

QVIO2: Quantized MAP-based Visual-Inertial Odometry

Yuxiang Peng, Chuchu Chen, and Guoquan Huang

Abstract—Energy-efficient visual-inertial motion tracking on SWAP-constrained edge devices (e.g., drones and AR glasses) is essential but challenging. Our previous work [1] introduced the first-of-its-kind quantized visual-inertial odometry (QVIO), utilizing either raw measurement quantization (zQVIO) or single-bit residual quantization (rQVIO). While QVIO has demonstrated significant data transfer reduction with competitive performance, it has limitations. Specifically, zQVIO directly quantizes raw measurements into multi-bit values, while requiring the ad-hoc inflation of measurement noise to account for quantization errors. On the other hand, rQVIO is limited to single-bit measurement with certain accuracy loss. This work introduces QVIO2 to address these issues. The proposed QVIO2 improves data quantization strategies and derives a Maximum A Posteriori (MAP) quantized estimator that rigorously handles both multi-bit and single-bit, raw and residual quantized measurements in a unified manner. These improvements lead to more communication-efficient and accurate systems. Additionally, we optimize the communication protocol to further reduce data transfer by eliminating unnecessary transmissions. Extensive numerical and experimental results demonstrate reduced communication requirements and improved accuracy. Compared to the previous QVIO system, zQVIO2 achieves the same accuracy with a 30% reduction in data transfer, while rQVIO2 improves accuracy without increasing data communication. In real-world scenarios, our new zQVIO2 and rQVIO2 have demonstrated nearly no accuracy loss with only 4.6 bits and 3.5 bits of data communication, achieving compression rates of $7\times$ and $9.1\times$.

I. INTRODUCTION AND RELATED WORK

Next-generation edge devices like AR/VR wearables require real-time processing of heterogeneous data for tasks such as eye tracking, gesture recognition, motion detection, speech processing, and head movement detection, necessitating substantial computational tasks to be handled on-device. Motion tracking, critical for these devices, is typically achieved through visual-inertial odometry (VIO) [2]–[9]. However, building robust VIO on these small form-factor platforms is challenging due to strict size, weight, and power (SWAP) constraints, with the primary difficulty often arising from data management rather than computation. For example, in SLAM and hand-tracking modules of Meta XR wearable devices, the major energy consumption is data access in RAM [10].

To reduce the data transfer, different on-sensor computing architecture has been presented. Gome et al. [11] introduced a distributed on-sensor compute architecture by performing initial image processing on the sensor and then sending pre-processed image data to the aggregator for further processing through energy-intensive MIPI interfaces, which dramatically minimizes data and energy consumption. However, challenges remain in efficiently handling the high volume data

demands of tasks like visual-inertial estimation. Moreover, in applications like MAVs where visual processing and state estimation are performed separately in SoC and MCU, the bandwidth-limited UART communication interface between two chips makes the challenge even more severe.

Data quantization, which reduces data size and thus communication, has been widely applied in various fields. For instance, it is commonly used to reduce the computational time and memory consumption of neural networks by employing low-bit representations for weights and activations [12]. Previous studies have demonstrated that networks can be quantized to 4-bit weights with minimal loss in accuracy [13]. In SLAM and robotics, quantization has been used to compress images [14]–[16] and reduce memory size. For example, Navion [17] combines lossy image compression and block-wise quantization, achieving a $4.4\times$ memory reduction and a $4.9\times$ power reduction with minimal computational overhead, enabling VIO on nano drones. Vector quantization (VQ) is also used to compress data and segment motion, supporting RGBD-SLAM in dynamic environments [18]. It has also been used to compress the environmental map (i.e., the point cloud) [19] to enable efficient sharing of the environmental information in the multi-robot systems [20] or life-long localization [21]. Quantization of descriptors [22]–[25], is another practice in visual SLAM, reducing memory overhead while preserving matching quality.

From the perspective of estimation theory, using quantized measurements requires the estimator to be derived with mathematical rigor. For example, SOI-KF [26] employs a single-bit per observation in a Kalman filter (KF) setup for wireless sensor networks, achieving performance comparable to traditional KFs while significantly reducing communication complexity and power consumption. It models the quantized residual as a Gaussian tail distribution and approximates the minimum mean squared error (MMSE) estimator. This approach has been extended to multi-bit quantization techniques for cooperative localization (CL) in multi-robot systems, allowing effective estimation under strict communication constraints [27]–[29].

Despite the prevalence of quantization, the first of its kind quantized VIO (QVIO) was introduced in our previous work [1], which preprocesses images at the sensor level and transmits only quantized measurements or residuals to the host for visual-inertial estimation. Specifically, we developed two QVIO estimators. The first, measurement-quantized VIO (zQVIO), directly quantizes visual measurements into a fixed number of bits and applies an EKF update for fusion, significantly reducing data transfer between processors. However, zQVIO lacks a performance guarantee, relying on trial-and-error for bit selection and additional noise inflation. As such, we also propose residual-quantized VIO (rQVIO) that quantizes residuals on the co-processor and transmits only a single bit per residual to the host for estimation. Although the communication protocol is more complicated, the estimation

This work was partially supported by the University of Delaware (UD) College of Engineering, NSF (SCH-2014264, IIS-2410019), Google ARCore, and Meta Reality Labs.

The authors are with the Robot Perception and Navigation Group (RPNG), University of Delaware, Newark, DE 19716, USA. Email: {yxpeng, ccchu, ghuang}@udel.edu

problem can be framed as a MAP estimation problem, solved recursively using an EKF-like update. However, the MAP problem formulation is limited to single-bit measurement with unavoidable accuracy loss. While our baseline QVIO [1] performed well, in this work, we further improve these estimators and design the quantized MAP-based zQVIO2 and rQVIO2. Specifically, zQVIO2 maintains the same accuracy while reducing data transfer by approximately 30%, while rQVIO2 delivers improved accuracy without increasing data communication. The contributions of this work can be summarized as follows:

- We design the improved quantization strategies. For zQVIO, we propose the differential measurement quantization to reduce communication bandwidth without extra errors. For rQVIO, we extend from single-bit to multi-bit measurements for more accurate estimation.
- We derive a MAP quantized estimator that rigorously processes single- and multi-bit measurements for both measurement and residual quantization strategies.
- We optimize the communication protocol to reduce state update communication overhead for rQVIO to mitigate the increase of bits for residual measurements. Extensive numerical and experimental results show reduced communication requirements and improved accuracy.

