

# Lossless SIMD Compression of LiDAR Range and Attribute Scan Sequences

Jeff Ford<sup>1</sup> and Jordan Ford<sup>2</sup>

**Abstract**—As LiDAR sensors have become ubiquitous, the need for an efficient LiDAR data compression algorithm has increased. Modern LiDARs produce gigabytes of scan data per hour (Fig. 1) and are often used in applications with limited compute, bandwidth, and storage resources.

We present a fast, lossless compression algorithm for LiDAR range and attribute scan sequences including multiple-return range, signal, reflectivity, and ambient infrared. Our algorithm—dubbed “Jiffy”—achieves substantial compression by exploiting spatiotemporal redundancy and sparsity. Speed is accomplished by maximizing use of single-instruction-multiple-data (SIMD) instructions. In autonomous driving, infrastructure monitoring, drone inspection, and handheld mapping benchmarks, the Jiffy algorithm consistently outperforms competing lossless codecs while operating at speeds in excess of 65M points/sec on a single core. In a typical autonomous vehicle use case, single-threaded Jiffy achieves 6x compression of centimeter-precision range scans at 500+ scans per second. To ensure reproducibility and enable adoption, the software is freely available as an open source library<sup>3</sup>.

## I. INTRODUCTION

LiDAR sensors are rapidly becoming commonplace in applications ranging from infrastructure monitoring to autonomous cars and drones. At the same time, LiDAR sensors are rapidly improving in range, precision, spatial resolution, and scan rate. Additional outputs, including multi-return range, intensity, reflectivity, and ambient light, result in a flood of LiDAR data that must be processed, transmitted, and stored, consuming scarce resources of bandwidth, memory, time, and power. Fast, lossless LiDAR compression is needed to address the deluge (Fig. 1).

General purpose algorithms have been used to compress heterogeneous sensor data losslessly, but without prior knowledge of the data, they are limited. Algorithms specifically designed for LiDAR scan data have been proposed to exploit the structure of LiDAR data. These algorithms have primarily focused on lossy compression of LiDAR range scans by reusing concepts from image, video, or point cloud compression. Thus far, they have proven too slow, too lossy, or too ineffective for wide adoption.

This work leverages the sparsity and spatiotemporal coherence of LiDAR scan sequences to develop Jiffy—a fast, lossless codec for high-resolution, multiple-attribute LiDAR scans (Fig. 2). Single-threaded performance is ensured by maximizing use of single-instruction-multiple-data (SIMD)

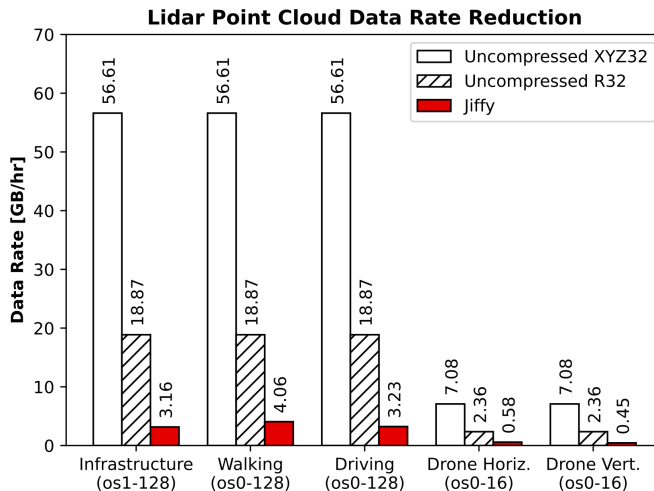


Fig. 1. The Jiffy algorithm uses techniques adapted from image encoding and database compression to achieve lossless compression of LiDAR range scan sequences by a factor of 12-18x vs. uncompressed 32-bit Cartesian point cloud format, or 4-6x vs. uncompressed 32-bit range image format in applications including autonomous driving, infrastructure monitoring, and handheld & drone mapping.

instructions. In multi-attribute benchmarks from real-world use cases, Jiffy consistently achieves best-in-class lossless compression with single-CPU throughput in excess of 65 million points per second.

