

Radar-Inertial Ego-Velocity Estimation for Visually Degraded Environments

Andrew Kramer, Carl Stahoviak, Angel Santamaria-Navarro,
Ali-akbar Agha-mohammadi and Christoffer Heckman

Abstract— We present an approach for estimating the body-frame velocity of a mobile robot. We combine measurements from a millimeter-wave radar-on-a-chip sensor and an inertial measurement unit (IMU) in a batch optimization over a sliding window of recent measurements. The sensor suite employed is lightweight, low-power, and is invariant to ambient lighting conditions. This makes the proposed approach an attractive solution for platforms with limitations around payload and longevity, such as aerial vehicles conducting autonomous exploration in perceptually degraded operating conditions, including subterranean environments. We compare our radar-inertial velocity estimates to those from a visual-inertial (VI) approach. We show the accuracy of our method is comparable to VI in conditions favorable to VI, and far exceeds the accuracy of VI when conditions deteriorate.

I. INTRODUCTION

Accurate and reliable estimates of ego-velocity are crucial for closed-loop control of autonomous mobile robots during navigation operations. This is especially true for fast-moving robots like micro aerial vehicles (MAVs). Robot body-frame velocities are commonly estimated using some combination of visual, LiDAR, inertial and/or GPS sensors. Accurate ego-velocity estimates are intrinsic to any number of simultaneous localization and mapping (SLAM) methods that have been developed. For those methods that rely on visual data, the quality of the ego-velocity estimate is quickly degraded in darkness, feature-poor environments, and so forth. Here, we consider the subterranean environment as a motivating example.

Robust autonomy in subterranean environments is currently a popular research topic. NASA is planning to explore caverns on the moon and Mars [1], while DARPA is conducting its Subterranean Challenge¹. Most state-of-the-art methods for body-frame velocity estimation are significantly impaired in conditions common to subterranean environments, e.g. GPS data is unavailable and cameras cannot capture useful information in complete darkness. Thermal imaging

A. Kramer, C. Stahoviak and C. Heckman are with the Department of Computer Science at University of Colorado, Boulder, Colorado 80309, USA {andrew.kramer; carl.stahoviak; christoffer.heckman}@colorado.edu

A. Santamaria-Navarro and A. Agha-mohammadi are with the Jet Propulsion Laboratory, California Institute of Technology, 4800 Oak Grove Dr, Pasadena, CA 91109, USA {angel.santamaria.navarro; aliakbar.ghamohammadi}@nasa.jpl.gov

This work was supported through the DARPA Subterranean Challenge, cooperative agreement number HR0011-18-2-0043. Part of this research was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration.

¹<https://www.subtchallenge.com>

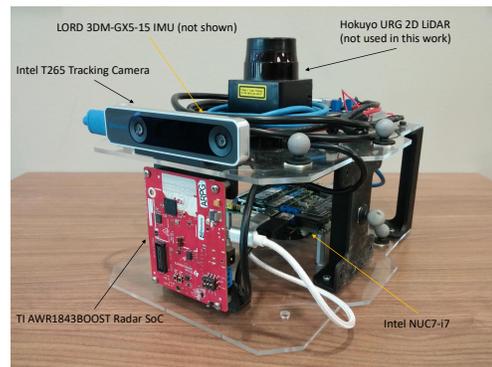


Fig. 1. Experimental sensor rig consisting of 77–81 GHz FMCW radar, Intel RealSense T265 tracking camera, LORD microstrain IMU and Intel NUC onboard computer.

is a popular option for robots operating in darkness, smoke or fog [2], but these methods track temperature gradients, which are often not present in subterranean environments. Thus, it is clear that there exists a need for a reliable and efficient method for ego-velocity estimation which is capable of generalizing across diverse environments and operating conditions.

Millimeter wave radar is an attractive option for subterranean environments. It does not require light or temperature gradients to operate. Additionally, automotive-grade system-on-chip (SoC) radars have low power requirements. However, radar measurements are adversely affected by sensor noise and radar-specific corruptions of data, e.g. multipath reflections and binning of spatial and Doppler measurements. So, while it is certainly possible to estimate a robot’s body-frame velocity from standalone radar data using robust optimization techniques [3], these estimates are not sufficiently accurate for reliable control. Additionally, the antenna pattern of the radar will naturally provide more accuracy in some dimensions than others. In order to overcome the shortcomings of radar as a standalone sensor, it is beneficial to fuse radar and inertial measurements.

This work presents a method for ego-velocity estimation that uses an automotive-grade radar SoC (Texas Instruments AWR1843) and a MEMS IMU. We use Doppler velocity measurements to estimate the body-frame velocity of the sensor at the time of each radar measurement. We jointly estimate ego-velocities over a sliding window of the previous K radar measurements in a nonlinear optimization framework, using IMU measurements to constrain the change in velocity between radar measurements. The addition of inertial data



Fig. 2. Example of the subterranean analog environment in which we tested: the steam tunnels beneath Folsom Field at CU Boulder.

helps to smooth the high noise that would be present if we were estimating body-frame velocity from radar data alone. Conversely, velocity estimates from the radar are drift-free, so the radar information allows us to estimate the biases of the IMU.

