# NAFGO Final Presentation

Three Tensors S.R.L., Romanian InSpace Engineering S.R.L.

The work is carried out under ESA NAVISP Element 1, devoted to the development of innovative PNT systems, technologies, algorithms and techniques (Ref. No.: NAVISP1-TN-3T-088-1)

23.06.2026

# Outline

# Table of Contents

# Overall description

Investigate innovative navigation algorithms exploiting FGO, both for a standalone GNSS unit and in combination with other sensors (INS), and for different receiver grades (including high- and low-quality modules).

The activity will benchmark:

- both the potential accuracy improvement and the computational cost of the proposed algorithms (in view of realtime implementations)
- design and implementation of a software-based FGO navigation (FGO-Nav) module
- integrated in a flexible proof-of-concept (PoC) testbed composed of at least two GNSS receivers (a low-end and a high-end receiver) and an INS

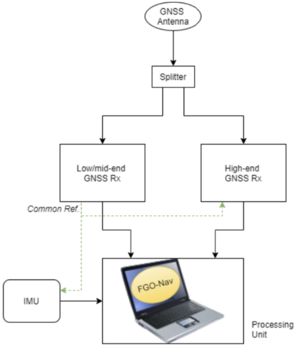

# Table of Contents

# What is a Factor Graph?

A Factor Graph (FG) is a bipartite graph $G = (V, F, E)$:

- $V$: hidden variable nodes (states to be estimated)
- $F$: factor nodes (probabilistic constraints from measurements / dynamics)
- $E$: edges connecting variables to factors

Under independent Gaussian noise, the MAP estimate of a state trajectory $\mathcal{X} = \{x_0, \dots, x_n\}$ solves the nonlinear least-squares

$$J(\mathcal{X}) = \|x_0 - \bar{x}_0\|^2_{\Sigma_0} + \sum_{k=1}^{n} \|h_k(x_k) - z_k\|^2_{\Sigma_R^k} + \sum_{k=1}^{n} \|f_k(x_{k-1}) - x_k\|^2_{\Sigma_Q^k},$$

with the Mahalanobis norm $\|r\|^2_\Sigma = r^\top \Sigma^{-1} r$.

- Decomposable cost ⇒ efficient Gauss–Newton / Levenberg–Marquardt solvers
- iSAM2 re-linearises only affected variables as new data arrive

- Marginalisation removes old states while preserving their information via a condensed prior ⇒ bounded memory & computation

# Tightly-coupled GNSS / IMU integration scheme

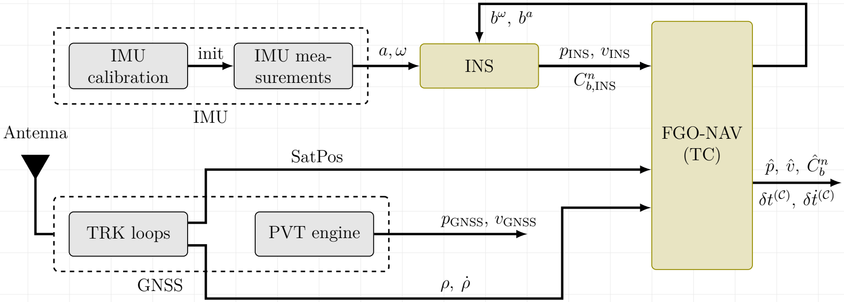

**Why tightly coupled?**

- Maximal information fusion from the propagation source (IMU)
- Maximal control over GNSS constellations, frequencies, models, residuals
- Inherent outlier damping and increased solution availability

**Cost:** convergence sensitivity, higher computational load, requires raw-measurement receivers.

# An illustrative example: Factor graph architecture (tightly coupled)

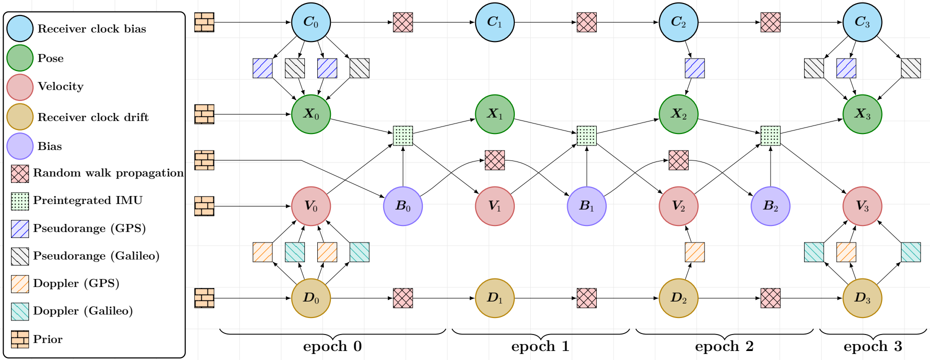

## State vector

$$x_k \triangleq \left( p_k^n,\ v_k^n,\ C_{b,k}^n,\ b_k^a,\ b_k^\omega,\ \delta t_k^{(\mathcal{C})},\ \delta \dot{t}_k^{(\mathcal{C})} \right)$$

## Variable nodes

- Position / velocity / attitude
- IMU biases (accel, gyro)
- Receiver clock bias / drift (per constellation)

# Table of Contents

# Software architecture

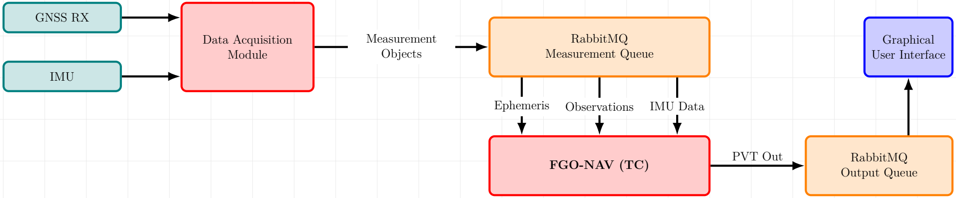

- partition into three main repos: Data acquisition (Sensors or Files), FGO and GUI
- design for pseudo real-time data acquisition and offline file parsing (asynchronous queuing using RabbitMQ)
- communication through messages (Ephemeris, nav, GUI, etc.)

# Data Acquisition block

- **Sensor interface and Drivers**: sampling rate, filtering, PPS synchronization and packet framing (via SPI)
- **Timestamping**: every incoming packet is stamped with both the host system clock, and where available, the sensor's own timestamp
- **Message publication (RabbitMQ)**: publish each merged, timestamped packet as a JSON message
- **Configuration**: all parameters (serial ports, baud rates, sensor configurations) live in json configuration files than can be updated in GUI.
- **Database logging**: all published data is stored in a MongoDB database.

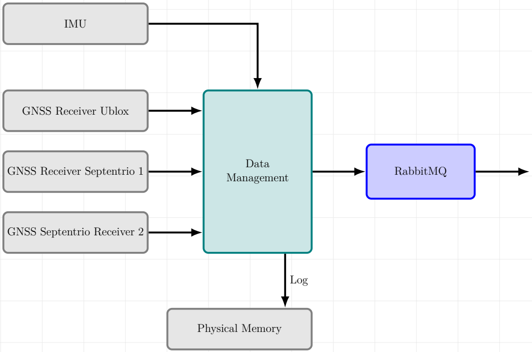

# Dataset Reader block

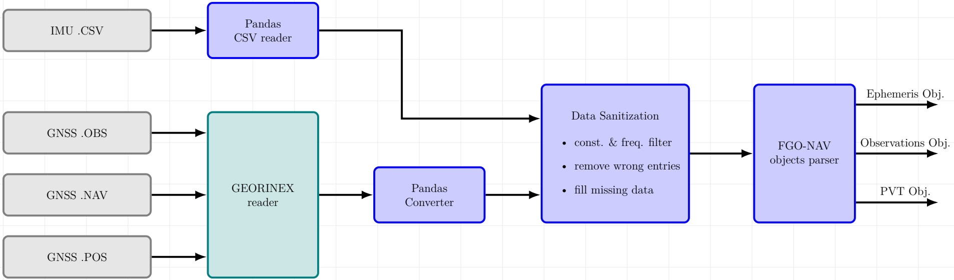

- **GEORINEX Reader** uses RINEX 3.x format to read satellite measurements and ephemeris messages, converting them into tabular data
- **Pandas CSV reader (IMU)** reads raw inertial meas. and sensor-specific covariances from CSVs
- **Pandas Converter** transforms the raw data into standardized Pandas DFs, for downstream processing.
- **Data Sanitization** performs validation and cleaning, filters by satellite constellation and frequency, removes corrupted entries, and handles missing values using interpolation or placeholder filling
- **FGO-NAV objects parser** Packages the sanitized data into structured messages used by the FGO-NAV system (ephemeris messages, GNSS observation epochs, and IMU measurements)