## II. RELATED WORK

Some of the earliest general purpose, lossless compression algorithms, Bzip2[1] and DEFLATE (Zlib)[2] have been used to compress binary files for decades. DEFLATE forms the basis for many application-specific file formats (e.g. ZIP[3], DOCX[4], PNG[5], OpenEXR[6]).

These early compression algorithms were developed before memory became a bottleneck, before many-core computers became common, and before standard SIMD instruction sets were widely available. Recently, new algorithms have been designed to take advantage of the hardware in modern CPUs. LZ4 is a fast dictionary matching algorithm that runs at GB/sec throughput, but produces modest compression ratios[7]. ZStandard (Zstd) leverages recent developments in Asymmetric Numeral Systems (ANS) and SIMD implementation to achieve improved compression at much faster speeds[8]. The popular Robotics Operating System (ROS) initially adopted Bzip2 as its data logging compression codec, and later added LZ4 as a faster option[9]. Both

<sup>1</sup>Jeff Ford, ComplexIQ, Atlanta, GA 30517, USA  
jeff@complexiq.com

<sup>2</sup>Jordan Ford, Department of Electrical Engineering, Carnegie Mellon University, Pittsburgh, PA 15217, USA  
jsford@andrew.cmu.edu

<sup>3</sup>Software is available here: <http://github.com/jsford64/jiffy-compression>

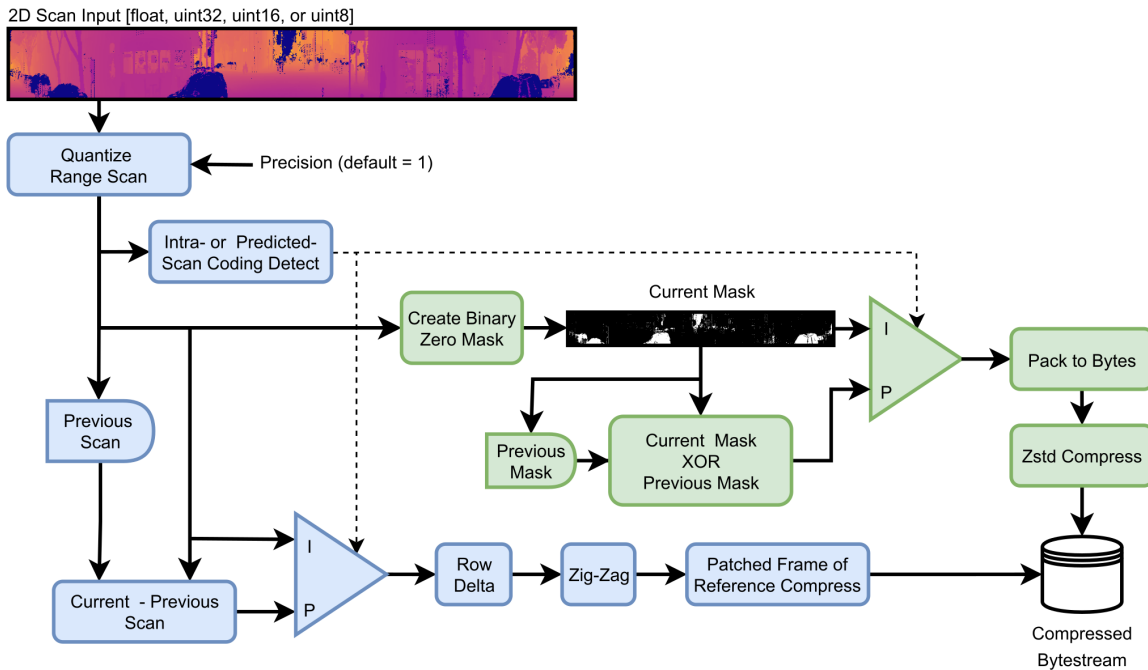


Fig. 2. The Jiffy algorithm losslessly compresses high bit-depth, multi-attribute LiDAR scan data. Incoming scans are quantized to the precision of the sensor. A bitmask locating zero-valued samples is extracted from the scan and compressed separately (green) from the remaining scan data (blue). A trial compression heuristic selects an encoding method: intra-scan (I-scan) coding is selected when the previous scan differs significantly from the current scan, otherwise, predicted-scan (P-scan) coding is selected.