This paper is organized as follows. Section II reviews related work in the area of ego-velocity estimation. Section III details our velocity estimation method. Section III-A briefly describes how Doppler velocity and inertial constraints are combined to accurately estimate the sensor platform’s body-frame velocity. Section III-B describes how body-frame velocity can be estimated from millimeter wave radar data and how the Doppler residual is formulated in our optimization. Section III-C explains the IMU kinematics used in our problem and how the inertial constraint is formulated. We then highlight the process for experimentally validating our method in Section IV. Finally, we discuss our results and conclusions in Sections V and VI, respectively.

II. RELATED WORK

While radar is well-established in the automotive industry and has been used for various tasks in vehicle autonomy including collision avoidance, automated braking, lane keeping, autonomous parking, etc. [4], very few methodologies for using radar as a primary sensor for ego-motion estimation have been presented in the literature. The most consequential of these have focused on odometry and SLAM using only millimeter wave radar. [5] presents methods for radar landmark detection, scan matching, and odometry in diverse and challenging conditions; this work is extended in [6]. Additionally [7] and [8] have presented methods for radar-based SLAM using arrays of automotive-grade radar sensors. Lastly, [9] made use of Doppler velocity measurements for body-frame velocity estimation as part of their odometry system.

Additionally, good results for ego-velocity estimation have been obtained using stereo camera configurations and RGB-D sensor systems [10], [11], [12]. The use of a smart camera (optical flow), an IMU and a range sensor is proposed in [13] and [14]. Similarly, methods like [15] take advantage of a LiDAR or a combination of LiDAR and IMU measurements as in [16]. Although these approaches are able to efficiently

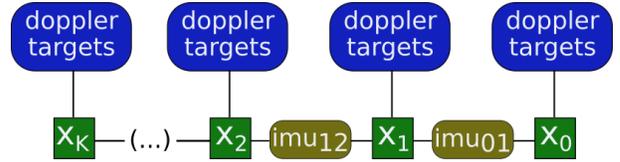


Fig. 3. Factor graph representation of the radar-inertial velocity estimation system. States from N previous timesteps are jointly estimated using Doppler targets and sets of IMU measurements as constraints

estimate ego-velocity, they are ill-equipped for use in environments with challenging sensing conditions.

Like our method, several previous state estimation systems have been formulated as maximum a posteriori estimation problems represented by factor graphs [17]. These include the popular visual method ORB-SLAM [18], [19] and the visual-inertial method OKVIS [20] from which our radar-inertial method draws considerable inspiration.

Our method represents a significant advance over previous methods in several important ways. First, our method is applicable to 3 dimensional environments, while the previously mentioned radar state estimation methods are only applicable to 2D environments. This means that our method is usable on micro aerial vehicles and other robots that do not operate in planar environments. Second, our method fuses radar and inertial measurements, providing highly accurate estimates and constraining the large uncertainty normally associated with radar-based methods. Previous radar state estimation methods have used radar as their sole sensor and robotic sensor fusion methods with radar have only been minimally explored. Lastly, previous radar-based state estimation methods have either depended on highly specialized scanning sensors [5], [6], [9] or arrays of several automotive sensors requiring precise extrinsic calibration [7], [8]. Our method requires only one single-board radar sensor and an IMU, and thus the sensor package we employ is simpler than those used in previous methods.

III. METHODOLOGY

A. System Structure

Our approach estimates the body-frame velocity of the sensor platform over a sliding window [21] of K previous radar measurements. These velocities are linked by integrated accelerometer measurements from the IMU. Figure 3 shows a factor-graph representation of our system’s structure.

Accelerometer measurements are affected by both bias \mathbf{b}_a and gravity \mathbf{g}_W . Velocity estimates from radar are bias free, so we can compensate for the accelerometer biases by simply including them in the state vector. Compensating for gravity is more complicated, however. To do this we need to estimate the IMU’s attitude (pitch and roll), which we represent as the orientation quaternion \mathbf{q}_{WS} . In order to estimate the IMU’s attitude, we need to use gyro measurements and to do this we must also estimate the gyro biases \mathbf{b}_g . The full state vector is then given as $\mathbf{x} = [\mathbf{v}_S^T, \mathbf{q}_{WS}^T, \mathbf{b}_g^T, \mathbf{b}_a^T]^T$.

We formulate our radar-inertial ego-velocity estimation as an optimization over the cost function

$$J(\mathbf{x}) := \underbrace{\sum_{k=1}^K \sum_{d \in \mathcal{D}^k} e_d w_d}_{\text{Doppler term}} + \underbrace{\sum_{k=1}^{K-1} \mathbf{e}_s^{kT} W_s^k \mathbf{e}_s^k}_{\text{inertial term}} \quad (1)$$

where K is the number of past radar measurements for which states are estimated, \mathcal{D}^k is the set of targets returned from the radar measurement at time k , e_d is the Doppler velocity error, \mathbf{e}_s is the IMU error. The error terms are weighted by the information matrix W_s in the case of the IMU errors; and the normalized intensity of the corresponding radar target

$$w_d^j = \frac{i_j}{\sum_{d \in \mathcal{D}} i_d} \quad (2)$$

in the case of the Doppler velocity measurements where w_d^j is the weight for target j in scan \mathcal{D} and i_j is the intensity of target j . In the following sections we detail the formulation of our Doppler and IMU measurement constraints.

