

Reviewer 1:

Comment 1.1: CalMAPLab is a Ford Transit 250 retrofitted as a research-grade mobile air-quality laboratory, built around a strong contemporary instrument payload (Vocus PTR-ToF-MS, dual Aeris Mira Ultras, LI-7200, CAPS, EcoPhysics nCLD, 2BTech O₃, Fidas, AE33, CPC, Spider-MAGIC) and supported by VanDAQ, a custom open-source data acquisition and user interface framework. The vehicle and electrical integration are modern, and with 500+ hours of test driving the platform clearly works as intended. The manuscript has three problems the current draft does not adequately address. The novelty claim does not hold up against an existing literature the authors do not cite. The most distinctive piece of work, VanDAQ, is documented at a depth that does not match its novelty. The hardware description is precise enough to convince a reader the system works but not precise enough to let another group build one.

Response: We thank reviewer 1 for a thorough and constructive review. Addressing these comments has produced a stronger and more accessible manuscript.

Comment 1.2: The novelty claim is not consistent with the prior platform-description literature. The Introduction names Drewnick et al. (2012) and Xia et al. (2020) as the only platform descriptions providing detailed open documentation, and the abstract and conclusions advance the manuscript as a "transferable blueprint" filling a gap. This reading of the field is incomplete. Bush et al. (2015, AMT) describe a compact mobile observatory of nearly identical scope and target species (CO₂, CO, CH₄, NO_x, O₃, aerosol, meteorology) with a comparable instrument payload, deployed across the western US, and is uncited. Whitehill et al. (2024, AMT) describe Aclima's Ford Transit-based mobile calibration laboratory. They operated in the same San Francisco Bay Area sampling domain in which CalMAPLab is tested, with a published cross-comparison against fixed reference sites, and is also uncited. The Aerodyne Mobile Laboratory (AML) and the broad literature it has generated (Kolb et al., Yacovitch et al., Herndon et al.), and the NOAA CSD mobile van, are not engaged with either. Several of these prior platforms address exactly the topics the present manuscript frames as documentation gaps: synchronised multi-instrument acquisition, GPS-merged data products, and plume-mapping demonstrations. The blueprint framing should be revised to engage with these precedents. The authors should state plainly what is incrementally new, most plausibly the integration of a Vocus PTR-ToF-MS into a community-scale Ford Transit deployment paired with a published open-source DAQ stack, rather than imply a void that is not actually present.

Response: We thank the reviewer for pointing out unreferenced literature related to mobile laboratory data collection. We have added the recommended citations. We have also clarified the main novelty of this manuscript through rewording the

introduction. Our main novelty point is the detail and characterisation of all major components within a mobile laboratory and their integration into a single complete system. We agree that each of the examples provided by the reviewer contain in-part the same components as described in this manuscript. But a complete open design, including how each component interacts, the set of decisions made during construction, and full characterisation is not presented. The rewording of the introduction and a statement on the novelty of VanDAQ is copied here as well:

“While mobile monitoring has become more widespread, building and operating a mobile laboratory remains a complex engineering challenge. Powering sensitive, power-hungry scientific instrumentation requires robust electrical systems that can operate reliably in the field, often under variable environmental conditions. Thermal management and vibration isolation are critical to maintaining measurement stability and protecting precision instrumentation. Additional complexities exist in synchronizing multiple 1 Hz data streams with precise GPS timing, often via diverse communication protocols. Each of these challenges have been tackled in the mobile measurement literature, both in dedicated descriptions and in the campaigns built upon them (Bukowiecki et al., 2002; Popovici et al., 2018; Wagner et al., 2021; Bush et al., 2015; Whitehill et al., 2024; Wild et al., 2017; Yacovitch et al., 2015). Comparatively few document all major components, their integration into a complete system and their real-world performance together in one place (Drewnick et al., 2012; Xia et al., 2020). The data acquisition and management software is a particular gap where, if it is described at all, it typically depends on proprietary tools/software and is not publicly available. Therefore, reproducibility remains hard and the barrier to entry is high.

