
  - ω-SLAM: mean ATE **−30.4%**, max ATE **−80.0%**

- **Data-ready Fixed masks vs Realtime OASIS** — Jetson Orin NX
  - 4×4: mean ATE **−31.8%**, max ATE **−66.7%**
  - 6×6: mean ATE **−11.5%**, max ATE **−21.8%**

CONCLUSIONS

OASIS (ours): Online, adaptive **masking** that funnels compute to the most informative image regions so embedded ORB-SLAM3 **meets realtime deadlines**.

Mechanism: Approximates full-frame content to **reduce per-frame load—no SLAM reconfiguration** required.

Metrics: Mean ATE (robustness) and Max ATE (predictability / worst-case).

Jetson Orin NX results:

- **−38% mean ATE**, **−85% max ATE**, **0% frame drops** under realtime constraints.

Takeaway: OASIS delivers **predictable, realtime SLAM** on embedded hardware **without sacrificing accuracy**.

FUTURE WORK

Smarter, adaptive masking

- Replace a fixed center mask with **frame-adaptive region selection**.
- Use **context-aware / learning cues** (e.g., prioritize center at high speeds; emphasize edges/details at low speeds).

Real-world scale & robustness

- **Validate outdoors** and at larger spatial scales across **varied lighting/conditions**.

Hardware acceleration / Computational offloading

REFERENCES

[1] Carlos Campos, Richard Elvira, Juan J. Gómez Rodríguez, José M. M. Montiel, and Juan D. Tardós. 2021. ORB-SLAM3: An Accurate Open-Source Library for Visual, VisualśInertial, and Multimap SLAM. IEEE Transactions on Robotics 37, 6 (2021), 1874ś1890. doi:10.1109/TRO.2021.3075644

[2] Michael Burri, Janosch Nikolic, Pascal Gohl, Thomas Schneider, Joern Rehder, Sammy Omari, Markus W Achtelik, and Roland Siegwart. 2016. The EuRoC micro aerial vehicle datasets. The International Journal of Robotics Research 35, 10 (2016), 1157ś1163. arXiv:https://doi.org/10.1177/0278364915620033 doi:10.1177/0278364915620033

EMBEDDED SYSTEMS WEEK