II. QUANTIZED MAP ESTIMATION

In this section, we outline the design of a quantized estimation problem by addressing two key challenges: *how to quantize measurements to reduce data transfer without adding error*, and *how to estimate states using quantized measurements rigorously*. To simplify the discussion, we consider the general nonlinear measurement function in the following. Given the measurement $\mathbf{z} = [\dots z_i \dots]$, the measurement function and its linearization are ¹:

$$\mathbf{z} = \mathbf{h}(\mathbf{x}) + \mathbf{n} \Rightarrow \mathbf{r} = \mathbf{z} - \mathbf{h}(\hat{\mathbf{x}}) \simeq \mathbf{H}\hat{\mathbf{x}} + \mathbf{n} \quad (1)$$

$$\mathbf{r} = [\dots r_i \dots]^\top, \mathbf{H} = [\dots \mathbf{H}_i \dots]^\top \quad (2)$$

where \mathbf{r} and \mathbf{H} are the residual and measurement Jacobian; \mathbf{n} is the white Gaussian measurement noise $\mathbf{n} \sim \mathcal{N}(\mathbf{0}, \mathbf{R})$.

A. Quantization Strategies

We previously introduced two quantization strategies [1]: measurement quantization (zQ), which directly quantizes raw measurements into a specified number of bits, and residual quantization (rQ), which reduces measurement residuals to a single bit. In the following, we outline how we can improve these two. At a high level, the quantized measurement can be expressed as:

$$\mathbf{b}(\mathbf{z}) = [\dots \mathbf{b}_i \dots], \mathbf{b}_i = [\dots b_i^j \dots] \quad (3)$$

where \mathbf{b}_i is the multi-bit quantized measurement and b_i^j is a 1-bit scalar measurement.

1) *Differential Measurement Quantization*: Naively, we can quantize the raw measurement directly to $\mathbf{b}(\mathbf{z})$ to multi-bit with the residual derived as:

$$\mathbf{r}^m := \mathbf{b}(\mathbf{z}) - \mathbf{h}(\hat{\mathbf{x}}) \simeq \mathbf{H}\hat{\mathbf{x}} + \mathbf{n}_b \quad (4)$$

where \mathbf{n}_b denotes noise after quantization, which is a mixture of original sensor measurement noise and quantization




Fig. 1: Illustration of the proposed multi-bit residual quantization: When $K=2$ bits are used, the residual range is divided into $2^K = 4$ intervals.

noise [1]. We now present a novel differential quantization method. Assume we have in total $k = 3$ measurements, $\mathbf{Z} = [\mathbf{z}_1 \dots \mathbf{z}_3]$. For the first measurement \mathbf{z}_1 , the same quantization strategy as our previous work [1] is applied, where $\mathbf{B}_1 := \mathbf{B}(\mathbf{z}_1)$. Then, for the following sequential measurement, we adopt the differential quantization scheme. For example, the communicated quantized measurement for \mathbf{z}_3 is formulated as $\mathbf{B}_3 := \mathbf{B}(\mathbf{z}_3 - \mathbf{B}(\mathbf{z}_2 - \mathbf{B}(\mathbf{z}_1)))$. The quantized measurement \mathbf{z}_3^q can be recovered as:

$$\mathbf{z}_3^q := \sum_{k=1}^3 \mathbf{B}_k \quad (5)$$

which can be used to triangulate features and calculate residuals. This method reduces communication bandwidth by focusing on the difference between measurements, which is generally much smaller than the raw data. Additionally, it avoids introducing extra quantization errors because the point where the difference is calculated is precisely known, and only the difference itself is quantized.

2) *Multi-bit Residual Quantization*: Inspired by [31], [32], we quantize the measurement residual as follows: Instead of only using 1-bit to quantize residual (i.e., 0 indicates negative and 1 is positive), we now allow for multi-bits to preserve more information. Given K bits, the range of $(-\infty, +\infty)$ can be divided into 2^K separate intervals. K bits can encode the interval where the residual lies, such as $t_i < r < t_j$. As illustrated in Figure 1, where $K = 2$, the range is divided into 4 intervals. If the quantized measurement for one scalar whitened residual is $b_1 = 1, b_2 = 0$, then this indicates the residual is in the range of $(0, t_1)$.

B. Quantized State Estimation

We now analytically derive the quantized MAP estimator that supports both the preceding rQ and zQ quantization schemes. For the sake of clarity, we assume a scalar measurement $h(\cdot)$ in our explanation. Given the quantized measurement (3), we formulate the quantized estimation as the following MAP optimization problem [32]:

$$\begin{aligned} \hat{\mathbf{x}} &= \arg \max_{\mathbf{x}} p(\mathbf{x}|\mathbf{b}(\mathbf{z})) \\ &= \arg \max_{\mathbf{x}} p(\mathbf{x}) \prod_{i=1}^m p(b_i^1, b_i^2, \dots, b_i^j | \mathbf{x}) \end{aligned} \quad (6)$$

where m is the number of measurements and j is the number of bits used for this measurement. Intuitively, for both rQ and zQ methods, those bit indicates the range of the measurements. The range of the rQ is pre-defined thresholds for the residuals, while the range of the zQ is the resolution defined by the number of bits.

Interestingly, for both rQ and zQ methods, we can derive $p(b_i^1, b_i^2, \dots, b_i^j | \mathbf{x})$ to be expressed in terms of the Q-function, and the MAP (6) can be written as:

$$\hat{\mathbf{x}} = \arg \max_{\mathbf{x}} p(\mathbf{x}|\mathbf{b}(\mathbf{z}))$$

¹Note that throughout the paper $\hat{\mathbf{x}}$ is used to denote the *current* best estimate of a random variable \mathbf{x} with $\tilde{\mathbf{x}} = \mathbf{x} \boxminus \hat{\mathbf{x}}$ denotes the error state. The “ \boxplus ” and “ \boxminus ” operations map elements to and from a given manifold and equate to simple “+” and “-” for vector variables [30].