B. Estimating Ego-Velocity From Doppler Velocity Measurements

A radar measurement consists of a set of targets \mathcal{D} . Each $d \in \mathcal{D}$ consists of $[r_S, v_R, \theta_S, \phi_S]^T$, the range, Doppler (radial) velocity, azimuth, and elevation for target d . The Doppler velocity measurement v_R is equal to the magnitude of the projection of the relative velocity vector between the target and sensor \mathbf{v}_S onto the ray between sensor origin and the target \mathbf{r}_S . This is simply the dot product of the target's velocity in the sensor frame and the unit vector directed from the sensor to the target

$$v_R = \mathbf{v}_S \left(\frac{\mathbf{r}_S}{\|\mathbf{r}_S\|} \right)^T \quad (3)$$

We assume the targets in the scene are stationary and only the sensor platform is moving. In this case each radar target can provide a constraint on our estimate of the sensor rig's velocity in the body-frame. The velocity error for each radar target is then

$$e^k(\mathbf{x}^k, \mathbf{d}^{i,k}) = v_R^{i,k} - \mathbf{v}_S^k \left(\frac{\mathbf{r}_S^{i,k}}{\|\mathbf{r}_S^{i,k}\|} \right)^T \quad (4)$$

where x^k is the state at time k and $\mathbf{d}^{i,k}$ is the i^{th} target in the set of radar measurements at time k . As previously noted, radar measurements are affected by non-Gaussian noise and radar scans often contain false target data. These challenges are addressed by using the Cauchy robust norm with the Doppler residual.

C. Formulation of the IMU Constraint

1) *IMU Kinematics*: In our system, the IMU's accelerometer readings are used to measure the system's change in body-frame velocity between radar measurements. Our IMU model is very similar to those used in OKVIS [20] and MSCKF [22], except we do not use the IMU to measure the change in the system's full pose, only its velocity

and attitude. The states are propagated via the following differential equations:

$$\begin{aligned} \dot{\mathbf{q}}_{WS} &= \frac{1}{2} \boldsymbol{\Omega}(\omega_S) \mathbf{q}_{WS} \\ \dot{\mathbf{v}}_S &= \mathbf{a}_S + \mathbf{C}_{SW} \mathbf{g}_W - (\omega_S) \times \mathbf{v}_S \\ \dot{\mathbf{b}}_g &= \mathbf{w}_{b_g} \\ \dot{\mathbf{b}}_a &= -\frac{1}{\tau} \mathbf{b}_a + \mathbf{w}_{b_a} \end{aligned} \quad (5)$$

The elements of $\mathbf{w} := [\mathbf{w}_g^T, \mathbf{w}_a^T, \mathbf{w}_{b_g}^T, \mathbf{w}_{b_a}^T]^T$ are zero-mean, uncorrelated Gaussian white noise, \mathbf{g}_W is the gravity vector in the world frame, and \mathbf{C}_{SW} is the rotation matrix specifying the rotation from the world frame of reference to the sensor frame. The gyro and accelerometer measurements, $\tilde{\omega}_S$ and $\tilde{\mathbf{a}}_S$ respectively, are defined as the true acceleration and angular rate with added bias and white noise

$$\begin{aligned} \tilde{\omega}_S &= \omega_S + \mathbf{b}_g + \mathbf{w}_{b_g} \\ \tilde{\mathbf{a}}_S &= \mathbf{a}_S + \mathbf{b}_a + \mathbf{w}_{b_a} \end{aligned} \quad (6)$$

and the matrix $\boldsymbol{\Omega}$ is formed from the estimated angular rate as

$$\boldsymbol{\Omega}(\omega_S) := \begin{bmatrix} -\omega_S \\ 1 \end{bmatrix}^{\oplus} \quad (7)$$

The \oplus operator is defined in [23].

In order to optimize over this model we need the linearized, discrete time version of the state equations in Eq. (5). First, the continuous time state transition matrix is

$$\mathbf{F}_c(\mathbf{x}) = \begin{bmatrix} \mathbf{0}_{3 \times 3} & \mathbf{0}_{3 \times 3} & \tilde{\mathbf{C}}_{WS} & \mathbf{0}_{3 \times 3} \\ -\tilde{\mathbf{C}}_{WS}[\mathbf{g}_W]^\times & -[\omega_S]^\times & -[\mathbf{v}_S]^\times & -\mathbf{I}_{3 \times 3} \\ \mathbf{0}_{3 \times 3} & \mathbf{0}_{3 \times 3} & \mathbf{0}_{3 \times 3} & \mathbf{0}_{3 \times 3} \\ \mathbf{0}_{3 \times 3} & \mathbf{0}_{3 \times 3} & \mathbf{0}_{3 \times 3} & -\frac{1}{\tau} \mathbf{I}_{3 \times 3} \end{bmatrix} \quad (8)$$

where the $[\cdot]^\times$ operator denotes the skew symmetric matrix associated with the cross product of the vector. Next, an approximate discretization of \mathbf{F}_c is found via Euler's method

$$\mathbf{F}_d(\mathbf{x}, \Delta t) = \mathbf{I} + \mathbf{F}_c(\mathbf{x}) \Delta t \quad (9)$$