# FGO-NAV block

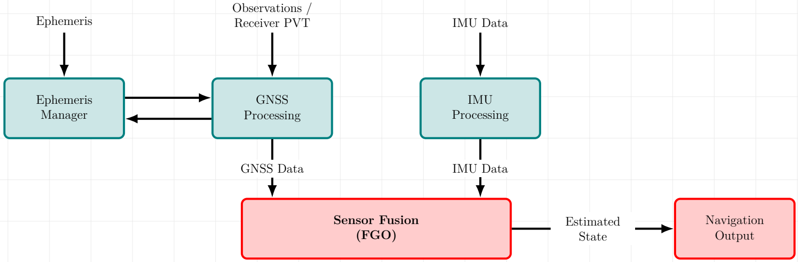

- **Ephemeris Manager** manages an Ephemeris Table storing up-to-date ephemeris, clock corrections, other broadcast navigation info
- **GNSS Processing** converts the ephemeris data into satellite pos (ECEF), performs iono-free combinations (dual freq.) and corrects for group delay (single freq.).
- **IMU Processing** performs lever arm compensation, time alignment, measurement batching for IMU preintegration.
- **Sensor Fusion Module** implements FGO, computes measurements residuals, Jacobians and noise covariances; performs outlier detection and filtering
- **Navigation Output Module** processes and outputs relevant information to the user or to the GUI process

# Algorithm design

The algorithm architectures considered follow the loosely-coupled and tightly-coupled architectures introduced in the DJF:

| Id | Type | Short description |
|---|---|---|
| FGO-NAV-GNSS | GNSS-only FGO | - GNSS pseudorange and Doppler measurements constrain navigation states<br>- Constant velocity motion model can be assumed between states. |
| FGO-NAV-LC | IMU/GNSS loosely coupled | - IMU (pre-)integrated via INS to FGO-NAV<br>- GNSS PVT outputs to FGO-NAV. |
| FGO-NAV-RTK-LC | IMU/GNSS-RTK loosely coupled | - IMU (pre-)integrated via INS to FGO-NAV<br>- GNSS RTK corrected PVT outputs to FGO-NAV. |
| FGO-NAV-PPP-LC | IMU/GNSS-PPP loosely coupled | - IMU (pre-)integrated via INS to FGO-NAV<br>- GNSS PPP corrected PVT outputs to FGO-NAV. |
| FGO-NAV-TC | IMU/GNSS tightly coupled | - IMU (pre-)integrated via INS to FGO-NAV<br>- GNSS pseudorange and Doppler to FGO-NAV. |

# Loosely-coupled GNSS and IMU factor graph architecture

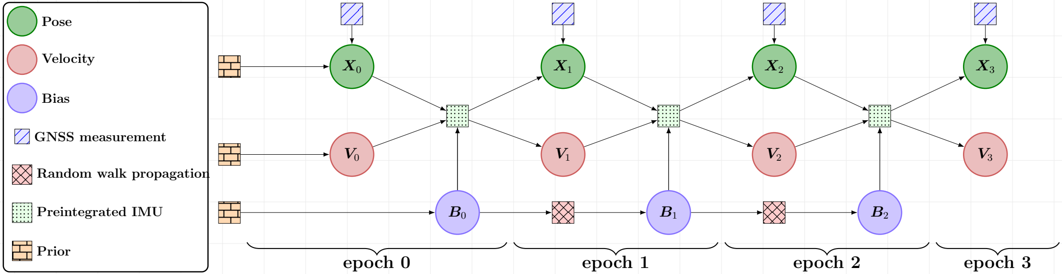

# Tightly coupled GNSS and IMU factor graph architecture

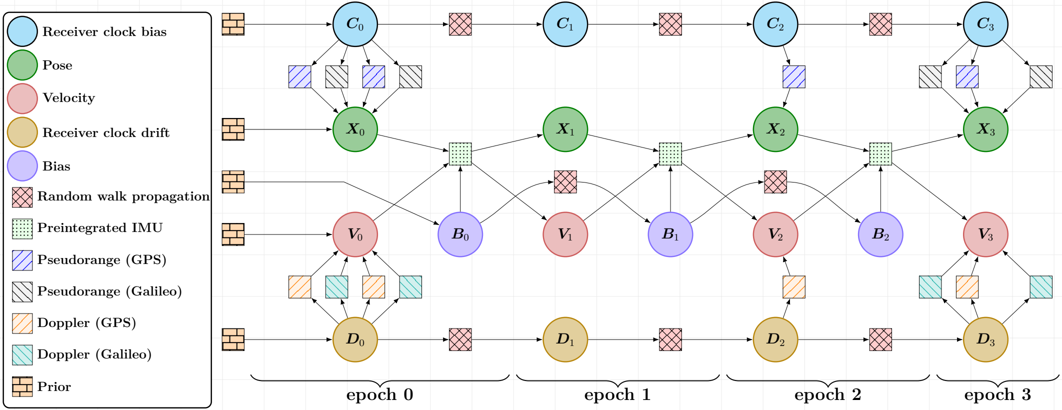

# New factors and processing rules in RTFGO-TC / RTFGO-LC

## New / extended factors

- **Heading prior factor** from velocity (`atan2(vE, vN)`), with velocity-dependent $\sigma$.
- **Yaw-freeze prior** when stationary.
- **ZUPT factor** with adaptive $\sigma$ (LC and TC, Sept and u-blox).

## Measurement-model rules

- Lever-arm application disabled at low speed (heading undefined).
- PVTData gated when reported covariance / arrival order is suspect (rejected as a factor instead of degrading the graph).

## Clock continuity

- Clock reinit after tunnel outage with protection window.
- Clock propagation extended to the full time window (TC clock-bias / drift factors).

## Initialization

- **WLMS bootstrap** now seeds the first key and feeds the RTFGO-TC heading prior the same way LC has always done — closed the TC-init gap.
- Fixed `pvt_data` / `clock_drift` initialization that was poisoning the early LC keys.

Tensors

# FGO-NAV Dataset Experiments

Tightly-Coupled Architecture:

- The most promising architecture: it improves availability and outlier rejection
- For UrbanNav dual-frequency measurements are sparse and DF performance is to be evaluated on another dataset and on the testbed
- Influenced by various configurations: use Doppler, elevation mask, SNR mask, PR weighting - needs tuning
- Needs a proper initialization and some time to observe heading

Loosely-Coupled Architecture:

- Struggles in UrbanNav because of low GNSS availability (WLMS rejection via chi-square test or residual thresholding)
- RTK-LC and PPP-LC also provide low GNSS availability and the IMU worsens sometimes the error as it accumulates bias quickly
- Too little control in FGO as we rely on the RX configs or RTKLIB

# Loosely Coupled Results

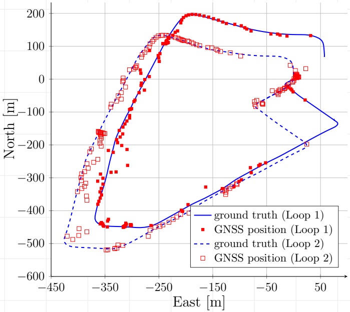

Ground truth and GNSS (Loops 1 and 2)

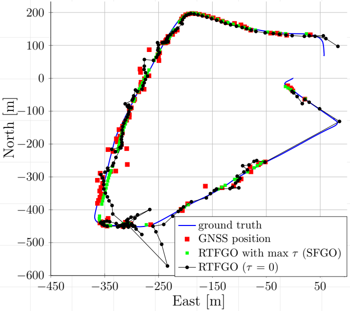

Estimated trajectories (Loop 1)

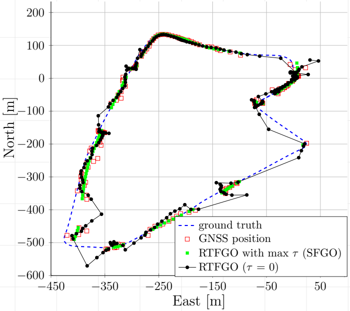

Estimated trajectories (Loop 2)

# Statistics and observations

Service availability:

$$A(\theta) = \frac{N_{\text{valid, error}\leq\theta}}{N_{\text{total}}} \times 100 \ [\%],$$

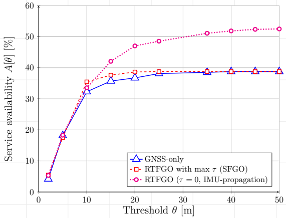

Marginalization lag impact on the RTFGO trajectory

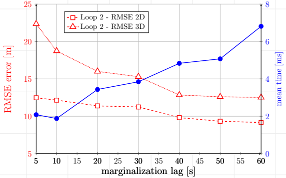

3D RMSE accuracy for Loop 1 and Loop 2

| Method | Configuration | Loop 1 RMSE_3D [m] | Loop 2 RMSE_3D [m] |
|---|---|---|---|
| GNSS-only | via RTKLIB | 29.20 | 13.24 |
| SFGO | [a] | — | 9.89 |
| RTFGO | Batch (max $\tau$) | 27.91 | 9.33 |
| RTFGO | RT ($\tau = 0$) | 33.31 | 11.83 |

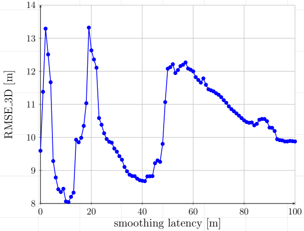

3D RMSE versus smoothing lag on the RTFGO trajectory, Loop 2.

---

[a] **11028340**.