The growing demand for mobile monitoring in response to environmental justice and community-scale research highlights the need for open, rigorously engineered, and field-validated mobile laboratory designs. In this context, we present the design and performance evaluation of the UC Berkeley **Mobile Air Pollution Laboratory** (CalMAPLab). This paper details the design process, engineering decisions, power management, system architecture, data post-processing, and an open-source data-acquisition and display framework (termed VanDAQ) necessary for the successful operation of the CalMAPLab. The ability to acquire and visualise data from a diverse range of instrumentation using open-source software (Python and PostgreSQL) rather than a reliance on proprietary means (e.g. LabView, IgorPro, DAQFactory) is a key contribution to the atmospheric sciences community. We anticipate this, in combination with detail on all other parts of the system, will reduce the barrier to entry for others interested in hyperlocal atmospheric research.”

Comment 1.3: An open-source, instrument-agnostic acquirer/collector/submitter pipeline that writes into a star-schema SQL database, drives a live operator dashboard, transmits one-minute submission files to a central server, and supports a map-based interface for adaptive plume tracing is a contribution worth publication on its own. The treatment in Section 2.7 does not give the reader enough to use it.

Response: We appreciate the acknowledgement of the efforts that went into the creation of the data acquisition system. While the full technical depth is best housed in the repository documentation, we agree the manuscript should give the reader enough to engage with it and direct them to the repository effectively. We have expanded Section 2.7 accordingly (see responses to Comments 1.4-1.9).

Comment 1.4: The configuration-file format that lets a single acquirer code body service many instruments is the central element of the instrument-agnostic claim, and the manuscript does not show an example.

Response: A description of the acquirer configuration with examples are provided in the repository documentation, and the reader has now been directed to this in the Section 2.7 text. See also response to Comment 1.8 for “adding a new instrument” tutorial.

“Configuration file format, including examples and how to set up a new instrument, are detailed in the GitHub repository documentation.”

Comment 1.5: The time-synchronisation regime across acquirer processes is not stated (NTP-disciplined? GPS-disciplined? what is the worst-case timestamp skew between two acquirers handling a 1 Hz event?).

Response: The system is SNTP-disciplined by systemd-timesyncd and synchronised against ntp.ubuntu.com. However, all acquirer processes run on the single data acquisition host within the van and assign a timestamp from the host’s operating system wall clock at the moment each measurement record is parsed and enqueued at 1 s resolution. Therefore, skew between acquirers is not clock drift between nodes, but rather latency and scheduling between host datetime calls. This skew is bounded by the serial read latency, process scheduling and queue handoff, all within the same second bin. The maximum skew for this set up is 1 s where two acquirers stamp adjacent seconds due to parsing some ms either side of an integer second. Across a 24 h period, we calculate that 99.5-99.7 % of pairwise

observations from the TSI CPC (polled), Licor CO₂ and Aeris CH₄/C₂H₆ (streamed) show identical timestamps with all remaining cases differing by exactly 1 s. This detail has been added to the Section 2.7 text as outlined below.

“Configuration file format, including examples and how to set up a new instrument are detailed in the GitHub repository documentation. The system is SNTP-disciplined and synchronized by system-timesyncd with time alignment across instruments established from deployment of acquirers on a single acquisition computer. Each acquirer process timestamps records using the shared host wall clock with 1 s resolution. Possible skew between acquirers is bounded by serial latency, process scheduling and queue handoff. We measure a maximum timestamp skew between pairwise observations of 1 s from acquirers stamping adjacent seconds 0.3 to 0.5 % of the time.”

Comment 1.6: The behaviour of the POSIX queue between acquirer and collector under back-pressure or process failure is not described, and it is unclear whether submission files persist to disk before acknowledgement (what happens when the system crashes or loses power?).

Response: The queue is a small, blocking, in-memory IPC buffer set at a max of 50 messages and 8 kb of data per message. There is no specific acknowledgement protocol between acquirer and collector. The submission files are written after the database commit in the collector on a timed interval for backup/export via upstream transfer and are not the primary durability of the system. The durability is provided by PostgreSQL following the collector commit. Should power loss occur, only data within the IPC and therefore uncommitted to PostgreSQL is lost. This description has been added to the Section 2.7 text.