and the covariance is propagated as

$$\begin{aligned} \mathbf{P}^{k+1} &= \\ &\mathbf{F}_d(\hat{\mathbf{x}}^k, \Delta t) \mathbf{P}^k \mathbf{F}_d(\hat{\mathbf{x}}^k, \Delta t)^T + \mathbf{G}(\hat{\mathbf{x}}^k) \mathbf{Q} \mathbf{G}(\hat{\mathbf{x}}^k)^T \end{aligned} \quad (10)$$

where $\mathbf{Q} := \text{diag}(\sigma_g^2, \sigma_a^2, \sigma_{b_g}^2, \sigma_{b_a}^2)$ and

$$\mathbf{G}(\mathbf{x}) = \begin{bmatrix} \tilde{\mathbf{C}}_{WS} & \mathbf{0}_{3 \times 3} & \mathbf{0}_{3 \times 3} & \mathbf{0}_{3 \times 3} \\ \mathbf{0}_{3 \times 3} & \mathbf{I}_{3 \times 3} & \mathbf{0}_{3 \times 3} & \mathbf{0}_{3 \times 3} \\ \mathbf{0}_{3 \times 3} & \mathbf{0}_{3 \times 3} & \mathbf{I}_{3 \times 3} & \mathbf{0}_{3 \times 3} \\ \mathbf{0}_{3 \times 3} & \mathbf{0}_{3 \times 3} & \mathbf{0}_{3 \times 3} & \mathbf{I}_{3 \times 3} \end{bmatrix} \quad (11)$$

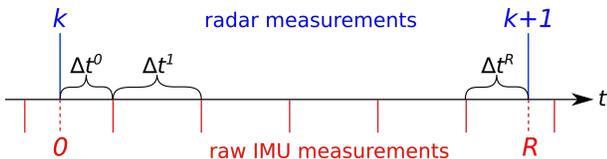


Fig. 4. Different rates of IMU and radar sensor. One IMU term uses all accelerometer and gyro readings between successive radar measurements. Additionally, gyro and accelerometer readings at times 0 and R are interpolated from the adjacent measurements.

2) *IMU Error*: The IMU provides measurements at many times the rate of the radar sensor. Further complicating matters, the IMU and radar measurements are not synchronized. This is illustrated in Fig. 4.

Between radar measurements at timesteps k and $k + 1$ several IMU measurements occur. We interpolate between the IMU measurements immediately before and after the radar measurements to obtain estimated IMU readings that align temporally with the radar measurements. The state at $k + 1$, $\hat{\mathbf{x}}^{k+1}$, is estimated by iteratively applying the propagation equations in Eq. (5) using the Euler forward method. The IMU error is then defined as

$$\mathbf{e}^k(\mathbf{x}^k, \mathbf{x}^{k+1}, \mathbf{z}^k) = \begin{bmatrix} 2 \left[\hat{\mathbf{q}}_{WS}^{k+1} \otimes \mathbf{q}_{WS}^{k+1-1} \right]_{1:3} \\ \hat{\mathbf{v}}_S^{k+1} - \mathbf{v}_S^{k+1} \\ \hat{\mathbf{b}}_g - \mathbf{b}_g \\ \hat{\mathbf{b}}_a - \mathbf{b}_a \end{bmatrix} \quad (12)$$

where the \otimes operator is as defined in [23]. The Jacobian of the error with respect to the state at $k + 1$ is defined as

$$\frac{\partial \mathbf{e}^k}{\partial \hat{\mathbf{x}}^{k+1}} = \begin{bmatrix} \left[\hat{\mathbf{q}}_{WS}^{k+1} \otimes \mathbf{q}_{WS}^{k+1-1} \right]_{1:3,1:3}^\oplus & \mathbf{0}_{3 \times 9} \\ \mathbf{0}_{9 \times 3} & \mathbf{I}_{9 \times 9} \end{bmatrix} \quad (13)$$

The Jacobian with respect to the state at time k is somewhat more difficult to calculate because the IMU error term is found by iteratively applying the IMU integration. Differentiating the error with respect to the state at time k requires use of the chain rule, leading to

$$\frac{\partial \mathbf{e}^k}{\partial \hat{\mathbf{x}}^k} = \left(\prod_{i=0}^R \mathbf{F}_d(\hat{\mathbf{x}}^i, \Delta t^i) \right) \frac{\partial \mathbf{e}^k}{\partial \hat{\mathbf{x}}^{k+1}} \quad (14)$$

IV. EXPERIMENTS

A. Sensor Setup

To demonstrate our method we use an automotive grade radar-on-chip sensor, Texas Instruments AWR1853. The sensor operates in the 77-81 GHz band and identifies targets within a field-of-view (FOV) of approximately ± 75 degrees azimuth and ± 20 degrees elevation. The sensor produces a maximum of 160 targets per measurement at a rate of 10 Hz. For inertial sensing we employ a LORD Microstrain 3DM-GX5-15 IMU. The extrinsic transform between the IMU and radar coordinate frames was manually measured.



Fig. 5. Data collection in the Mars yard at the NASA Jet Propulsion Laboratory.