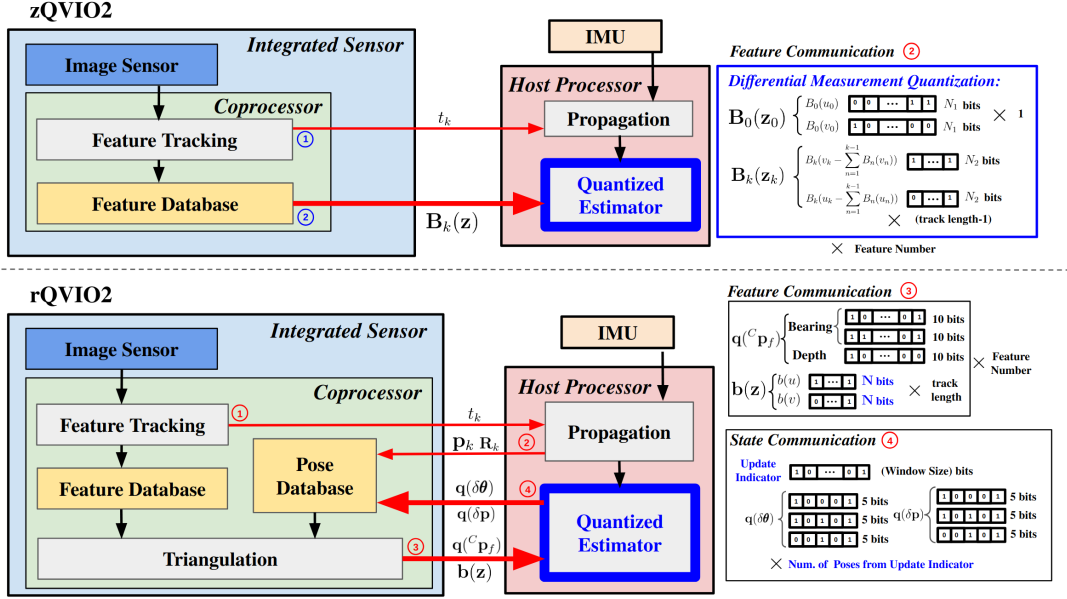


Fig. 2: System architecture of the measurement-quantized VIO 2.0 (zQVIO2) and residual-quantized VIO 2.0 (rQVIO2). Note that the new unified quantized MAP estimator is highlighted in blue, as compared to the baseline QVIO [1].

$$= \arg \max p(\mathbf{x}) \prod_{i=1}^m P\{t_i^1 < \frac{z_i^q - h(\mathbf{x}) - n}{\sigma} < t_i^2 | \mathbf{x}\} \quad (7)$$

$$= \arg \max p(\mathbf{x}) \prod_{i=1}^m (\mathbf{Q}(\chi_i^1) - \mathbf{Q}(\chi_i^2)) \quad (8)$$

where σ is a scalar of noise standard deviation that normalizes the measurement noise. t_i^1 and t_i^2 in rQVIO are the pre-defined thresholds where the residuals lie, while in zQVIO they are the minus and half quantization resolution $-\frac{r_q}{2}$ and $\frac{r_q}{2}$. Intuitively, if the quantization resolution is r_q , then given quantized measurement z_i^q , the raw measurement z_i should fall in the range of $(z_i^q - \frac{r_q}{2}, z_i^q + \frac{r_q}{2})$, thus $(z_i^q - z_i)$ should fall in the range of $(-\frac{r_q}{2}, \frac{r_q}{2})$. Then for the zQ-based estimator, we can derive:

$$z \chi_i^1 = \frac{1}{\sigma} (z_i^q - h(\mathbf{x}) - \frac{r_q}{2}) \quad (9)$$

$$z \chi_i^2 = \frac{1}{\sigma} (z_i^q - h(\mathbf{x}) + \frac{r_q}{2}) \quad (10)$$

The rQ-based estimator is written as:

$$r \chi_i^1 = \frac{1}{\sigma} (h(\bar{\mathbf{x}}) - h(\mathbf{x})) + t_1 \quad (11)$$

$$r \chi_i^2 = \frac{1}{\sigma} (h(\bar{\mathbf{x}}) - h(\mathbf{x})) + t_2 \quad (12)$$

where $\bar{\mathbf{x}}$ is the quantized state in the coprocessor, which is also known in the host processor and can be taken as pseudo quantized measurements z_i^q . Detailed derivation can be found in Appendix I. Take logarithm on Eq. (8), the MAP problem is equivalent to:

$$\hat{\mathbf{x}} = \arg \max \left[\|\mathbf{x} \ominus \bar{\mathbf{x}}\|_{\mathbf{P}}^2 + \sum_{i=1}^m 2 \ln (\mathbf{Q}(\chi_i^1) - \mathbf{Q}(\chi_i^2)) \right]$$

where we have assumed a Gaussian prior. To solve for the above MAP estimate, we can adopt the standard EKF equations to approximately, recursively and efficiently compute the estimate and covariance (e.g., see [1] for details).

III. QVIO2: QUANTIZED-VIO 2.0

In the VIO system, the state vector \mathbf{x} at time t_k consists of the current navigation states \mathbf{x}_{I_k} , and a set of historical IMU pose clones \mathbf{x}_C (see [33]):

$$\mathbf{x} = [\mathbf{x}_{I_k}^T \ \mathbf{x}_C^T]^T, \quad \mathbf{x}_C = [\mathbf{x}_{T_k}^T \ \dots \ \mathbf{x}_{T_{k-c}}^T]^T \quad (13)$$

$\mathbf{x}_{I_k} = [{}^I_G \bar{q}^T \ G \mathbf{p}_{I_k}^T \ G \mathbf{v}_{I_k}^T \ \mathbf{b}_g^T \ \mathbf{b}_a^T]^T$ (14) where ${}^I_G \bar{q}$ is the unit quaternion corresponding to the rotation matrix ${}^I_G \mathbf{R}$ that represents the rotation from the global frame $\{G\}$ to the IMU frame $\{I\}$; $G \mathbf{p}_{I_k}$, $G \mathbf{v}_{I_k}$ are the IMU position, velocity; \mathbf{b}_g and \mathbf{b}_a are the gyroscope and accelerometer biases; $\mathbf{x}_{T_i} = [{}^I_G \bar{q}^T \ G \mathbf{p}_{T_i}^T]^T$.