“The unified and packaged measurement data are loaded by the various acquirer processes into a single blocking POSIX interprocess communications queue (max messages = 50, max data per message = 8 kb) for further processing downstream.”

“Durability of the system is primarily provided by the PostgreSQL database. Following power loss, only data within the IPC and uncommitted to the database is lost.”

Comment 1.7: The database write pattern, per-row inserts versus COPY, and how throughput scales across 1 Hz × 500 h × N instruments, is not measured.

Response: We use batched INSERT rather than per-row INSERT, but agree that the write-behaviour has not been presented. Within the collector logs, we calculate a sustained insert throughput of ~ 1300 rows s^{-1} , compared to a daily arrival rate of ~ 160 rows s^{-1} (from roughly 160 1 Hz instrument parameter pairs, including some engineering data) indicating significant operational headroom for continuous acquisition. We have revised the Section 2.7 text to report this throughput information, and clarified that live ingest currently uses batched multi-row INSERT, with COPY reserved for future bulk-loading if necessary.

“Ingestion performance was quantified from production collector logs with an average sustained throughput of ~ 1300 rows s^{-1} compared to an average arrival rate of ~ 160 rows s^{-1} from roughly 160 1 Hz instrument-parameter pairs (including some engineering data). This indicated significant operational headroom for continuous acquisition with the current instrument payload.”

“loads the data items into the mobile server’s PostgreSQL database via batched row INSERT”

Comment 1.8: The repository link is given but the manuscript does not point to a license, a tagged release for the paper, a CITATION.cff and Zenodo DOI, an "adding a new instrument" tutorial, or any test coverage.

Response: We have added explicit licensing terms (BSD 3-Clause) to the Code Availability section. We have created a tagged release corresponding to this manuscript and archived it on Zenodo (DOI: 10.5281/zenodo.20705283), added a CITATION.cff file and included an “Adding a new instrument” tutorial in the repository documentation.

We have added a pytest-based unit test suite (43 tests) covering instrument string parsing, alarm rules, NMEA handling, collector database insertion and serial line buffering. Tests run on every push and pull request via GitHub Actions on Ubuntu with Python 3.10 with results and a report produced for each run (available as a CI artifact). Hardware-dependent paths and integration tests requiring live PostgreSQL are excluded from CI and documented for optional local execution. In this setup, we address reproducible verification of core ingest logic (43% coverage on acquirers.py and 41% on van_collector.py) with lower repository wide coverage but opportunities to expand in the future.

“The data acquisition and post-processing code is distributed under a BSD 3-clause License and available from <https://github.com/CalMAPLab>, with version 1.1.0 archived on Zenodo (DOI: 10.5281/zenodo.20705283).”

Comment 1.9: Until these are in place the open-source framing is a label rather than a deliverable. I would suggest re-anchoring the manuscript on VanDAQ as its central contribution, with the hardware payload providing context.

Response: By addressing Comments 1.4 to 1.8, we believe the open-source framing has now been delivered. We thank the reviewer for the re-anchoring suggestion and have revised the manuscript to make VanDAQ more central, in part through enhanced descriptions of features (Comments 1.4 to 1.8), prominence within the introduction (Comment 1.17), and toning back some of the blueprint language (Comment 1.10 and 1.15).

Comment 1.10: The hardware sections are detailed enough to be persuasive but not detailed enough to be reproducible. For a paper whose stated goal is a transferable blueprint, several of the engineering decisions that most affect a replicator's outcome are described only in outline.

Response: Thank you for the comment. This has been addressed by increasing the description of the engineering decisions in relation to Comment 1.11 to 1.13, and softening the blueprint wording throughout the manuscript in response to Comment 1.15.

Comment 1.11: Fig. 3 lists per-instrument flow rates, but inlet residence time is given only as "2–3 seconds." A 1 Hz platform claiming meter-level resolution warrants a tabulation of residence time as a function of line ID, length, and flow, and the laminar-flow figure ($Re = 1200$ for the particle inlet) should be accompanied by an entrance-length calculation given the inlet cone geometry.