To estimate groundtruth body-frame velocity of the sensor rig, we use measurements from a Vicon motion capture system and an IMU onboard the vehicle platforms. The Vicon system provides drift-free pose measurements. However, the transform between the Vicon system's coordinate frame and the vehicle's coordinate frame is unknown. Additionally, the Vicon measurements are subject to both noise and communication latency between the Vicon system and the host system. Thus, IMU measurements are used to estimate the transform between the Vicon coordinate frame and the vehicle's body-frame, the timestamp offsets between the Vicon system and host system, and to smooth noise in the Vicon measurements, similar to [24].

B. Evaluation Procedure

To evaluate our method, we create a new dataset of radar and inertial data using the sensors previously described. In addition to radar and inertial data, we record Vicon data when available for groundtruth. Lastly, for comparison we use the visual inertial odometry (VIO) output from an Intel Realsense T265.

We conducted experiments using the handheld sensor rig shown in Fig. 1. Several of these were done in the steam tunnels beneath Folsom Field at CU Boulder to simulate a subterranean environment, and outdoors in the Mars yard at NASA's Jet Propulsion Laboratory. These environments are pictured in Fig. 2 and Fig. 5, respectively. These runs do not include groundtruth data so the performance of the two methods is compared qualitatively. Here we seek to demonstrate that when VIO works our radar-inertial method's performance is comparable to VIO and when VIO fails our method continues to work.

The steam tunnel runs are between 60 and 120 seconds in length and are conducted in bright and dark conditions. The Mars yard runs were roughly 60 seconds in length. These were done with the sensor rig close to the ground, as would be the case for a planetary rover. Also, the Mars yard experiments were conducted near sunset. This created challenging lighting conditions for VIO with large brightness gradients, shadows, and a very bright sky near the horizon which would make it difficult to obtain properly exposed images.

TABLE I
CONDITIONS IN WHICH EXPERIMENTS WERE RUN. EACH LISTED
EXPERIMENT WAS RUN THREE TIMES.

Location	Motion	Lighting	Platform	Groundtruth
CU Vicon Space	Fast	Bright	Quadrotor	Vicon
CU Vicon Space	Slow	Bright	Quadrotor	Vicon
CU Vicon Space	Fast	Dim	Quadrotor	Vicon
CU Vicon Space	Fast	Dark	Quadrotor	Vicon
CU Steam Tunnels	Slow	Bright	Handheld	None
CU Steam Tunnels	Slow	Dark	Handheld	None
JPL Mars Yard	Slow	Bright	Handheld	None

TABLE II
MEAN AND STD DEVIATION OF RMS VELOCITY ERROR IN M/S FOR
QUADROTOR EXPERIMENTS ($N = 3$)

Lighting Movement		bright slow		bright fast		dim fast		dark fast	
		μ	σ	μ	σ	μ	σ	μ	σ
VIO	x	.10	.025	.20	.021	.20	.030	.56	.012
	y	.26	.012	.19	.014	.23	.021	.31	.018
	z	.17	.021	.31	.030	.31	.009	.40	.013
RI	x	.10	.011	.21	.006	.19	.016	.16	.013
	y	.16	.015	.28	.009	.25	.017	.17	.020
	z	.11	.025	.28	.006	.25	.024	.14	.017

The same sensor suite was also mounted on a quadrotor for experiments in our motion capture space. These quadrotor runs varied between 30 and 60 seconds in length, featured both slow, smooth motions and fast, aggressive motions, and were conducted in bright, dim, and dark conditions. Each of these experiments was run 3 times in the same conditions and over the same path. All of these runs have groundtruth from motion capture. For these runs, we are able to quantify the accuracy of VIO and our radar-inertial method in terms of the root mean squared error (RMSE) between the ego-velocity estimate and the groundtruth:

$$\mathbf{v}_{\text{RMSE}} = \sqrt{\frac{\sum_{i=1}^N (\mathbf{v}_{\text{est}}^i - \mathbf{v}_{\text{gt}}^i)^2}{N}} \quad (15)$$

The runs included in our dataset are summarized below in table I.

V. RESULTS

Table II lists the mean and standard deviation of the RMSE in estimated ego-velocity from radar-inertial (RI) and VIO along all body-frame axes for our quadrotor experiments. From these results, it is clear the accuracy of RI is comparable to VIO when the scene is brightly or moderately lit. However, in dark conditions VIO’s performance deteriorates considerably while RI’s accuracy is unaffected.

For the included plots, VIO estimates are plotted in red, RI in blue, and groundtruth in green. Figure 6 shows the error of velocity estimates from RI and VIO for the quadrotor experiment for which VIO performed best. Fig. 6 shows that even when conditions are optimal for VIO, RI’s accuracy is comparable to VIO. However, if the input radar measurements to our RI system are persistently incorrect