As in the standard MSCKF-based VIO [34], we propagate the state over time based on the nonlinear IMU kinematics:

$$\mathbf{x}_{I_{k+1}} = \mathbf{f}_I(\mathbf{x}_{I_k}, I \mathbf{a}_k, I \boldsymbol{\omega}_k, \mathbf{n}_I) \quad (15)$$

where \mathbf{n}_I consists the noises for the measurement and model. We then linearize this kinematic model to propagate the state estimate and covariance.

Assume the camera measures a feature \mathbf{f} at timestamp t_k , where \mathbf{f} denotes its 3D position in the global frame. The corresponding bearing measurement is given by:

$$\mathbf{z}_k = \mathbf{h}(\mathbf{x}_k) + \mathbf{n}_k \simeq \mathbf{H}_x^* \tilde{\mathbf{x}}_{I_k} + \mathbf{H}_f^* \mathbf{f} + \mathbf{n}_k \quad (16)$$

where $\mathbf{z}_k = [u, v]^T$ is the raw uv pixel coordinate; \mathbf{n}_k is the zero-mean white Gaussian raw pixel noise. We refer to [33], [34] for more details about MSCKF update.

In what follows, we explain how to leverage the new quantization strategies and the quantized MAP estimator introduced in Section II to develop the Quantized VIO 2.0.

A. zQVIO2: Measurement Quantization VIO 2.0

As described in Section II-A.1, zQ method directly quantizes the raw pixel measurements using a given number of bits. This led to a simple and clean communication protocol between the co-processor and host computer, as shown in Figure 2 (top). The initial preprocessing of visual and IMU data remains the same. In the following steps, only the quantized measurements are transmitted to the host processor,

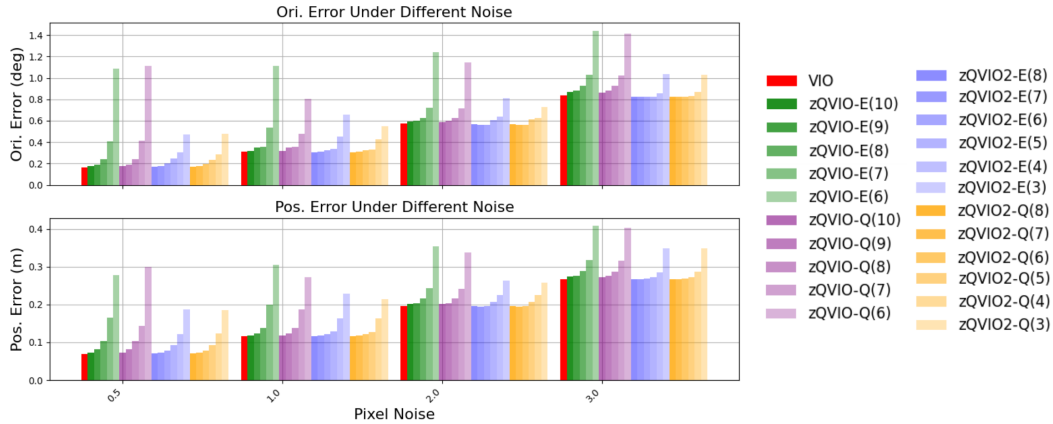


Fig. 3: Estimation error comparison across varying measurement noise levels and bit measurement for zQVIO.




Fig. 4: Estimation error comparison across varying measurement noise levels and bit measurement for rQVIO.

where both feature triangulation and state estimation are carried out using these quantized inputs. Compared to the baseline zQVIO [1], two key innovations are introduced:

- We employ the new *differential quantization* method, instead of simply quantizing the raw measurements into multiple bits, to enable further data compression without additional errors.
- The baseline zQVIO blindly performs EKF updates with the quantized measurements and was unable to adequately compensate for the potentially significant quantization errors in practice. Although we modelled these errors as additional noise using the Kullback–Leibler divergence (KLD), this approach was still suboptimal. In contrast, the new zQVIO2 employs a quantized MAP estimator with mathematical rigor, eliminating the need to inflate the noise to account for quantization error.

B. rQVIO2: Residual Quantization VIO 2.0

As shown in Figure 2 (bottom), rQVIO 2.0 has a more complex architecture. First, the co-processor performs feature detection and tracking, sending the timestamp to the host processor. The host uses IMU measurements to propagate the state and returns the IMU pose to assist with triangulation. Next, the co-processor triangulates features into 3D and sends the quantized initial guess for 3D features and its quantized residual measurements back to the host for state update. Finally, the quantized state is updated, and the corrected IMU poses are sent back to the co-processor. A detailed explanation of this architecture can be found [1]. It is important to highlight here the key innovations. First, compared to [1], we design multi-bit quantized measurements with rigorous mathematical derivations, as explained in Section II-B. While the baseline rQVIO demonstrated that 1-bit measurements delivered impressive performance, accuracy loss is inevitable with significant information loss. The new estimator allows users to further enhance accuracy

by enabling higher-bit quantized residuals. Moreover, we develop a novel update indicator as shown in Figure 2. A significant data transfer is typically required to send the state correction term back to the co-processor and update the IMU pose. However, we found this transfer is not always necessary, especially when the correction is minimal. By setting a threshold for the state correction, the update indicator checks whether the correction term needs to be transferred and reduces unnecessary data communication.

IV. NUMERICAL STUDY

We employ the OpenVINS [34] to produce realistic bearings and inertial measurements and implement rQVIO2 and zQVIO2. In each timestamp, 200 bearing measurements are generated. Regarding the estimator configuration, we maintain a window size of 15. All features are MSCKF features and will be used to do updates when they lose tracking or reach the maximum window size. In the following section, VIO serves as the baseline without any measurement quantization. zQVIO and rQVIO refer to the original systems, while zQVIO2 and rQVIO2 represent the newly proposed systems. The number indicates the bit design of the quantized measurement (e.g., zQVIO2(3) means the measurement is quantized to 3 bits).