# Statistics and observations (II)

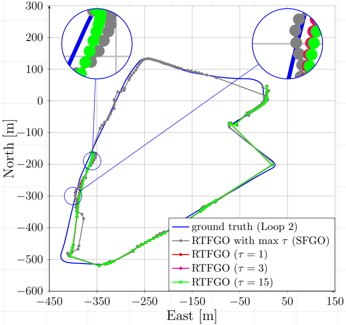

Smoothing latency impact on the RTFGO trajectory estimated for Loop 2

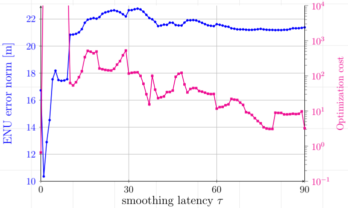

Position error evolution versus smoothing latency $\tau$ for the position at (02:43:35.000 GPST)

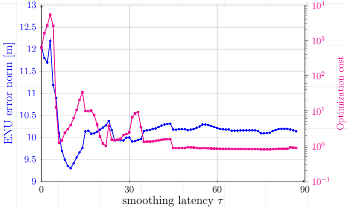

Position error evolution versus smoothing latency $\tau$ for the position at (02:43:48.000 GPST)

# RTK/PPP-LC

## RTK-LC

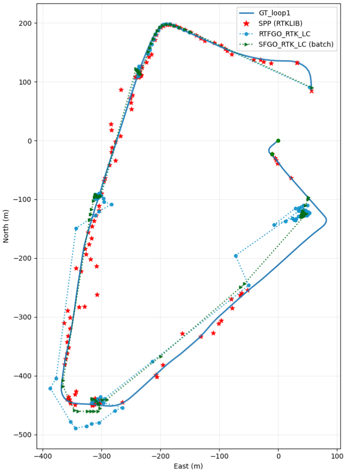

## PPP-LC

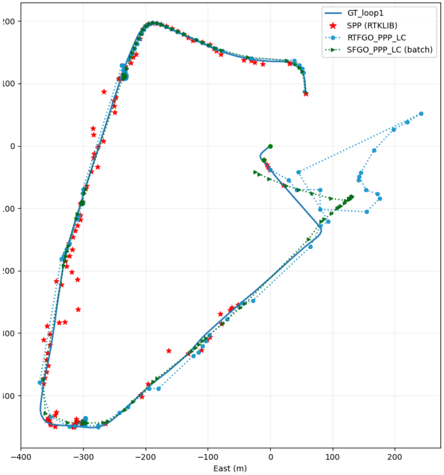

# Tightly-Coupled

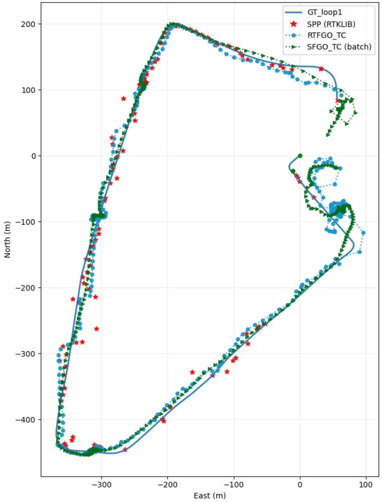

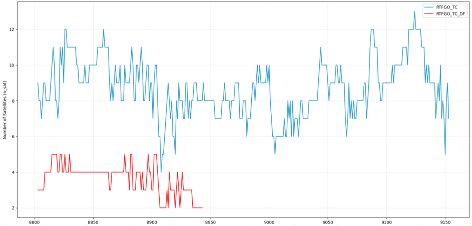

# GUI Dashboard

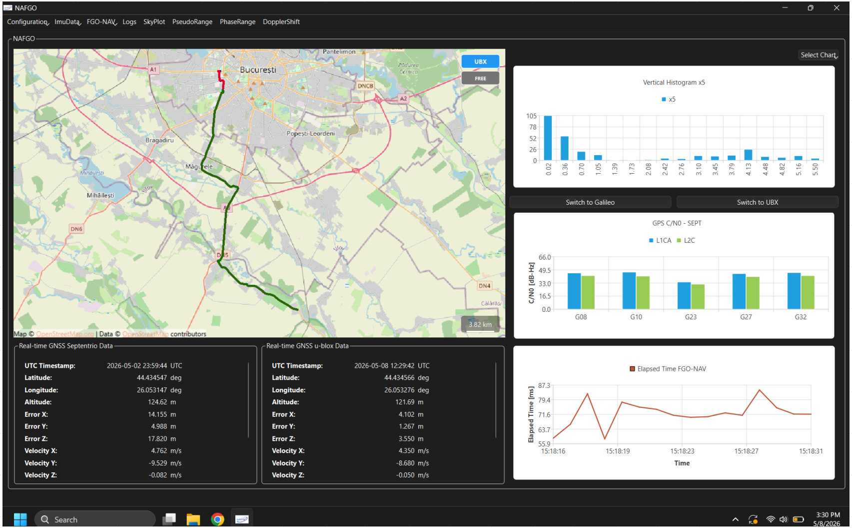

Real-time GNSS Septentrio Data

**UTC Timestamp:** 2026-05-02 23:59:44 UTC
**Latitude:** 44.434547 deg
**Longitude:** 26.053147 deg
**Altitude:** 124.62 m
**Error X:** 14.155 m
**Error Y:** 4.988 m
**Error Z:** 17.820 m
**Velocity X:** 4.762 m/s
**Velocity Y:** -9.529 m/s
**Velocity Z:** -0.082 m/s

Real-time GNSS u-blox Data

**UTC Timestamp:** 2026-05-08 12:29:42 UTC
**Latitude:** 44.434566 deg
**Longitude:** 26.053276 deg
**Altitude:** 121.69 m
**Error X:** 4.102 m
**Error Y:** 1.267 m
**Error Z:** 3.550 m
**Velocity X:** 4.350 m/s
**Velocity Y:** -8.680 m/s
**Velocity Z:** -0.050 m/s

# GUI – Map Widget

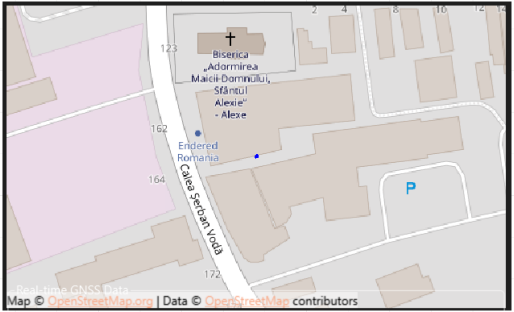

Map view widget

Receiver's location

# GUI – $C/N_0$ and SkyPlot

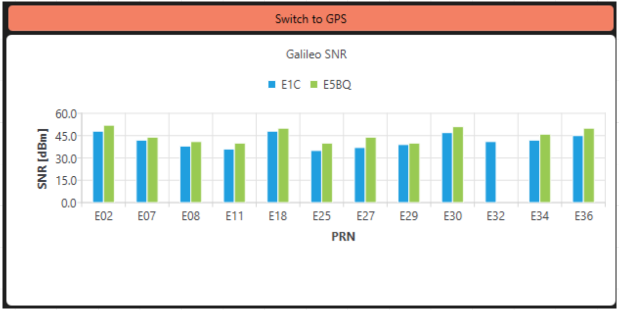

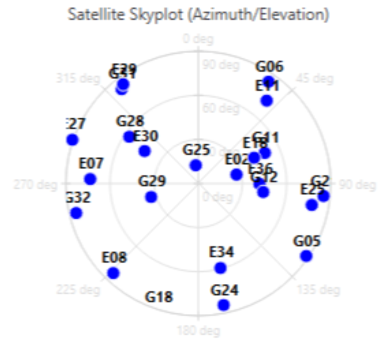

# Table of Contents

# Testbed design – block diagram

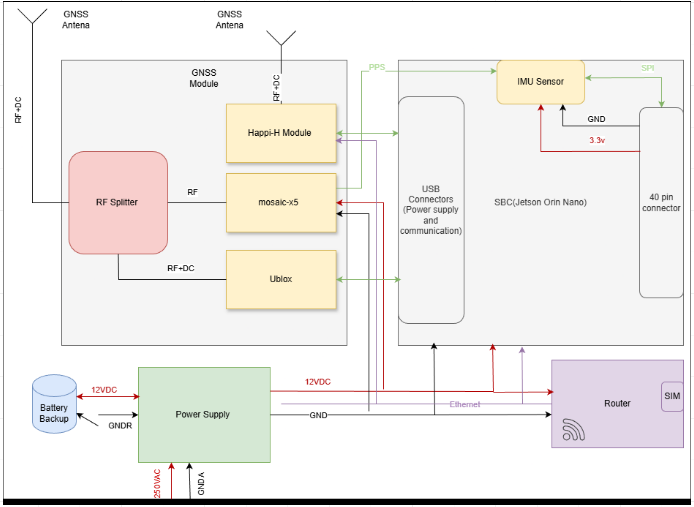

# Testbed - overview

- **Detailed list of components**:
  - a low-end dual-frequency/multi-constellation COTS GNSS receiver;
  - a high-end dual-frequency/multi-constellation COTS GNSS receiver with RTK capabilities;
  - a dedicated COTS GNSS base station to provide the RTK corrections;
  - a tactical IMU.