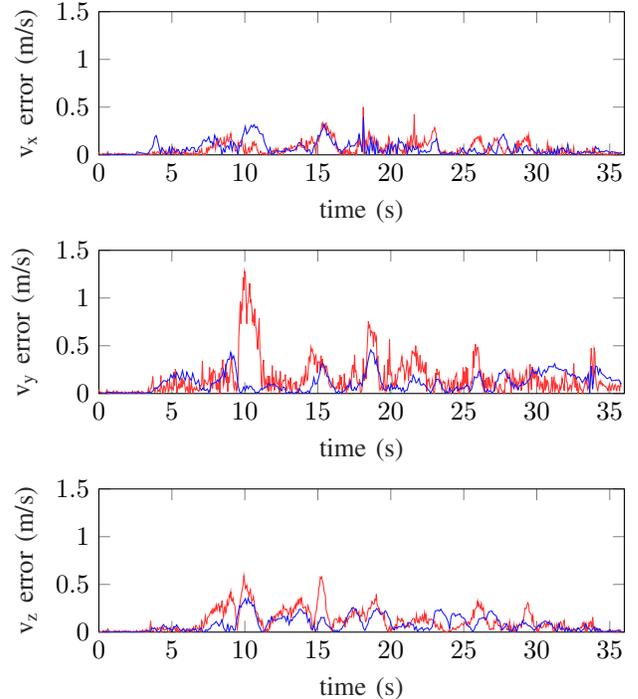


Fig. 6. Body-frame velocity estimate error from our method and VIO for an example run in ideal sensing conditions for VIO. Errors for VIO are plotted in blue and ours are plotted in red.

for multiple seconds then the IMU constraints will not be sufficient to smooth out the noise in the radar measurements. This can be seen around the 10 second mark in the v_y error plot of figure 6. This issue may be addressed in future work by adding a term to the optimization that approximates the influence of measurements that have passed out of the sliding window.

Figure 7 shows the estimated ego-velocity components for an example run conducted in dark conditions in our motion capture space. The velocity estimates from RI track the groundtruth closely, while the estimates from VIO often have large deviations. This demonstrates that the proposed method continues to function normally in dark conditions while VIO’s performance suffers. Additionally, Fig. 8 shows the ego-velocity estimates from RI and VIO taken with our handheld rig in the subterranean environment in dark conditions. RI consistently produces results in the dark conditions, while VIO cuts out completely for about 10 seconds at the 20 second mark. Groundtruth is not available for this experiment, so it is not possible to say which method is more accurate, but Fig. 8 shows that RI continues functioning where VIO fails completely. This behavior was typical of all runs in this experiment.

Figure 9 shows the estimated ego-velocity components for an experiment in the JPL Mars yard. In this experiment, the sensor rig is moved steadily forward in the x direction with small movements in the y and z directions. Subjectively, radar-inertial velocity estimates reflect the platform’s true motion, while VIO is noisy and has difficulty tracking in

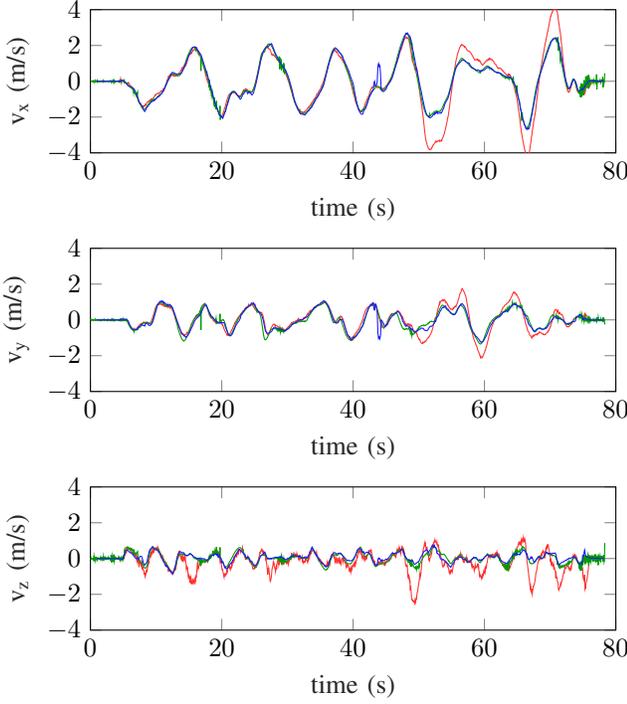


Fig. 7. Estimated and groundtruth velocity components for a run with the quadrotor rig in a motion capture space in dark conditions. VIO estimates are plotted in red, radar-inertial in blue, and groundtruth in green.

the x and z directions. This demonstrates how RI performs well outdoors in conditions that present difficulties for VIO.