A. zQVIO2: Better Accuracy without Extra Inflation

To demonstrate the improvements of zQVIO2, we conducted a comprehensive numerical study to evaluate the effectiveness of two key innovations: the differential measurement quantization method and the use of a quantized estimator, as introduced in Section II-A.1. A summary of the

TABLE I: zQVIO System Setups for Evaluation

	zQVIO-E	zQVIO-Q	zQVIO2-E	zQVIO2-Q
Differential Meas.	✗	✗	✓	✓
Quantized Estimator	✗	✓	✗	✓

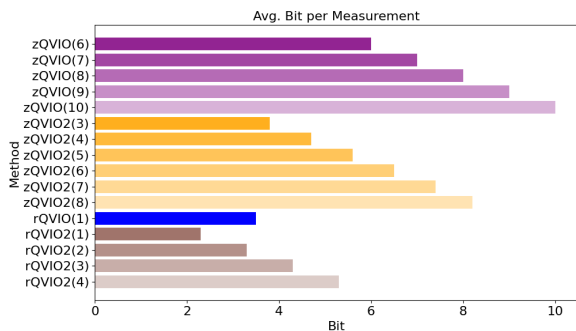


Fig. 5: Avg. number of bits per measurement in 1-pixel noise.

estimators is provided in Table I. Specifically, “-E” refers to using the standard EKF as the estimator, as we discussed in our previous work [1], where we showed the need to inflate noise to compensate for quantization errors. In contrast, “-Q” denotes the use of the quantized estimator, as highlighted in Section II-A.1. We varied the number of bits used to quantize raw measurements and compared the performance of zQVIO2 against both the non-quantized system (VIO) and the previous zQVIO, as illustrated in Figure 3. We also report the communication burden for each estimator, as reported in Figure 5. We calculate the total number of bits communicated and then divide it by the total number of measurements to get the average bit required for each measurement.

To demonstrate the effectiveness of using differential quantized measurements, we first compare zQVIO-Q (purple) and zQVIO2-Q (orange) in Figure 3. Both systems use the same quantized estimator, with the only difference being that zQVIO2 employs differential quantization, while zQVIO does not. Clearly, zQVIO2-Q achieves better performance with fewer bits. For instance, zQVIO-Q(6) nearly doubles the error of zQVIO2-Q(6). A similar trend can be seen when comparing zQVIO-E and zQVIO2-E. Additionally, in Figure 5, it is evident that zQVIO2 (orange) achieves a smaller average bit count per measurement to achieve similar accuracy compared with zQVIO (purple). Through this fair comparison, we clearly demonstrate the benefit of using differential quantized measurements. Next, we show that the development of the quantized estimator not only avoids the need for noise inflation but also models the measurements more accurately, improving estimation performance. This is evident in the comparison between zQVIO2-E (blue) and zQVIO2-Q (orange) in Figure 3. Both systems use the same measurements, but zQVIO2-Q achieves better accuracy due to the quantized estimator used. A similar observation can be made when comparing zQVIO-E (green) and zQVIO-Q (purple), further highlighting that while noise inflation can account for quantization errors, the quantized estimator offers superior performance thanks to its mathematical rigor. Finally, the system that leverages both the differential quantized measurement method and the quantized estimator, zQVIO2-Q (orange), achieves the best overall performance. It only requires 5-6 bits to reach a similar performance to VIO, cutting the number of bits transferred in half (from 9-10 bits to 5-6) compared with previous work. Notably, when the noise is large (e.g., 3-pixel noise), zQVIO2 only begins to show accuracy loss at zQVIO2(4), as sensor noise dominates over quantization noise.

TABLE II: Average Absolute Trajectory Error (ATE) in degrees/meters for different estimators with 1 pixel noise.

Method	zQVIO2(5)	zQVIO2(6)	rQVIO2(2)	rQVIO2(3)
ATE (deg/m)	0.33 / 0.13	0.33 / 0.12	0.35 / 0.13	0.32 / 0.12
Bit/Meas.	5.6	6.5	3.3	4.3

B. rQVIO2: Better Accuracy with Lower Data Transfer

We then perform the numerical study for the rQVIO2. The first improvement is rQVIO2 allows multi-bit residual quantization which can further reduce the accuracy drop from the data quantization [see Section III-B]. To demonstrate this, we compare rQVIO2 with different number of bits used for quantized measurement (i.e., $\mathbf{b}(z)$ in Figure 4) with different measurement noise. The results clearly show that, compared to the original rQVIO, which only allowed single-bit quantization, rQVIO2 significantly reduces the accuracy drop as the number of bits increases due to its ability to handle multi-bit measurements. Notably, rQVIO2 achieves comparable performance to baseline VIO with only 3-4 bits of quantized measurements. Another key improvement comes from the design of the update indicator, as described in Section III-B, which further reduces communication overhead. To demonstrate this, we examine the data communication burden for rQVIO2, as shown in Figure 5. When using only 1-bit quantized measurements, we compare rQVIO(1) and rQVIO2(1), where the latter utilizes the update indicator while rQVIO does not. It is evident that rQVIO2 reduces the average number of bits transferred by approximately 30%. Moreover, rQVIO2(2) uses almost the same communication but achieves much better accuracy than rQVIO(1).

C. Discussion

To compare the performance between zQVIO2 and rQVIO2, we present estimators with similar estimation performance (i.e., comparable ATE results) and evaluate their data transfer requirements, as shown in Table II. From the results, rQVIO2(2) and rQVIO2(3) achieve similar accuracy as zQVIO2(5) and zQVIO2(6), while the average bit required is 2 bits less, shows rQVIO2 is prevailing in achieving the same level of accuracy in more compressed bandwidth requirements compared with zQVIO2. However, we should note that rQVIO2 requires a more complicated co-processor design, where it requires the ability to perform feature triangulation. Moreover, a more complex communication protocol is required for rQVIO2, as state update communication is necessary, which might degrade system robustness [See Figure 2]. For computation complexity, similar to our previous work [1], both methods do not change measurement size. They only require calculating quantized residual and the scaling ratio for each measurement, thus incurring negligible computation overhead.