LZ4 and Bzip2 have been replaced by Zstd in ROS2[10]. These general purpose compressors are convenient for compressing heterogeneous sensor data because they do not require a specific model of the data, but without a model of the data, their compression ratios are often limited.

Point cloud compression algorithms like Draco[11], LASZip[12], and PCD[14] have been employed to compress LiDAR-derived point clouds and attributes. Google's Draco triangle mesh codec treats point clouds as degenerate meshes with no edges or faces. Floating-point vertices are quantized as integers, then ANS coding is applied, similar to Zstd. The PCD file format groups coordinates by axis, then applies fast LZF compression. The Motion Picture Experts Group (MPEG) has developed two codecs, G-PCC and V-PCC, for compressing point clouds[19]. The Geometry-Point Cloud Compressor (G-PCC) is based on voxel- and tree-based decomposition of 3D space and is best-suited for static point clouds. The Video-Point Cloud Compressor (V-PCC) tries to reuse the HEVC video codec. V-PCC searches for planar regions of a point cloud and projects geometry and attribute information onto them. It then packs planar segments into video frames and applies HEVC. Deep point cloud compression using hardware-intensive neural networks is an active area of research. A more complete review of point cloud compression can be found in Pereira et al.[17].

The range image representation is a method of representing LiDAR point measurements in spherical coordinates. Conversion from Cartesian XYZ coordinates to the range image representation can mean 3x compression of a LiDAR

scan. Azimuth and altitude angles are implied by the x,y position of a sample, leaving an x,y array of polar magnitudes (ranges). Conversion to range images enables the use of image and video compression codecs. Tu et al. applied JPEG2000 to LiDAR range images[21]. Ahn et al. developed a lossless codec for stationary survey scanners[20]. Their method recursively subdivides the image into patches and applies a variety of predictors to each patch to achieve reasonable compression at the expense of speed. Nenci et al. used many instances of H.264 video codecs to compress LiDAR sequences[18]. They achieved 10x lossless compression of Microsoft Kinect range images at 25 Hz using 8 parallel video streams, but their method does not scale to long range sensors. Their method requires 391 concurrent video streams to encode a 100 meter range LiDAR scan with one millimeter resolution.

Researchers interested in database compression have developed integer compression algorithms capable of losslessly compressing billions of unsigned 32-bit integers per second at compression ratios competitive with general purpose compressors like Zstd. Goldstein et al.[22] developed Frame Of Reference (FOR) compression, which groups values into frames (e.g. 128 integers). The minimum value in the frame is identified, then the values in that frame are re-coded as an offset from that minimum. Zukowski et al.[23] proposed Patched-FOR (PFOR), which stores values exceeding a thresholded number of bits in a separate list. Lemire et al. [24] implemented SIMD-FastPFOR (SIMDPFOR)— PFOR using SIMD acceleration.

### III. METHODOLOGY

#### A. Standardizing LiDAR Input

LiDAR scanners produce geometry and attribute information in formats that differ by manufacturer. Some sensors output range data as an ordered stream of floating point Cartesian coordinates. Others represent ranges as an integer number of millimeters and natively output range images. Signal, reflectivity and other attributes may be returned as floats, as n-byte unsigned integers, or not at all.

Jiffy expects all scan types in a canonical 2D format. Cartesian points must be converted to spherical coordinates, sorted by altitude and azimuth, then quantized to 1, 2 or 4 byte unsigned integers. Invalid or out-of-range measurements are set to a sentinel value of zero. The user-provided quantization precision is stored in the compressed bitstream and used to restore the original representation of a quantized scan. Attribute values are similarly arranged and quantized.