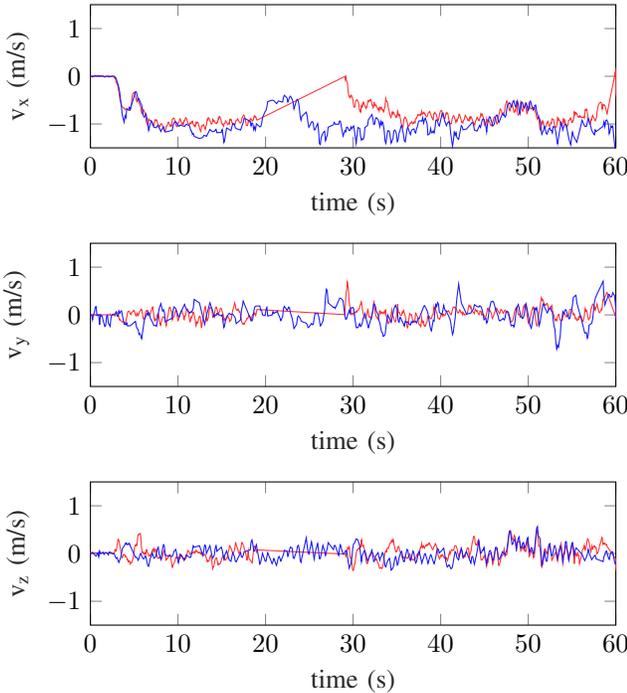


Fig. 8. Estimated body-frame velocity components from RI and VIO taken with the handheld rig in the subterranean test environment in dark conditions. VIO estimates are plotted in red and radar inertial estimates are plotted in blue. Note RI steadily produces estimates throughout the run, while VIO drops out between 20 and 30 seconds.

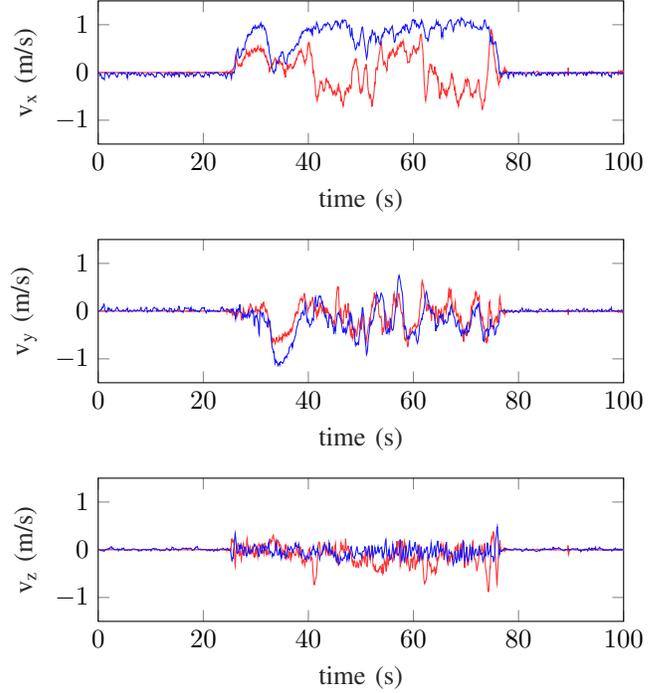


Fig. 9. Plots of the body-frame velocity components estimated by VIO and RI in the Mars yard at the JPL. VIO estimates are plotted in red and RI's estimates are plotted in blue. In this run the sensor rig moved steadily forward in the x direction, with small movements in the y and z directions.

VI. CONCLUSIONS

This work presents a method for estimating the 3D body-frame velocity of a radar-inertial sensor platform. We fuse Doppler velocity measurements from an SoC millimeter wave radar sensor with inertial measurements from an IMU. Radar is invariant to the kinds of perceptually challenging conditions that present problems for vision-based ego-velocity estimation methods. Radar-based methods will fail when a sufficient number of strong radar reflectors are not present in the environment; however, this work demonstrates that even in open outdoor environments such as JPL's Mars Yard a sufficient number of radar targets are detected for the proposed method to be successful. Additionally, the radar-inertial sensor suite is lightweight and has low power requirements making it an attractive alternative for platforms with constraints on their payload and power.

The accuracy of the presented approach is shown through indoor, outdoor and subterranean experiments via comparisons with a motion capture system (indoors) and a commercial VIO system. The resulting experiments demonstrate that the proposed method is comparable to the VIO approach for ego-velocity estimation in conditions favorable to VIO methods, and far exceeds VIO accuracy when conditions deteriorate.

VII. ACKNOWLEDGEMENT

The authors would like to thank Michael Ohradzansky for his help in outfitting the experimental quadrotor vehicle with the visual-radar-inertial sensor suite, and for his time spent piloting the quadrotor for data collections.