V. REAL-WORLD EXPERIMENTS

We further evaluate the two proposed quantized estimators on the EuRoC MAV dataset [35] using the left camera. We extracted 300 sparse point features and managed 15 clones, processing them with MSCKF features when they lost track or exceeded the window size.

The communication setup for state update is similar to our previous work [1] for rQVIO2, where 30 bits are used to communicate one 3D feature and 15 bits are used to

TABLE III: Average ATE in degrees/meters and the number of bits per measurement for different estimators.

	V101	V102	V103	V201	V202	V203	MH01	MH02	MH03	MH04	MH05	Bit/Meas.
VIO	0.57 / 0.07	1.96 / 0.06	1.92 / 0.06	0.86 / 0.08	1.29 / 0.08	1.35 / 0.15	1.98 / 0.19	1.19 / 0.27	1.31 / 0.17	0.91 / 0.25	0.61 / 0.34	32.0
zQVIO(8)	0.56 / 0.08	1.94 / 0.06	1.90 / 0.07	0.80 / 0.10	1.38 / 0.08	1.54 / 0.13	2.24 / 0.19	1.15 / 0.26	1.46 / 0.20	0.91 / 0.27	0.67 / 0.37	8.0
zQVIO(7)	0.66 / 0.10	2.01 / 0.06	1.89 / 0.08	1.06 / 0.10	1.52 / 0.09	1.49 / 0.14	2.44 / 0.20	1.17 / 0.28	1.54 / 0.21	1.13 / 0.36	0.68 / 0.38	7.0
zQVIO(6)	1.03 / 0.13	2.09 / 0.10	2.72 / 0.14	1.31 / 0.11	1.79 / 0.12	1.74 / 0.15	2.44 / 0.29	1.48 / 0.25	1.98 / 0.39	1.54 / 0.48	0.57 / 0.49	6.0
zQVIO2(5)	0.57 / 0.07	2.01 / 0.06	1.88 / 0.07	0.93 / 0.06	1.30 / 0.08	1.14 / 0.15	1.94 / 0.18	1.18 / 0.27	1.29 / 0.18	0.95 / 0.26	0.64 / 0.34	5.5
zQVIO2(4)	0.58 / 0.07	1.97 / 0.06	1.66 / 0.07	0.89 / 0.06	1.24 / 0.08	1.24 / 0.19	1.99 / 0.20	1.23 / 0.29	1.30 / 0.19	0.90 / 0.25	0.65 / 0.37	4.6
zQVIO2(3)	0.58 / 0.08	1.95 / 0.07	1.91 / 0.09	0.96 / 0.06	1.33 / 0.08	1.15 / 0.20	1.90 / 0.19	1.28 / 0.31	1.43 / 0.22	1.03 / 0.30	0.70 / 0.41	3.7
rQVIO(1)	0.87 / 0.09	1.82 / 0.06	1.81 / 0.06	0.95 / 0.13	1.19 / 0.09	1.34 / 0.18	1.83 / 0.16	1.14 / 0.20	1.13 / 0.15	0.73 / 0.20	0.81 / 0.33	3.4
rQVIO2(1)	1.14 / 0.12	1.80 / 0.07	1.88 / 0.07	0.97 / 0.15	1.22 / 0.09	1.40 / 0.17	2.25 / 0.24	1.11 / 0.23	1.40 / 0.21	0.77 / 0.19	0.71 / 0.28	2.6
rQVIO2(2)	0.59 / 0.06	1.89 / 0.06	1.88 / 0.06	0.79 / 0.12	1.27 / 0.08	1.30 / 0.17	1.78 / 0.17	1.01 / 0.24	1.26 / 0.18	0.84 / 0.21	0.65 / 0.34	3.5
rQVIO2(3)	0.57 / 0.07	1.92 / 0.06	1.90 / 0.06	0.87 / 0.11	1.33 / 0.08	1.40 / 0.13	1.73 / 0.14	1.07 / 0.26	1.25 / 0.16	0.86 / 0.23	0.63 / 0.34	4.5

communicate state update for rotation, while for position update, 15 bits are used instead of 9 to compensate for reduced state update communication frequency. The new clones are communicated using 16-bit half float for each element. The state update threshold is set to 0.01 m for position and 0.001 rad for rotation. As before, we apply a naive χ^2 test, ignoring pose uncertainty and rejecting measurements with residuals larger than 3. Note that rQVIO(1) has different bit/meas. compared with what was reported in our previous work because image tracking uses improved parameters. For zQVIO2, the first measurement communication is set to 10 bits. We vary the number of bits for differential measurements from 3 to 5. To avoid out-of-range issues and minimize quantization error, instead of quantizing differential measurements to a fixed range, we use 16 bits to communicate the largest differential measurements at the current time, all the other differential measurements are divided by it and then scaled to 0 to 1 before quantization. We also compared with our previous work [1] with rQVIO(1) and zQVIO using 6 to 8 bits, reported in Table III. The results show that while zQVIO experiences an accuracy drop starting at 7-8 bits, zQVIO2 maintains good accuracy down to 4.6 bits, thanks to the improvements we introduced.

In general, rQVIO2 has improved accuracy with more bits available as expected because more information is preserved. Notably, rQVIO2(2), using only 3.5 bits on average, achieves accuracy comparable to VIO. However, for some of the sequences, the increased bits lead to an accuracy drop, especially for sequences in the machine hall. We conjecture the reason is that in those sequences, the image has a larger illumination change, which can include more outliers in feature tracking. Using fewer bits incurs more information loss, however, it might also mitigate the impact of outliers, as less wrong information is absorbed into the system. This interesting finding can be further explored to develop a method to robustify estimator with noisy measurements or under a high ratio of outliers.