Response: We have addressed this in two parts. Firstly, lag times are determined by cross correlation of paired species across instruments which captures the combined effect of inlet residence time, internal instrument response, processing latency, and the database pipeline, as well as any pump degradation over time. This is why we reported an operational lag rather than purely geometric residence time, and we have clarified this in the Section 2.8 text. However, we agree that a geometric tabulation may still be useful and have added a column to Table A1 detailing residence time as a function of line internal diameter, length and flow rate for each inlet.

For the particle inlet we calculate the hydrodynamic entrance length under laminar flow as $L_n = 0.05 \times Re \times D = 0.76 \text{ m}$. This has been added to the Section 2.6 text accordingly.

“Flow to the gas rack (~ 5 slpm) and LI-7200 (~ 10 slpm) is provided through the internal pumps of the instruments and for a response time of ~1 second (see Table A1 for breakdown by instrument).”

“Flow is provided by the instrument's internal pumps for laminar flow ($Re = 1400$ at 0.43” ID tube diameter, 11 LPM, 20 °C and 1000 mbar with an entrance length of 0.76 m) up until the flow splitters.”

Comment 1.12: The Mira Ultras, CAPS, and Vocus are temperature-sensitive, and the sample lines run between an interior climate-controlled volume and a roof-mounted ambient inlet. The transient thermal response of the lines as the van moves between, say, a 35 °C parking lot and a 12 °C coastal fog bank is not addressed.

Response: We thank the reviewer for raising sample line thermal transients. The inlet draws ambient air into a van maintained at $\geq 23 \text{ °C}$. Analysis of the test dataset showed the median daily roof temperature range as measured by the Airmar weather station was 7.7 °C (min = 0.3 °C, max = 22.4 °C). Sample gas temperature is measured by the Licor 7200 CO₂. In comparison to ambient temperature, the median daily sample gas temperature range during operation was measured as 2.1 °C (min = 0.1 °C, max = 12.4 °C). This indicates a strong thermal buffering relative to roof ambient temperature and consistent sample line temperatures throughout operation. For uncorrected air density-sensitive measurements, biases are estimated at ~1%, rising to ~3 % on the days with the greatest temperature variability. Within our instrument payload, these are largely accounted for through measured cell temperature and subsequent internal corrections.

We also do not experience any condensation issues from sampling close to the dew point temperature. Sample dew point was at least 5 °C below the van internal temperature 100% of the time. However, we do note that should sampling take place in more humid environments, heating the sample line is an important consideration.

This discussion has been added to the Section 3.1 text, and a figure highlighting the sample line thermal buffering has been added to the SI.

“Some atmospheric measurements are sensitive to the density and therefore temperature of the sample gas. To assess thermal transients along the sampling path, we compare roof ambient temperature with sample gas temperatures reported at the cell inlet of the Licor 7200. On average, the median daily temperature range on the roof was 7.7 °C (min = 0.3 °C, max = 22.4 °C) whereas median daily ranges at the Licor was 2.1 °C (min = 0.1 °C, max = 12.4 °C, see Fig. S11 for more detail). These results indicate substantial thermal buffering of sample air relative to roof ambient, consistent with heating and residence time in the climate-controlled cabin rather than directly tracking outdoor temperature. For uncorrected density-sensitive measurements, biases are estimated at ~1% rising to ~3% on the days with greatest temperature variability. Within our instrument payload, these are largely accounted for through measured cell temperature and subsequent internal corrections. In addition, no condensation issues were encountered where sample dew point was at least 5 °C below the van internal temperature 100% of the time. However, we do note that should sampling take place in more humid environments, heating the sample line is an important consideration.”