- **Power supply specifications and power budget summary**

- **Key characteristics**:
  - Time synchronization (PPS signal);
  - RF antenna splitter, WiFi router;
  - GNSS and GSM antennas;
  - Single-board computer, Power management system, GUI

# Assembled testbed

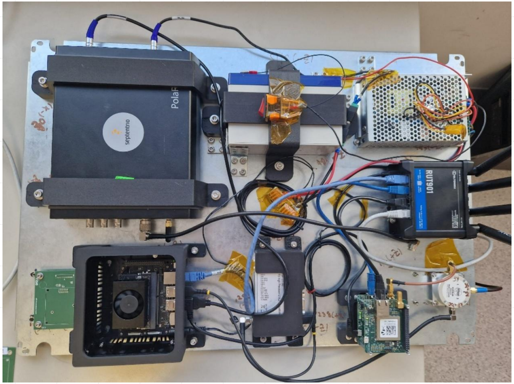

Testbed

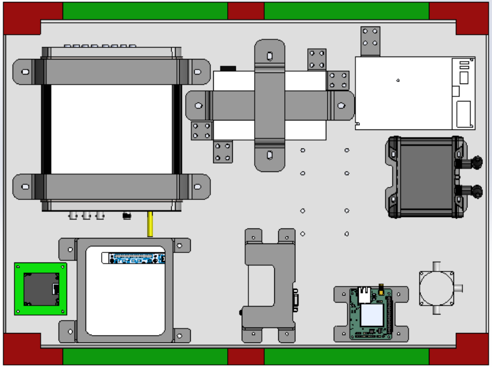

Wiring

# Measurement Campaign Classification

- **Mode:** most campaigns are dynamic driving measurements (we also have static ones).
- **Environment classes:** deep urban, light urban, sub-urban, and non-urban scenarios.
- **Road classes:** local streets, arterial roads, county roads, highways, and outside-city routes.
- **Configurations:** mainly C1, with one C2 campaign; later campaigns include improved synchronization and logging.
- **Weather and time:** daytime campaigns under clear or overcast conditions.
- **Data availability:** early days contain missing Mosaic PVT, Ublox raw data, or live FGO streams; later days include Chrony synchronization, raw Ublox data, and updated PVT timestamps.
- **Total duration: 25.51 hours** of measurements across all routes.

# Data collection campaign

The cumulative duration of all routes amounts to **25.51 hours.**

| DoY | Mode | Route / Area | Config | Dur. [min] | Env. | Weather | Time | Road Type | Notes |
|---|---|---|---|---|---|---|---|---|---|
| doy117 | dynamic | Outskirt Ferentari | C1 | 45.4 | light-urban | clear | daytime | local streets | No PVTData from Mosaic + No LIVE FGO data + No RawData from Ublox |
| doy118 | dynamic | Aviatorilor – Tineretului, through center | C1 | 120.1 | deep urban | clear | daytime | arterial roads | No PVTData from Mosaic + LIVE FGO DATA – UBLOX LC + No RawData from Ublox |
| doy119 | dynamic | A0, section South-East | C1 | 118.8 | light-urban | overcast | daytime | highways | No PVTData from Mosaic + LIVE FGO DATA – UBLOX TC + No RawData from Ublox |
| doy124 | dynamic | Timpuri Noi – Berceni – Calea Văcărești | C1 | 81 | deep urban | clear | daytime | local streets | No PVTData from Mosaic + No LIVE FGO data + No RawData from Ublox |
| doy125 | dynamic | Chiajna – Militari Residence | C1 | 114 | deep urban | clear | daytime | local streets | No PVTData from Mosaic + LIVE FGO DATA – UBLOX TC + No RawData from Ublox |
| doy126 | dynamic | Parcul Titan – Pantelimon – Obor – Unirii | C1 | 79.8 | deep urban | clear | daytime | arterial roads | Chrony clock sync + LIVE FGO DATA – UBLOX TC + No RawData from Ublox |
| doy127 | dynamic | Unirii, Obor, Victoriei tunnels – Grozăvești – Ghencea | C2 | 94.7 | deep urban | clear | daytime | arterial roads | Chrony clock sync + LIVE FGO DATA – UBLOX LC + RawData from Ublox |
| doy128 | dynamic | Comana | C1 | 113.3 | light-urban | clear | daytime | county roads | Chrony clock sync + LIVE FGO DATA – UBLOX LC + RawData from Ublox + New timestamp for PVTData on Ublox |
| doy129a | dynamic | Drumul Taberei Park | C1 | 29.6 | deep urban | overcast | daytime | arterial roads | Chrony clock sync + LIVE FGO DATA – UBLOX TC + RawData from Ublox + New timestamp for PVTData on Ublox |

# Data collection campaign (II)

The cumulative duration of all routes amounts to **25.51 hours.**

| DoY | Mode | Route / Area | Config | Dur. [min] | Env. | Weather | Time | Road Type | Notes |
|---|---|---|---|---|---|---|---|---|---|
| doy129b | dynamic | Herăstrău Park | C1 | 80.5 | deep urban | overcast | daytime | arterial roads | Chrony clock sync + LIVE FGO DATA – UBLOX TC + RawData from Ublox + New timestamp for PVTData on Ublox |
| doy131 | dynamic | Ghimpați | C1 | 59 | light-urban | clear | daytime | county roads | Chrony clock sync + LIVE FGO DATA – SEPTENTRIO LC + RawData from Ublox + New timestamp for PVTData on Ublox |
| doy132a | dynamic | Măgurele | C1 | 46.45 | sub-urban | clear | daytime | arterial roads | Chrony clock sync + LIVE FGO DATA – SEPTENTRIO LC + RawData from Ublox + New timestamp for PVTData on Ublox |
| doy132b | dynamic | 1 Decembrie – Jilava – Giurgiului Route | C1 | 75.73 | light-urban | clear | daytime | arterial roads | Chrony clock sync + LIVE FGO DATA – SEPTENTRIO TC (wlms) + RawData from Ublox + New timestamp for PVTData on Ublox |
| doy133 | dynamic | Bucharest – Bragadiru – Giurgiu Route | C1 | 163 | non urban | clear | daytime | outside the city | Chrony clock sync + LIVE FGO DATA – SEPTENTRIO TC + RawData from Ublox + New timestamp for PVTData on Ublox |
| doy134 | dynamic | Sălăjan – Ozana – Titan – Baicului – Colentina – Fabrica de Glucoză Street – Floreasca – City Center Route | C1 | 180.3 | deep urban | overcast | daytime | arterial roads | Chrony clock sync + Mosaic Clock steering + LIVE FGO DATA – UBLOX TC (wlms) + RawData from Ublox + New timestamp for PVTData on Ublox |
| doy135 | dynamic | Unirii – Splaiul Unirii – Pallady – Bucharest Ring Road – Berceni – Tineretului Route | C1 | 124.72 | light urban | clear | daytime | arterial roads | Chrony clock sync + Mosaic Clock steering + LIVE FGO DATA – UBLOX TC (wlms) + RawData from Ublox + New timestamp for PVTData on Ublox |

# Satellite map and estimated paths (Route 1 - DOY117)

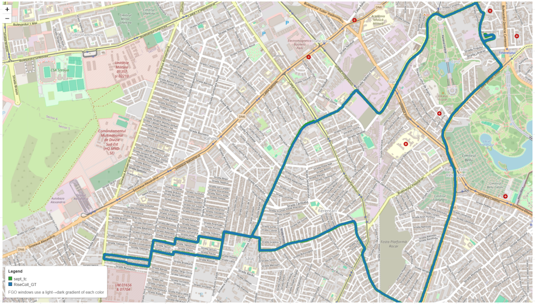

Route 1- satellite map and FGO path

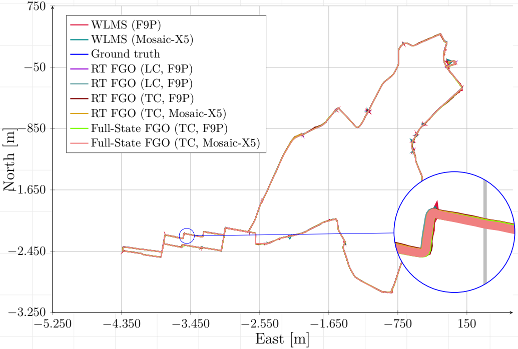

Route 1- trajectory estimation (ground truth, WLMS and FGO variations)