VI. CONCLUSIONS AND FUTURE WORK

Our previous work introduced the first quantized VIO system, demonstrating impressive performance but with some limitations [1]. In this work, significantly improved QVIO that pushes the limits by reducing data communication while achieving better performance. First, we improved the data quantization method. For zQVIO, we developed the differential measurement quantization method to data transmission without extra errors. For rQVIO, we extended it to handle multi-bit measurement to allow accuracy improvement with minimal increase in communication. Second, we developed

a quantized estimator derived from MAP that supports both single and multi-bit measurements for different quantization methods. Thanks to its mathematical rigor, it handles quantized data for zQVIO without the need to inflate measurement noise. Additionally, it supports multi-bit data for rQVIO, improving accuracy over the previous single-bit approach. Building on these improvements, we developed zQVIO2 and rQVIO2, while optimizing the system architecture and update indicator to reduce unnecessary data transfers and lower communication requirements. In our study, we show that zQVIO2 can achieve almost no accuracy loss with only 4.6 bits while rQVIO2 can reduce accuracy loss by introducing more bits and preserving accurate estimation using only 3.5 bits on average. In the future, we plan to explore applying the proposed approach to cooperative visual-inertial localization.

APPENDIX I

Q-FUNCTIONS IN QUANTIZED MAP

We now show how to derive the Q function for both zQVIO and rQVIO in Section II-B. Start from zQVIO, given quantized measurement z_i^q , and the quantization resolution r_q as in Eq. (9) and (10), we can derive the conditional probability of quantized measurements as:

$$\begin{aligned}
 & P\left\{t_i^1 < \frac{z_i^q - h(\mathbf{x}) - n}{\sigma} < t_i^2 \mid \mathbf{x}\right\} \\
 &= P\left\{-\frac{r_q}{2\sigma} < \frac{z_i^q - h(\mathbf{x}) - n}{\sigma} < \frac{r_q}{2\sigma} \mid \mathbf{x}\right\} \\
 &= P\left\{\frac{1}{\sigma}(z_i^q - h(\mathbf{x}) - \frac{r_q}{2}) < \frac{n}{\sigma} < \frac{1}{\sigma}(z_i^q - h(\mathbf{x}) + \frac{r_q}{2}) \mid \mathbf{x}\right\} \\
 &= P\left\{\frac{n}{\sigma} > \frac{1}{\sigma}(z_i^q - h(\mathbf{x}) - \frac{r_q}{2}) \mid \mathbf{x}\right\} \\
 &\quad - P\left\{\frac{n}{\sigma} > \frac{1}{\sigma}(z_i^q - h(\mathbf{x}) + \frac{r_q}{2}) \mid \mathbf{x}\right\} = \mathbf{Q}(z\chi_i^1) - \mathbf{Q}(z\chi_i^2)
 \end{aligned} \tag{17}$$

Similarly, given quantized state $\bar{\mathbf{x}}$, and the residual range (t_1, t_2) indicated by the quantized measurements as in Eq. (11) and (12), we can derive for the residual quantization:

$$\begin{aligned}
 & P\left\{t_i^1 < \frac{z_i^q - h(\bar{\mathbf{x}}) - n}{\sigma} < t_i^2 \mid \mathbf{x}\right\} \\
 &= P\left\{t_1 < \frac{h(\bar{\mathbf{x}}) - h(\mathbf{x}) - n}{\sigma} < t_2 \mid \mathbf{x}\right\} \\
 &= P\left\{\frac{1}{\sigma}(h(\bar{\mathbf{x}}) - h(\mathbf{x})) - t_1 < \frac{n}{\sigma} < \frac{1}{\sigma}(h(\bar{\mathbf{x}}) - h(\mathbf{x})) + t_2 \mid \mathbf{x}\right\} \\
 &= P\left\{\frac{n}{\sigma} > \frac{1}{\sigma}(h(\bar{\mathbf{x}}) - h(\mathbf{x})) - t_1 \mid \mathbf{x}\right\} \\
 &\quad - P\left\{\frac{n}{\sigma} > \frac{1}{\sigma}(h(\bar{\mathbf{x}}) - h(\mathbf{x})) - t_2 \mid \mathbf{x}\right\} = \mathbf{Q}(r\chi_i^1) - \mathbf{Q}(r\chi_i^2)
 \end{aligned} \tag{18}$$

We thus finish the derivation of the Q-function in the MAP optimization problem for quantized estimators.