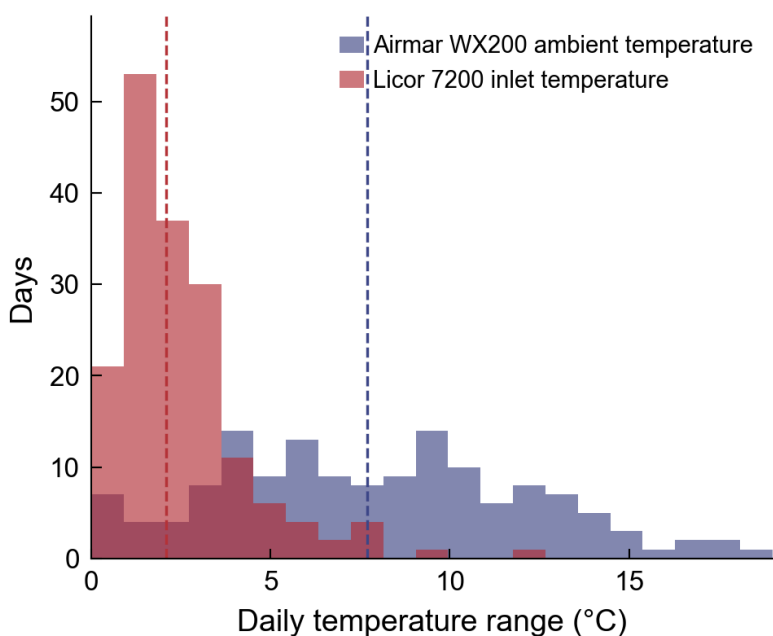


Figure S11: Distributions of temperature range (max-min) for individual drives as measured on the roof with the Airmar WX200 weather station and at the Licor 7200 inlet. The lower magnitude and narrower range of the Licor highlights strong thermal buffering within the sample line relative to ambient temperature.

Comment 1.13: The GitHub repository is referenced but the manuscript does not point to specific design files: rack-layout CAD, machining files for the inlet plates, a parts BOM with quantities. For a transferable blueprint these need to be either in the supplement or clearly catalogued in the repository, with the manuscript pointing readers to them.

Response: Having softened the blueprint framing (Comment 1.15), the expectation that every component is reproducible from supplied files is reduced. However, we have improved the documentation in Section 2.5 by detailing that the instrument racks are commercially purchased, off-the-shelf units rather than custom fabrications, and included photographs in the SI for rack specific layout should the reader be interested (see Comment 2.2).

“For vibration dampening, most instrumentation is mounted within two off the shelf 24U 19” server racks each bolted to a shock-rope isolation plate system.”

Comment 1.14: The "community-scale" framing is asked to carry more than the application supports. "Community-scale" appears in the title, abstract, and conclusions, and the West Oakland demonstration is the only engagement with a community context. As presented it is a technical demonstration of the live-mapping interface rather than community-engaged research. The methanethiol-based attribution of a CH₄ plume to a wastewater treatment plant is qualitatively compelling but is not quantified: no emission rate, no plume-decay analysis, no comparison against inventory. It is an existence proof rather than a finding. The framing should be revised to "neighborhood-scale" or "fine-spatial-scale."

Response: We thank the reviewer for raising concerns on the community-scale framing. While all of the work that we do in West Oakland and surrounding neighborhoods is community focussed, with close collaboration with local community groups and designing monitoring strategies around their concerns, this manuscript is mostly referring to the geographical area that is a community. Since this could be confusing for the reader, most community-scale wording has been changed to either hyperlocal or neighborhood-scale throughout the manuscript, including the title. However, we do keep some references to community-scale in relation to potential applications, where community air monitoring (particularly on the West Coast of the US) is a growing area of mobile monitoring deployment (see AB617 legislature).

Regarding the wastewater treatment plant analysis, this is an active field of research. More detailed and quantified science on this topic is currently in preparation, will be published separately, and is therefore out of scope for this

manuscript. We hope the reviewer continues to follow our work for more detail on this topic.

Comment 1.15: The paper currently reaches for three different framings, a novel platform, a community reference benchmark, and a replicable open design, and is stronger on the third than the first two. Committing more clearly to one, and accepting that a platform-description paper does not need to be all three, would tighten the manuscript considerably.