# Timeseries for horizontal RMSE (Route 1 - DOY117)

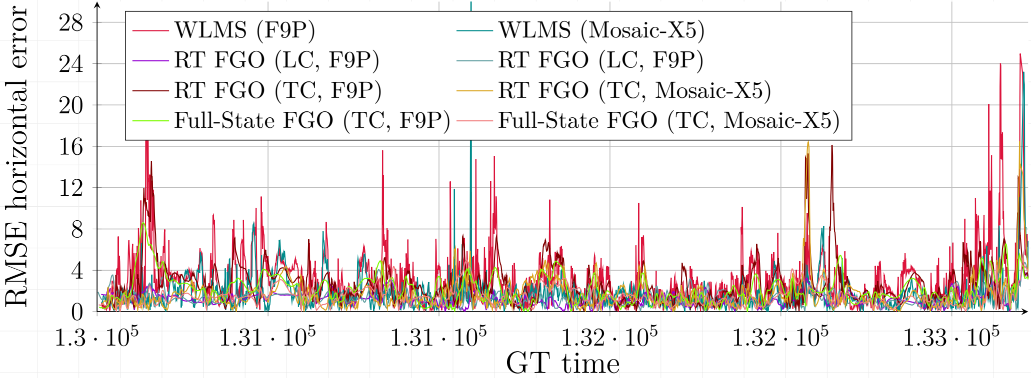

Tensors

# Timeseries for vertical RMSE (Route 1 - DOY117)

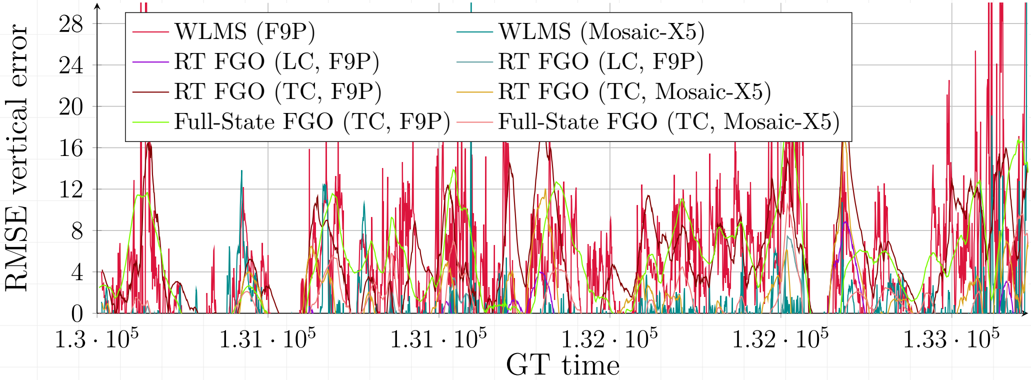

Tensors

# Timeseries for computation time (Route 1 - DOY117)

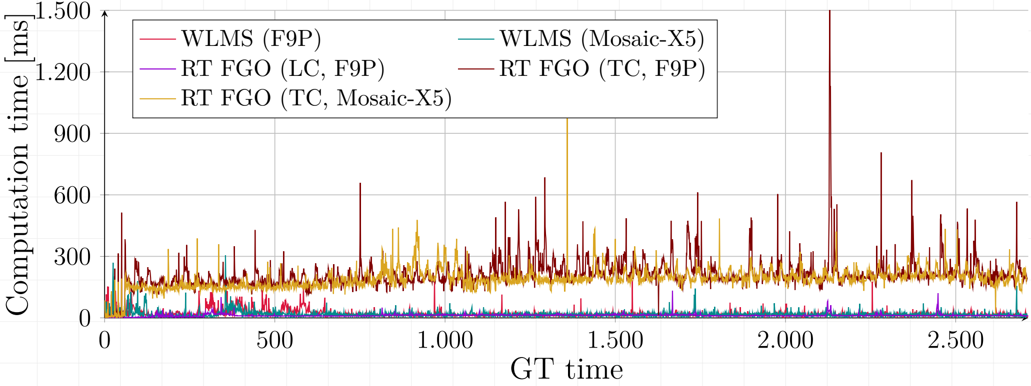

Tensors

# Histograms for horizontal RMSE (Route 1 - DOY117)

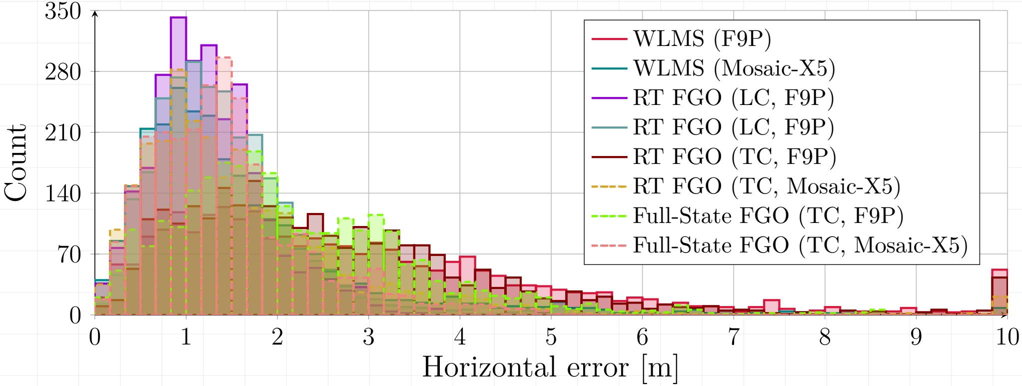

# Histograms for vertical RMSE (Route 1 - DOY117)

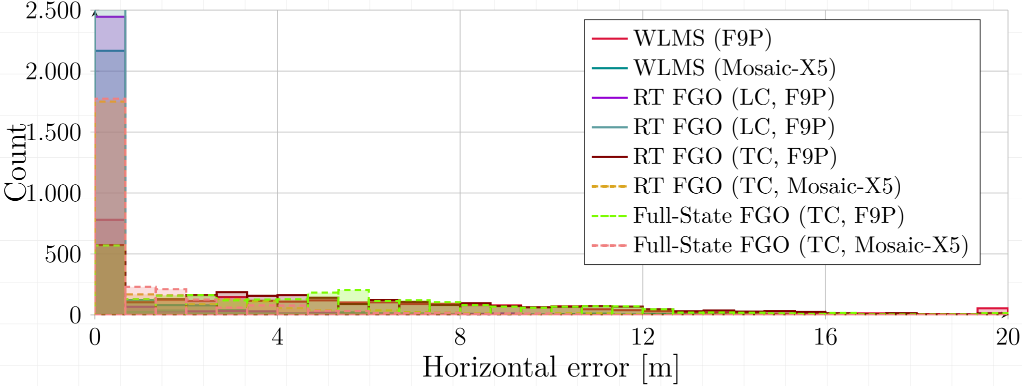

# Histograms for computation time (Route 1 - DOY117)

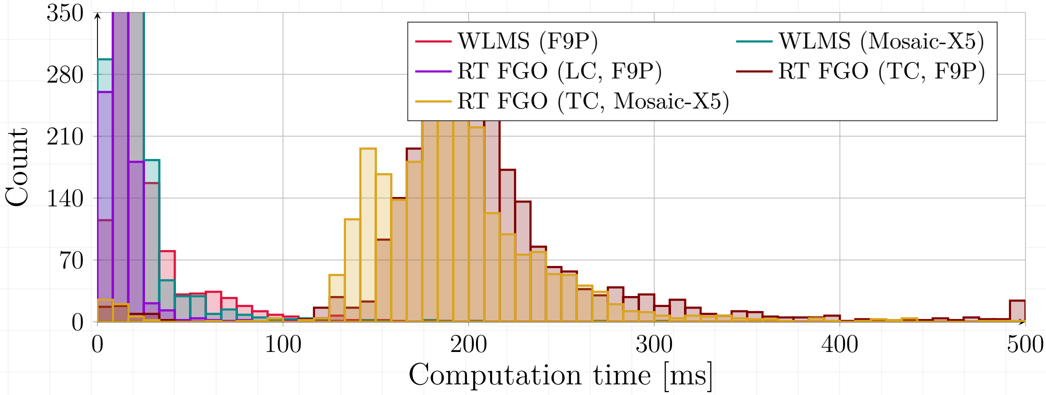

Tensors

# Overall statistics (Route 1 - DOY117)

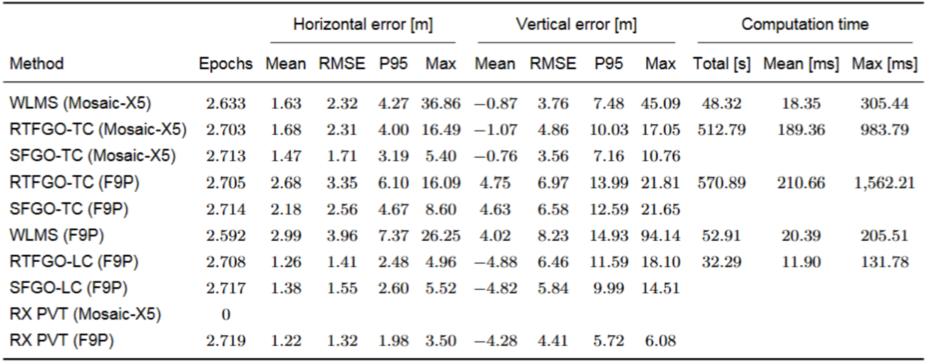

| Method | Epochs | Horizontal error [m] Mean | RMSE | P95 | Max | Vertical error [m] Mean | RMSE | P95 | Max | Computation time Total [s] | Mean [ms] | Max [ms] |
|---|---|---|---|---|---|---|---|---|---|---|---|---|
| WLMS (Mosaic-X5) | 2.633 | 1.63 | 2.32 | 4.27 | 36.86 | −0.87 | 3.76 | 7.48 | 45.09 | 48.32 | 18.35 | 305.44 |
| RTFGO-TC (Mosaic-X5) | 2.703 | 1.68 | 2.31 | 4.00 | 16.49 | −1.07 | 4.86 | 10.03 | 17.05 | 512.79 | 189.36 | 983.79 |
| SFGO-TC (Mosaic-X5) | 2.713 | 1.47 | 1.71 | 3.19 | 5.40 | −0.76 | 3.56 | 7.16 | 10.76 | | | |
| RTFGO-TC (F9P) | 2.705 | 2.68 | 3.35 | 6.10 | 16.09 | 4.75 | 6.97 | 13.99 | 21.81 | 570.89 | 210.66 | 1,562.21 |
| SFGO-TC (F9P) | 2.714 | 2.18 | 2.56 | 4.67 | 8.60 | 4.63 | 6.58 | 12.59 | 21.65 | | | |
| WLMS (F9P) | 2.592 | 2.99 | 3.96 | 7.37 | 26.25 | 4.02 | 8.23 | 14.93 | 94.14 | 52.91 | 20.39 | 205.51 |
| RTFGO-LC (F9P) | 2.708 | 1.26 | 1.41 | 2.48 | 4.96 | −4.88 | 6.46 | 11.59 | 18.10 | 32.29 | 11.90 | 131.78 |
| SFGO-LC (F9P) | 2.717 | 1.38 | 1.55 | 2.60 | 5.52 | −4.82 | 5.84 | 9.99 | 14.51 | | | |
| RX PVT (Mosaic-X5) | 0 | | | | | | | | | | | |
| RX PVT (F9P) | 2.719 | 1.22 | 1.32 | 1.98 | 3.50 | −4.28 | 4.41 | 5.72 | 6.08 | | | |