REFERENCES

- [1] Y. Peng, C. Chen, and G. Huang, "Quantized visual-inertial odometry," in *Proc. International Conference on Robotics and Automation*, Yokohama, Japan, May 2024.
- [2] G. Huang, "Visual-inertial navigation: A concise review," in *Proc. International Conference on Robotics and Automation*, Montreal, Canada, May 2019.
- [3] C. Chen, Y. Yang, P. Geneva, W. Lee, and G. Huang, "Visual-inertial-aided online mav system identification," in *Proc. of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, Kyoto, Japan., 2022.
- [4] S. Katragadda, W. Lee, Y. Peng, P. Geneva, C. Chen, C. Guo, M. Li, and G. Huang, "Nerf-vins: A real-time neural radiance field map-based visual-inertial navigation system," in *Proc. International Conference on Robotics and Automation*, Yokohama, Japan, May 2024.
- [5] C. Chen, Y. Yang, P. Geneva, and G. Huang, "Fej2: A consistent visual-inertial state estimator design," in *Proc. of the IEEE International Conference on Robotics and Automation*, Philadelphia, PA, 2022.
- [6] C. Chen, P. Geneva, Y. Peng, W. Lee, and G. Huang, "Monocular visual-inertial odometry with planar regularities," in *Proc. of the IEEE International Conference on Robotics and Automation*, London, UK., 2023.
- [7] Y. Peng, C. Chen, and G. Huang, "Ultrafast square-root filter-based VINS," in *Proc. International Conference on Robotics and Automation*, Yokohama, Japan, May 2024.
- [8] C. Chen, Y. Peng, and G. Huang, "Fast and consistent covariance recovery for sliding-window optimization-based vins," in *Proc. International Conference on Robotics and Automation*, Yokohama, Japan, May 2024.
- [9] —, "Visual-inertial state estimation with decoupled error and state representations," in *Proc. of International Workshop on the Algorithmic Foundations of Robotics*, Chicago, IL, 2024.
- [10] "Reality labs chief scientist outlines a new compute architecture for true ar glasses," <https://www.roadtovr.com/michael-abrash-iedm-2021-compute-architecture-for-ar-glasses/>.
- [11] J. Gomez, S. Patel, S. S. Sarwar, Z. Li, R. Capoccia, Z. Wang, R. Pinkham, A. Berkovich, T.-H. Tsai, B. De Salvo, and C. Liu, "Distributed on-sensor compute system for ar/vr devices: A semi-analytical simulation framework for power estimation," 2022. [Online]. Available: <https://arxiv.org/abs/2203.07474>
- [12] I. Hubara, M. Courbariaux, D. Soudry, R. El-Yaniv, and Y. Bengio, "Quantized neural networks: Training neural networks with low precision weights and activations," *The Journal of Machine Learning Research*, vol. 18, no. 1, pp. 6869–6898, 2017.
- [13] G. Xiao, J. Lin, M. Seznec, H. Wu, J. Demouth, and S. Han, "Smoothquant: Accurate and efficient post-training quantization for large language models," in *International Conference on Machine Learning*. PMLR, 2023, pp. 38087–38099.
- [14] Q. Picard, S. Chevobbe, M. Darouich, and J.-Y. Didier, "Image quantization towards data reduction: robustness analysis for slam methods on embedded platforms," in *2022 IEEE International Conference on Image Processing (ICIP)*. IEEE, 2022, pp. 4158–4162.
- [15] O. Christie, J. Rego, and S. Jayasuriya, "Analyzing sensor quantization of raw images for visual slam," in *2020 IEEE International Conference on Image Processing (ICIP)*. IEEE, 2020, pp. 246–250.
- [16] G. K. Wallace, "The jpeg still picture compression standard," *Communications of the ACM*, vol. 34, no. 4, pp. 30–44, 1991.
- [17] A. Suleiman, Z. Zhang, L. Carlone, S. Karaman, and V. Sze, "Navion: A 2-mw fully integrated real-time visual-inertial odometry accelerator for autonomous navigation of nano drones," *IEEE Journal of Solid-State Circuits*, vol. 54, no. 4, pp. 1106–1119, 2019.
- [18] Y. Sun, M. Liu, and M. Q.-H. Meng, "Improving rgb-d slam in dynamic environments: A motion removal approach," *Robotics and Autonomous Systems*, vol. 89, pp. 110–122, 2017.
- [19] M. Mera-Trujillo, B. Smith, and V. Fragoso, "Efficient scene compression for visual-based localization," in *2020 International Conference on 3D Vision (3DV)*. IEEE, 2020, pp. 1–10.
- [20] L. Zheng, K. Xu, J. Jiang, M. Wei, B. Zhou, and H. Cheng, "Real-time efficient environment compression and sharing for multi-robot cooperative systems," *IEEE Transactions on Intelligent Vehicles*, 2024.
- [21] M. Dymczyk, S. Lynen, T. Cieslewski, M. Bosse, R. Siegwart, and P. Furgale, "The gist of maps-summarizing experience for lifelong localization," in *2015 IEEE international conference on robotics and automation (ICRA)*. IEEE, 2015, pp. 2767–2773.
- [22] T. Sattler, M. Havlena, F. Radenovic, K. Schindler, and M. Pollefeys, "Hyperpoints and fine vocabularies for large-scale location recognition," in *Proceedings of the IEEE International Conference on Computer Vision*, 2015, pp. 2102–2110.
- [23] S. Lynen, T. Sattler, M. Bosse, J. A. Hesch, M. Pollefeys, and R. Siegwart, "Get out of my lab: Large-scale, real-time visual-inertial localization," in *Robotics: Science and Systems*, vol. 1, 2015, p. 1.
- [24] L. Baroffio, A. E. Redondi, M. Tagliasacchi, and S. Tubaro, "A survey on compact features for visual content analysis," *APSIPA Transactions on Signal and Information Processing*, vol. 5, p. e13, 2016.
- [25] D. Van Opendenbosch and E. Steinbach, "Collaborative visual slam using compressed feature exchange," *IEEE Robotics and Automation Letters*, vol. 4, no. 1, pp. 57–64, 2018.
- [26] A. Ribeiro, G. B. Giannakis, and S. I. Roumeliotis, "Soi-kf: Distributed kalman filtering with low-cost communications using the sign of innovations," *IEEE Transactions on Signal Processing*, vol. 54, no. 12, pp. 4782–4795, 2006.
- [27] E. D. Nerurkar, K. X. Zhou, and S. I. Roumeliotis, "A hybrid estimation framework for cooperative localization under communication constraints," in *2011 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2011, pp. 502–509.
- [28] E. D. Nerurkar and S. I. Roumeliotis, "A communication-bandwidth-aware hybrid estimation framework for multi-robot cooperative localization," in *2013 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2013, pp. 1418–1425.
- [29] —, "Hybrid maximum a posteriori estimation under communication constraints," in *2013 IEEE International Conference on Acoustics, Speech and Signal Processing*. IEEE, 2013, pp. 4898–4902.
- [30] C. Hertzberg, R. Wagner, U. Frese, and L. Schröder, "Integrating generic sensor fusion algorithms with sound state representations through encapsulation of manifolds," *Information Fusion*, vol. 14, no. 1, pp. 57–77, Jan. 2013.
- [31] E. J. Msechu, S. I. Roumeliotis, A. Ribeiro, and G. B. Giannakis, "Decentralized quantized kalman filtering with scalable communication cost," *IEEE Transactions on Signal Processing*, vol. 56, no. 8, pp. 3727–3741, 2008.
- [32] N. Trawny, S. I. Roumeliotis, and G. B. Giannakis, "Cooperative multi-robot localization under communication constraints," in *2009 IEEE International Conference on Robotics and Automation*. IEEE, 2009, pp. 4394–4400.
- [33] 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*, 2007.
- [34] P. Geneva, K. Eickenhoff, W. Lee, Y. Yang, and G. Huang, "Openvins: A research platform for visual-inertial estimation," in *Proc. of the IEEE International Conference on Robotics and Automation*, Paris, France, 2020. [Online]. Available: https://github.com/rpng/open_vins
- [35] 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*, vol. 35, no. 10, pp. 1157–1163, 2016.