REFERENCES

- [1] A. Aghamohammadi, K. Mitchel, and P. Boston, "Robotic exploration of planetary subsurface voids in search for life," in *American Geophysical Society Fall Meeting 2019*. San Francisco, CA: AGU, 2019. [Online]. Available: <https://agu.confex.com/agu/fm19/meetingapp.cgi/Paper/598427>
- [2] S. Khattak, C. Papachristos, and K. Alexis, "Keyframe-based direct thermal-inertial odometry," *CoRR*, vol. abs/1903.00798, 2019. [Online]. Available: <http://arxiv.org/abs/1903.00798>
- [3] C. Stahoviak, "An instantaneous 3d ego-velocity measurement algorithm for frequency modulated continuous wave (fmcw) doppler radar data," 2019.
- [4] J. Dickmann, N. Appenrodt, H. Bloecher, C. Brenk, T. Hackbarth, M. Hahn, J. Klappstein, M. Muntzinger, and A. Sailer, "Radar contribution to highly automated driving," in *2014 44th European Microwave Conference*, Oct 2014, pp. 1715–1718.
- [5] S. H. Cen and P. Newman, "Precise Ego-Motion Estimation with Millimeter-Wave Radar Under Diverse and Challenging Conditions," in *2018 IEEE International Conference on Robotics and Automation (ICRA)*. Brisbane, QLD: IEEE, May 2018, pp. 1–8. [Online]. Available: <https://ieeexplore.ieee.org/document/8460687/>
- [6] —, "Radar-only ego-motion estimation in difficult settings via graph matching," *arXiv:1904.11476 [cs]*, Apr. 2019, arXiv: 1904.11476. [Online]. Available: <http://arxiv.org/abs/1904.11476>
- [7] F. Schuster, M. Wrner, C. G. Keller, M. Haueis, and C. Curio, "Robust localization based on radar signal clustering," June 2016, pp. 839–844.
- [8] F. Schuster, C. G. Keller, M. Rapp, M. Haueis, and C. Curio, "Landmark based radar SLAM using graph optimization," in *2016 IEEE 19th International Conference on Intelligent Transportation Systems (ITSC)*, Nov. 2016, pp. 2559–2564.
- [9] D. Vivet, P. Checchin, and R. Chapuis, "Localization and Mapping Using Only a Rotating FMCW Radar Sensor," *Sensors (Basel, Switzerland)*, vol. 13, no. 4, pp. 4527–4552, Apr. 2013. [Online]. Available: <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC3673098/>
- [10] N. Michael, S. Shen, K. Mohta, V. Kumar, K. Nagatani, Y. Okada, S. Kiribayashi, K. Otake, K. Yoshida, K. Ohno, E. Takeuchi, and S. Tadokoro, "Collaborative mapping of an earthquake-damaged building via ground and aerial robots," *Journal of Field Robotics*, vol. 29, no. 5, pp. 832–841, 2012.
- [11] S. Shen, Y. Mulgaonkar, N. Michael, and V. S. A. Kumar, "Vision-based state estimation and trajectory control towards high-speed flight with a quadrotor," in *Robotics: Science and Systems*, 2013.
- [12] T. Tomic, K. Schmid, P. Lutz, A. Domel, M. Kassecker, E. Mair, I. L. Grix, F. Ruess, M. Suppa, and D. Burschka, "Toward a fully autonomous uav: Research platform for indoor and outdoor urban search and rescue," *IEEE Robotics Automation Magazine*, vol. 19, no. 3, pp. 46–56, Sep. 2012.
- [13] A. Santamaria-Navarro, J. Solà, and J. Andrade-Cetto, "High-frequency mav state estimation using low-cost inertial and optical flow measurement units," in *Intelligent Robots and Systems (IROS), 2015 IEEE/RSJ International Conference on*, Sept 2015, pp. 1864–1871.
- [14] A. Santamaria-Navarro, G. Loianno, J. Solà, V. Kumar, and J. Andrade-Cetto, "Autonomous navigation of micro aerial vehicles using high-rate and low-cost sensors," *Autonomous Robots*, vol. 42, no. 6, pp. 1263–1280, Aug 2018.
- [15] J. Zhang and S. Singh, "Loam: Lidar odometry and mapping in real-time," in *Robotics: Science and Systems*, 2014.
- [16] P. Geneva, K. Eickenhoff, Y. Yang, and G. Huang, "Lips: Lidar-inertial 3d plane slam," in *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Oct 2018, pp. 123–130.
- [17] F. R. Kschischang, B. J. Frey, and H. Loeliger, "Factor graphs and the sum-product algorithm," *IEEE Transactions on Information Theory*, vol. 47, no. 2, pp. 498–519, Feb. 2001.
- [18] R. Mur-Artal, J. M. M. Montiel, and J. D. Tardós, "Orb-slam: A versatile and accurate monocular slam system," *IEEE Transactions on Robotics*, vol. 31, pp. 1147–1163, 2015.
- [19] R. Mur-Artal and J. D. Tardós, "Orb-slam2: An open-source slam system for monocular, stereo, and rgb-d cameras," *IEEE Transactions on Robotics*, vol. 33, pp. 1255–1262, 2017.
- [20] S. Leutenegger, S. Lynen, M. Bosse, R. Siegwart, and P. Furgale, "Keyframe-based visualinertial odometry using nonlinear optimization," *The International Journal of Robotics Research*,

- vol. 34, no. 3, pp. 314–334, Mar. 2015. [Online]. Available: <https://doi.org/10.1177/0278364914554813>
- [21] G. Sibley, “A Sliding Window Filter for SLAM,” p. 17.
- [22] A. I. Mourikis and S. I. Roumeliotis, “A Multi-State Constraint Kalman Filter for Vision-aided Inertial Navigation,” in *Proceedings 2007 IEEE International Conference on Robotics and Automation*. Rome, Italy: IEEE, Apr. 2007, pp. 3565–3572. [Online]. Available: <http://ieeexplore.ieee.org/document/4209642/>
- [23] T. Barfoot, J. R. Forbes, and P. Furgale, “Pose estimation using linearized rotations and quaternion algebra,” *Acta Astronautica*, vol. 68, pp. 101–112, Feb. 2011.
- [24] M. Burri, J. Nikolic, P. Gohl, T. Schneider, J. Rehder, S. Omari, M. W. Achtelik, and R. Siegwart, “The euroc micro aerial vehicle datasets,” *The International Journal of Robotics Research*, 2016. [Online]. Available: <http://ijr.sagepub.com/content/early/2016/01/21/0278364915620033.abstract>