# Table of Contents

# Analysis of the results

**Comparisons between methods and computation time**:

- Raw PVT versus realtime LC FGO
- WLMS versus realtime TC FGO
- Realtime method used each run

**FGO-fused RTK / PPP (a selection)**

**Details for urban segments (FGO variants**

# Raw PVT vs Realtime LC FGO — per-day horizontal RMSE (m)

| DOY | Best PVT | RMSE [m] | Best RTFGO-LC | RMSE [m] |
|---|---|---|---|---|
| 117 | RX PVT (F9P) | **1.323** | RTFGO-LC (F9P) | 1.407 |
| 118 | RX PVT (F9P) | **3.482** | RTFGO-LC (F9P) | 4.174 |
| 119 | RX PVT (F9P) | 4.892 | RTFGO-LC (F9P) | **2.070** |
| 124 | RX PVT (F9P) | **2.847** | RTFGO-LC (F9P) | 3.759 |
| 125 | RX PVT (F9P) | **1.867** | RTFGO-LC (F9P) | 2.034 |
| 126 | RX PVT (F9P) | **1.771** | RTFGO-LC (F9P) | 3.062 |
| 127 | RX PVT (Mosaic-X5) | **2.408** | RTFGO-LC (F9P) | 4.076 |
| 128 | RX PVT (F9P) | **1.161** | RTFGO-LC (F9P) | 1.190 |
| 129 p1 | RX PVT (F9P) | **1.270** | RTFGO-LC (F9P) | 1.344 |
| 129 p2 | RX PVT (F9P) | 2.163 | RTFGO-LC (F9P) | **2.023** |
| 131 | RX PVT (Mosaic-X5) | 1.709 | RTFGO-LC (Mosaic-X5) | **1.433** |
| 132 p1 | RX PVT (F9P) | **2.006** | RTFGO-LC (F9P) | 2.114 |
| 132 p2 | RX PVT (F9P) | **1.924** | RTFGO-LC (F9P) | 2.100 |
| 133 | RX PVT (F9P) | 1.555 | RTFGO-LC (F9P) | **1.538** |
| 134 | RX PVT (F9P) | **2.610** | RTFGO-LC (F9P) | 2.679 |
| 135 | RX PVT (Mosaic-X5) | **2.140** | RTFGO-LC (F9P) | 2.656 |

# WLMS vs Realtime TC FGO — per-day horizontal RMSE (m)

| DOY | Best WLMS | RMSE [m] | Best RTFGO-TC | RMSE [m] |
|---|---|---|---|---|
| 117 | WLMS (Mosaic-X5) | 2.318 | RTFGO-TC (Mosaic-X5) | **2.308** |
| 118 | WLMS (Mosaic-X5) | 5.982 | RTFGO-TC (Mosaic-X5) | **4.708** |
| 119 | WLMS (Mosaic-X5) | 2.543 | RTFGO-TC (Mosaic-X5) | **2.245** |
| 124 | WLMS (F9P) | 7.514 | RTFGO-TC (F9P) | **6.462** |
| 125 | WLMS (F9P) | 4.605 | RTFGO-TC (F9P) | **4.302** |
| 126 | WLMS (Mosaic-X5) | **2.976** | RTFGO-TC (Mosaic-X5) | 3.300 |
| 127 | WLMS (Mosaic-X5) | **3.498** | RTFGO-TC (Mosaic-X5) | 3.949 |
| 128 | WLMS (Mosaic-X5) | **1.758** | RTFGO-TC (Mosaic-X5) | 2.020 |
| 129 p1 | WLMS (Mosaic-X5) | 2.627 | RTFGO-TC (Mosaic-X5) | **2.063** |
| 129 p2 | WLMS (Mosaic-X5) | 2.728 | RTFGO-TC (Mosaic-X5) | **1.835** |
| 131 | WLMS (Mosaic-X5) | **1.940** | RTFGO-TC (F9P) | 3.082 |
| 132 p1 | WLMS (F9P) | 7.234 | RTFGO-TC (F9P) | **6.584** |
| 132 p2 | WLMS (Mosaic-X5) | **2.813** | RTFGO-TC (Mosaic-X5) | 3.801 |
| 133 | WLMS (Mosaic-X5) | **1.627** | RTFGO-TC (Mosaic-X5) | 1.886 |
| 134 | WLMS (Mosaic-X5) | 5.056 | RTFGO-TC (Mosaic-X5) | **4.706** |
| 135 | WLMS (Mosaic-X5) | 2.683 | RTFGO-TC (Mosaic-X5) | **2.210** |

Per-day winner in bold. RTFGO-TC wins on 10/16 tests; WLMS wins on 6/16. Largest RTFGO-TC gain: doy118 (5.98 m → 4.71 m, −1.27 m); largest WLMS lead: doy131 (1.94 m vs 3.08 m, −1.14 m).

# Realtime on-vehicle runs — per-day result and compute cost

| DOY | Realtime variant | #epochs | Horiz RMSE [m] | Mean elapsed [ms] |
|---|---|---|---|---|
| 117 | — | — | — | — |
| 118 | — | — | — | — |
| 119 | RTFGO-TC (F9P) | 6418 | 31.678 | 395.9 |
| 124 | — | — | — | — |
| 125 | RTFGO-TC (F9P) | 2056 | 87.715 | 396.7 |
| 126 | RTFGO-TC (F9P) | 3481 | 30.878 | 440.6 |
| 127 | RTFGO-LC (F9P) | 1655 | **2.109** | 72.5 |
| 128 | RTFGO-LC (F9P) | 5945 | **1.103** | 81.0 |
| 129 p1 | RTFGO-TC (F9P) | 1645 | 3.517 | 443.3 |
| 129 p2 | RTFGO-TC (F9P) | 4253 | 3.353 | 381.8 |
| 131 | RTFGO-LC (Mosaic-X5) | 185 | 2.393 | **11.0** |
| 132 p1 | RTFGO-LC (Mosaic-X5) | 1234 | 2.430 | 51.5 |
| 132 p2 | WLMS (Mosaic-X5) | 3419 | 2.960 | 48.1 |
| 133 | RTFGO-TC (Mosaic-X5) | 2752 | 2.581 | 691.7 |
| 134 | WLMS (F9P) | 8508 | 8.171 | 59.7 |
| 135 | WLMS (F9P) | 7001 | 4.397 | 50.8 |

Mean elapsed time = average per-epoch FGO compute. LC variants $\leq$ 100 ms/epoch; F9P RTFGO-TC sits around 400 ms, Mosaic-X5 RTFGO-TC ~690 ms. RTFGO-TC realtime runs on F9P diverged on days 119/125/126 (RMSE $\gtrsim$ 30 m) — likely a windowing / reorder-buffer pathology.

# Raw vs FGO-fused RTK / PPP — doy117 & doy118 (horiz RMSE, m)

**doy117**

| Solution | Horiz RMSE [m] |
|---|---|
| RX PVT (F9P) | **1.323** |
| RTKLIB RTK (no IMU) | 1.389 |
| RTFGO-LC + RTK (Mosaic-X5) | 1.351 |
| CSRS-PPP (no IMU) | 1.571 |
| RTFGO-LC + PPP (Mosaic-X5) | 1.430 |

**doy118**

| Solution | Horiz RMSE [m] |
|---|---|
| RX PVT (F9P) | **3.482** |
| RTKLIB RTK (no IMU) | 3.992 |
| RTFGO-LC + RTK (Mosaic-X5) | 6.680 |
| CSRS-PPP (no IMU) | 5.080 |
| RTFGO-LC + PPP (Mosaic-X5) | 6.229 |

- Raw RTK / PPP available only on doy117 and doy118 (post-processed from Mosaic-X5 RINEX, no IMU).
- On doy117 the FGO-fused traces match raw RTK and slightly beat raw PPP.
- On doy118 both FGO-fused traces are markedly *worse* than the raw solutions — worth a dedicated diagnostic before the review.