Response: We thank the reviewer for this helpful framing comment and believe that we are now committing to a tighter framing with novelty tied to the open design and in depth documentation of all major components relative to prior builds (Comments 1.2 and 1.17). The community reference benchmark has been reframed around the hyperlocal and neighborhood geographical scales afforded by this laboratory build (Comment 1.14).

Comment 1.16: The section ordering (goals → instruments → vehicle → power → layout → inlet → DAQ → post-processing → application) reads cleanly, and the prose is generally precise. Section 2.2 substantially restates Table A1 and could be cut to roughly half its current length by deferring instrument specs to the table and confining the prose to the design choices (e.g., Fidas vs alternative PM monitors, Spider-MAGIC vs conventional SMPS).

Response: The purpose of Section 2.2 was to explain why each instrument and species was chosen to be monitored, and include details specific to individual instruments that would not apply across instruments in the table e.g. Vocus PTR-ToF-MS specifications vs Spider Magic ionisation source. We do not feel that the information given in the table (weight, time resolution, dimensions, communications protocol) is restated in the text. However, we have removed the repetition of instrument manufacturer details from the text to trim the section.

Comment 1.17: A short "novelty" framing at the close of Section 1 would set reader expectations. At present the strongest contribution (VanDAQ) is not encountered until Section 2.7.

Response: A short novelty statement on VanDAQ, and the additional detail on all other parts of the system has been added to the Introduction in response to Comment 1.2.

Comment 1.18: "kW/hour" on line 414 is not a valid charging-rate unit; Fig. 6d correctly uses kWh for the cumulative quantity.

Response: Thank you for spotting this mistake. This has been changed to kW units.

“Overall charging rate was very consistent (see Fig. 6d) at ~ 3.2 kW corresponding to a 8 hour full charge time.”

Comment 1.19: The $Re = 1200$ figure for the particle inlet should be accompanied by the diameter, viscosity, and velocity used.

Response: The flow regime discussion in Section 2.6 now includes the diameter, viscosity and velocity used in the calculation.

“Flow is provided by the instrument's internal pumps for laminar flow ($Re = 1400$ at 0.43" ID tube diameter, 11 LPM, 20 °C and 1000 mbar with an entrance length of 0.76 m) up until the flow splitters.”

Comment 1.20: Hyphenation is inconsistent (real time vs real-time), and "battery-electric" capitalisation drifts between Methods and Conclusions.

Response: All hyphenation instances have been changed to real-time and all incorrect capitalisations of “battery” have been replaced with lower case letters.

Comment 1.21: The reference list is generally fit for purpose but should be expanded to engage with the prior platforms identified above – a deeper literature review may be necessary (Bush 2015, Whitehill 2024, the AML literature, the NOAA CSD van).

Response: The reference list and literature engagement has been expanded accordingly (see Comment 1.2).

Comment 1.22: I recommend major revisions. The platform is real, the engineering is solid, and the paper has a place in the AMT literature. To make this a better manuscript, the authors should

- (1) reframe the novelty claim against the prior platform-description literature,
- (2) substantially deepen Section 2.7 so the open-source claim is operationally meaningful,
- (3) close the reproducibility gaps in thermal management, pumping, inlet, and construction-level documentation, and

(4) either dial back the "community-scale" framing or back it with a quantitative scientific demonstration.

A more focused paper anchored on VanDAQ and the open-design framing, with the hardware as supporting context, would in my view be a stronger contribution than the current attempt to claim all three framings at once.

Response: We thank the reviewer for the constructive summary and for recognising the platform's place in the AMT literature. We have acted on all four recommendations: (1) the novelty claim has been reframed against the prior platform literature (Comments 1.2 and 1.21); (2) Section 2.7 and the GitHub repository has been substantially deepened so the open-source contribution is operationally meaningful (Comments 1.3-1.9); (3) we have closed gaps in thermal management, inlet and construction level documentation (Comments 1.10-1.13 and 1.19); (4) we have dialed back the community scale framing to a more geographic-based context (Comment 1.14 and 1.15).