# doy117 — all solutions ranked by horizontal RMSE (m)

| # | Family | Solution | Receiver | Horiz RMSE [m] |
|---|---|---|---|---|
| 1 | Raw PVT | RX PVT | F9P | **1.323** |
| 2 | RTFGO-LC + RTK | RTFGO-LC-RTK | Mosaic-X5 | 1.351 |
| 3 | RTK (no IMU) | RTKLIB kinematic | Mosaic-X5 | 1.389 |
| 4 | RTFGO-LC | RTFGO-LC | F9P | 1.407 |
| 5 | RTFGO-LC + PPP | RTFGO-LC-PPP | Mosaic-X5 | 1.430 |
| 6 | SFGO-LC (batch) | SFGO-LC | F9P | 1.554 |
| 7 | PPP (no IMU) | CSRS-PPP kinematic | Mosaic-X5 | 1.571 |
| 8 | SFGO-TC (batch) | SFGO-TC | Mosaic-X5 | 1.713 |
| 9 | RTFGO-TC | RTFGO-TC | Mosaic-X5 | 2.308 |
| 10 | WLMS | WLMS | Mosaic-X5 | 2.318 |
| 11 | SFGO-TC (batch) | SFGO-TC | F9P | 2.560 |
| 12 | RTFGO-TC | RTFGO-TC | F9P | 3.352 |
| 13 | WLMS | WLMS | F9P | 3.964 |

- The top cluster (1.3–1.6 m) is dominated by LC fusion and the externally-aided solutions (RTK, PPP) — the loosely-coupled pipeline tracks the receiver PVT closely and inherits its quality.
- TC solutions (raw-observable FGO and WLMS) sit a clear step below, and every F9P TC entry trails its Mosaic-X5 counterpart, isolating the receiver-grade effect from the algorithm choice.

# doy117 — urban segments, GNSS vs FGO trajectory

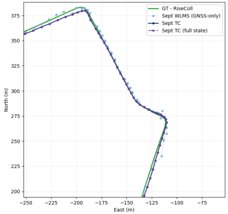

Calea Şerban Vodă — tight urban canyon

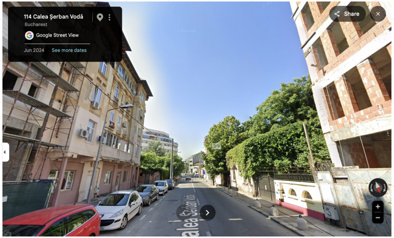

Trajectory at same location: TC tracks GT, WLMS noisy

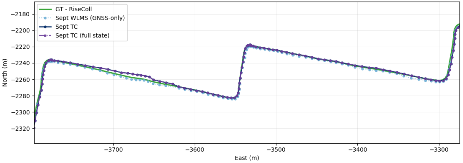

Str. Brătila — open suburban, low-rise

All solutions converge;TC fails at the heading change

# doy117 — heading and gyro-bias estimates

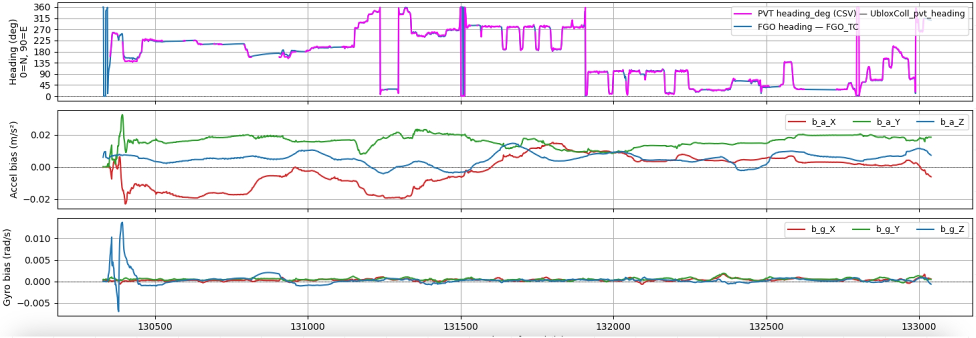

- Top panel: vehicle heading (FGO yaw vs PVT-derived heading from `atan2(vE, vN)`).
- Bottom panel: estimated gyro biases (x/y/z) over the run.
- Grey bands mark low-speed epochs where the velocity-derived heading is unreliable.

# doy117 — marginalization window sweep (RTFGO-TC, Mosaic-X5)

| Window [s] | Horiz RMSE [m] | Mean compute [ms / epoch] |
|---:|---:|---:|
| 10 | 2.524 | 36.2 |
| 15 | 2.311 | 50.0 |
| 20 | 2.307 | 65.8 |
| 30 | **2.306** | 118.2 |
| 60 | 2.308 | 186.3 |
| 120 | 2.307 | 357.0 |
| 300 | 2.307 | 1418.9 |

- Same run (SeptColl, RTFGO-TC) with only the marginalization horizon swept.
- Going from 10 s to 15 s yields the only real accuracy gain (2.52 m → 2.31 m, −0.21 m); beyond 15 s the horizontal RMSE plateaus around 2.31 m.
- Compute cost, however, keeps growing almost linearly: a 30× longer window costs ∼40× more compute (36 ms → 1.42 s/epoch) for no measurable accuracy improvement.
- The knee sits at **w** ≈ **15–20 s**: longer windows simply buy more latency.

# doy118 — Calea Victoriei (urban canyon)

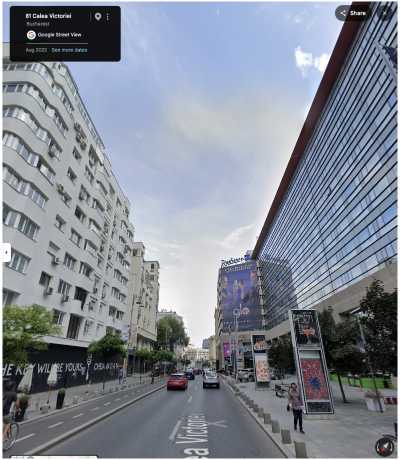

Street view — Calea Victoriei

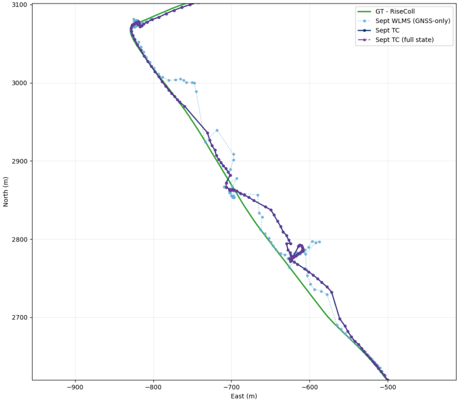

Trajectory in the same segment

Tensors

# doy118 — Strada Lipscani (old town, tree-lined)

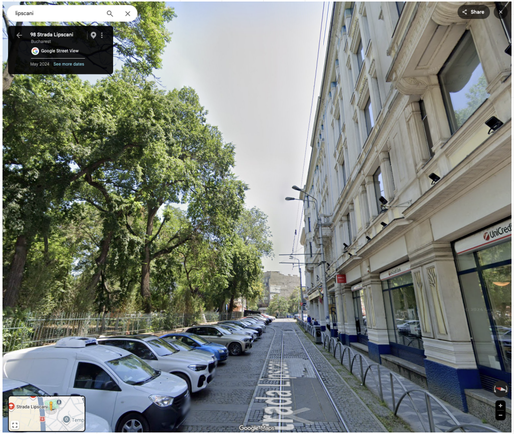

Street view — Strada Lipscani

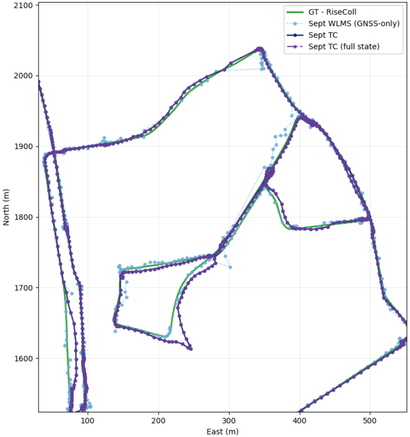

Trajectory in the same segment

# doy118 — Titulescu intersection (glass-faced towers)

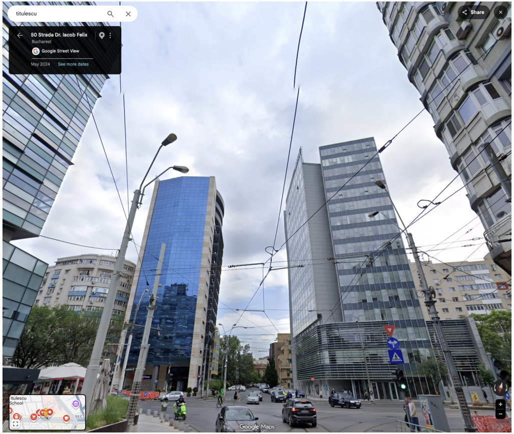

Street view — Titulescu intersection

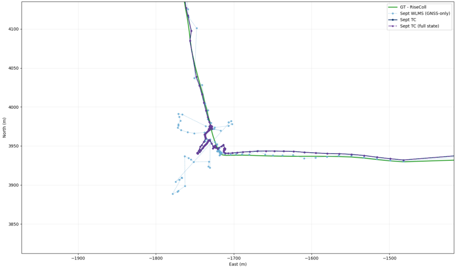

Trajectory in the same segment

# doy118 — tunnel exit

Street view — tunnel exit

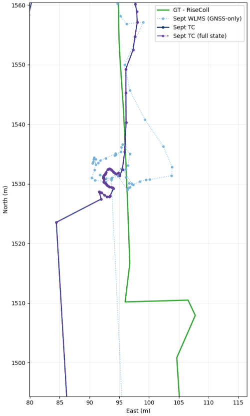

Trajectory in the same segment

# Table of Contents

# Conclusions (1)

- Successfully delivered a **NAFGO proof-of-concept testbed** combining GNSS receivers, IMU, data acquisition, GUI monitoring, dataset replay, and FGO-based navigation processing;
- More than **25.5 hours of driving** data were collected across different road types, receiver configurations, and operational conditions, with data prepared for a future public dataset;
- FGO proved technically viable for GNSS/IMU navigation, especially in TC configurations, where it can **improve availability, outlier rejection, and robustness in degraded GNSS conditions**;
- real-time FGO-LC generally tracks the receiver PVT solution closely, while it is also computationally viable, typically **below 100 ms on the Jetson** per epoch, but it also inherits much of the receiver PVT quality and limitations;
- Implemented **real-time (causal)** and batch-mode FGO;
- FGO-TC matches WLMS on typical epochs and tames the worst ones — across all 16 scenarios, the mean horizontal error over the hardest 1% of epochs drops 34% (33.4 m → 22.0 m).

Tensors

# Conclusions (2)

- Ublox F9P **low-cost receiver showed competitive performance** in most scenarios, indicating that low-cost hardware can be viable when paired with robust algorithms (but, raw-observable TC processing still benefits from the higher-quality Mosaic-X5 receiver);
- IMU integration is **useful during GNSS degradation or outage**, but the current GNSS + IMU-only setup **remains sensitive to heading initialization and observability at low speeds**;
- **The absence of an external heading reference is a major limitation**. Future versions should consider magnetometers, map constraints, vision, LiDAR, or other heading/pose aiding sources (even the availability of the ground-truth heading, to check effects on performance);
- For FGO, **marginalization window length has a clear accuracy/latency trade-off**. The useful knee appears around 15-20 s; longer windows increase computing cost significantly without a measurable significant accuracy improvement (more than 30 epochs do not affect performance);
- A few real-time FGO-TC runs diverged under challenging dynamics, indicating that **buffering, windowing, reordering, and robustification of the solution still require further important investigation**;
- **External RTK/PPP fusion showed mixed results**: on some runs, PVT and FGO-fused solutions matched or improved raw post-processed solutions, while on others they degraded;
- Performance variations suggest that the proposed pipeline needs dedicated diagnostics and that **current processing commercial pipelines are heavily optimized for accuracy** after years of real-world deployment.

# Lessons learned

- Ublox F9P comparable with Septentrio Mosaic X5: algorithms robust enough to handle small differences;
- IMU useful for corrections and in case of GNSS loss, sensitive to initialization and heading error;
- Current solution is pure GNSS + IMU, difficult to ensure good observability of the heading angle ⇒ use of magnetometer, dual antenna GNSS heading or map information
- FGO-LC requires handling of the RX PVT behavior, which is usually unknown and has to be observed by testing
- FGO-TC can compete with internal RX PVT solution but only if the same DSP chains are used
- To clarify the Ground Truth ambiguous source, alternative RTK base signals can be used;

# Challenges

- **No external heading reference** — vehicle heading has to be derived from velocity (`atan2(vE, vN)`), which is undefined at rest and noisy at low speed.
- **Missing Septentrio raw PVT on several runs** — a clock-sync / steering issue on the Mosaic-X5 delays PVT output at the start of some datasets, so the Sept raw PVT row is empty on those days.
- **GT-vs-estimate comparison requires lever-arm projection** — projecting from antenna to IMU body requires an attitude. Using the FGO's own yaw is circular (the estimate under test would define the truth frame). We adopted the velocity-derived heading and drop low-speed epochs from RMSE; both modes have been compared.
- **Suspect GT RTK behaviour** — high covariance reported in some fix epochs; degraded GT inflates apparent estimator error.
- **No raw-observable multipath rejection** on the TC pipelines yet.
- The GTSAM framework allows to **handle** 7% **loss of IMU packages** (the pre-integration factor integrates packages arriving non-uniformly at an increased 50Hz frequency)

# Future Work (beyond NAFGO)

**Extending NAGFO:**

- Add multiple new sensing modalities.
- Maps, images, Lidar, BT, Wi-Fi.
- Integrate Machine Learning techniques (e.g., KalmanNet).
- Goal is to maximize localization accuracy.

**Robust NAFGO**:

- Alg. for resilient NAFGO.
- Robust testbed.
- Mitigate jamming, spoofing, etc. attacks.

**Low-cost, efficient NAFGO**

- Cheaper hardware components for almost the same performance.
- Efficient approximate algorithms.

**GNSS + LoRa**

- Add localization in env. where there is currently no solution.
- Cheap, approximate, low maintenance.
- Emergency situations.

**GNSS + LoRa Mesh Network**

- Robust mesh-based localization.
- Network of off-grid devices, they coordinate to improve localization accuracy.
- Decentralized localization.

**Funding opportunities**

- NAVISP-EL2.
- Other ESA-funded project.
- Other sources.

# Thank you!

Questions?

`contact@three-tensors.com`

Tensors

Acknowledgment: The work is carried out under ESA NAVISP Element 1, devoted to the development of innovative PNT systems, technologies, algorithms and techniques
(Ref. No.: NAVISP1-TN-3T-088-1).

# Table of Contents

# Project schedule

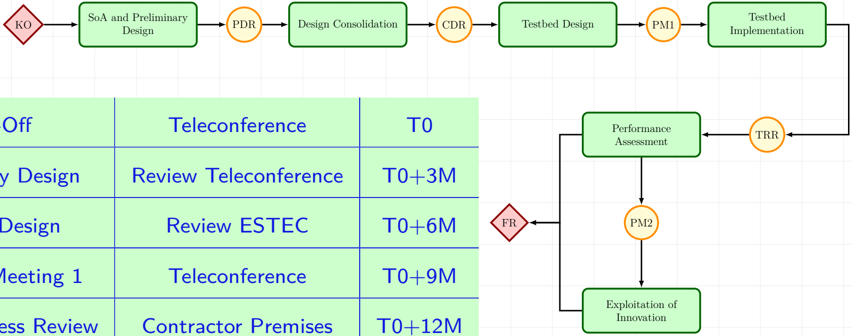

| | | | |
|---|---|---|---|
| KO | Kick-Off | Teleconference | T0 |
| PDR | Preliminary Design | Review Teleconference | T0+3M |
| CDR | Critical Design | Review ESTEC | T0+6M |
| PM1 | Progress Meeting 1 | Teleconference | T0+9M |
| TRR | Test Readiness Review | Contractor Premises | T0+12M |
| FR | Final Review | Teleconference | T0+18M |

# Deliverables list

| Doc ID | Title | Milestone |
|---|---|---|
| [D-01] | Design Definition File (DDF) | V. 1, end of Task 1<br>V. 2, end of Task 2<br>V. 3, end of Task 3 |
| [D-02] | Design Justification File (DJF) | V. 1, end of Task 1<br>V. 2, end of Task 2<br>V. 3, end of Task 3 |
| [D-03] | Verification and performance assessment report | V. 1, end of Task 2<br>V. 2, end of Task 3<br>V. 3, end of Task 4 |
| SW-UM | SW User Manual | end of Task 3 |
| HW-UM | HW User Manual | end of Task 3 |
| AB | Abstract | Final Review |
| FPH | Final Presentation Handout | Final Review |
| ESR | Executive Summary Report | Final Review |
| FR | Final Report | Final Review |
| CCD | Contract Closure Documentation | Final Review |
| WPI | Website Project Information | Kick Off Meeting |
| BIR | Benefit and Impact Report | Final Review |

# Table of Contents

## Testbed in car

# Installation and setup

**Mechanical design**:

- Testbed enclosed in a $700 \times 500 \times 245\,\mathrm{mm}$ ABS plastic enclosure, rated IP65 (rain-proof operation)
- 1 mm thick bent metal sheet fastened to the enclosure corners (base plate)
- Mounting via drilled holes (per CAD layout) + rivet nuts for threaded fixation
- Custom 3D-printed (PLA) supports for components without mounting holes

**Wiring**:

- Component placement optimized for short cable routes and to ensure connector mating clearance
- Free volume inside the enclosure (outside the testbed footprint) reserved for routing wires/connectors
- Cable management: Kapton tape for fixing wires; optional drilled holes for zip ties

**Assembled testbed**:

- All components fit within the enclosure; fit-check performed with the fully assembled testbed inside the waterproof enclosure

# Mechanical design

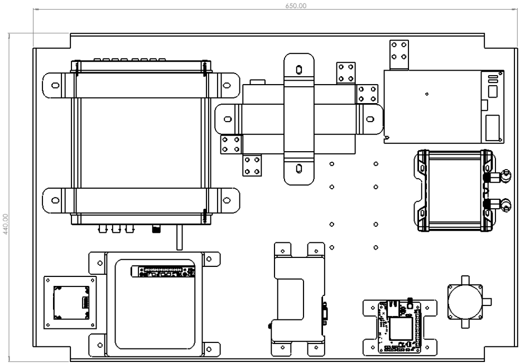

x-y plane

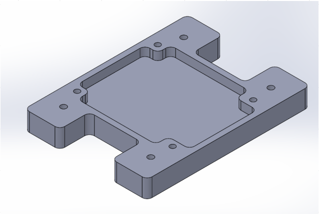

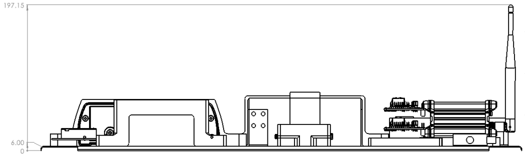

z (height)