#### B. Quantization

Quantization precision is an important parameter of the Jiffy encoder. Range measurements are quantized losslessly by default, but lossy quantization can be used to improve compression. Pereira[17] defines quantization as lossless when the quantization precision is less than or equal to the resolution of the sensor. By default, the Jiffy encoder quantizes range measurements using the sensor resolution—one millimeter for all sensors evaluated here (Table I).

In some applications the sensor resolution is orders of magnitude more precise than is necessary (e.g. a streaming visualizer), and aggressive quantization can be used to substantially improve compression. It is straightforward to set the quantization precision, and the maximum quantization error is always less than half of the precision.

#### C. Intra-Scan (I-Scan) Compression

Delta compression exploits spatial redundancy by subtracting neighboring samples within a scanline. Jiffy subtracts each sample in a scanline from its left-side neighbor and compresses the results.

Delta encoding is effective on range data because the difference between adjacent samples is typically much smaller than the samples themselves. However, because Jiffy encodes 'out of range' (too far or too close) with a zero-valued sentinel, transitions into- and out-of-range can produce very large deltas, reducing compression. To eliminate the large deltas caused by these transitions, Jiffy creates a binary mask of one bit per sample. The mask records a zero bit for in-range samples and a one for out-of-range samples. The bitmask is packed into bytes and compressed using Zstd. After creating the out-of-range bitmask, all out-of-range samples are removed from the 2D scan, creating a flattened 1D vector. The 1D vector can now be delta-encoded without introducing large differences caused by transitions into and out of range. Masking of zero values is not required for attribute scan types since they do not use zero to indicate valid or invalid measurements. However, Jiffy applies zero

masking to all scan types, producing a positive effect on compression ratio.

The SIMD-PFOR algorithm is designed to compress unsigned integers, but delta encoding produces many negative values. Signed integers can be cast to unsigned and sent directly to SIMD-PFOR for compression, but negative integers cast to extremely large unsigned values because two's-complement sign extension sets all of their highest bits. Since SIMD-PFOR is designed to pack small unsigned integers with many leading zeros in the upper bits, this leads to poor compression. Jiffy uses ZigZag encoding to represent negative differences without the problems caused by sign extension (Equation 1).

$$\text{ZigZag}(x) = 2|x| + \begin{cases} 1 & \text{if } x < 0 \\ 0 & \text{otherwise} \end{cases} \quad (1)$$

Jiffy performs standalone (Intra-scan or I-scan) compression of individual LiDAR scans by masking and removing out-of-range samples, delta encoding the remaining values, ZigZag encoding the deltas, and SIMD-PFOR compressing the result. I-scan decoding is performed by a straightforward reversal of the encoding procedure.

#### D. Predicted-Scan (P-Scan) Compression

When LiDAR sensors are stationary or moving slowly, successive scans exhibit a high degree of similarity that can be exploited to improve compression. Stationary infrastructure monitoring LiDARs typically overlook a largely static scene. Autonomous vehicles are frequently stationary, waiting for passengers or idling in traffic. In these situations, it is redundant to encode near-identical scans of a static scene. Jiffy provides an inter-scan compression mode called Predicted-scan (P-scan). The 'predicted' nomenclature is borrowed from MPEG video compression naming conventions. Like Intra-scan encoding, a compressed bitmask is created to represent zero-valued samples in the input scan. In predicted-scan (P-scan) mode, the bitmask for the current frame is XORed with the bit mask from the previous frame to remove redundancy from successive bitmasks. The XORed mask result is compressed using Zstd and appended to the output bitstream. Jiffy's P-scan compression mode uses temporal differencing to avoid encoding redundant scans. After masking the current frame and the previous frame with the current bitmask, the scan differences are encoded using the I-scan compression method.

#### E. Automatic I/P Mode Selection

For some applications, I-scan only or P-scan only encoding might make sense. For example, a fast-moving drone could use I-scan only encoding for best results, whereas a pole-mounted monitoring sensor would benefit from using P-scan only. However, applications such as a car in stop-start traffic, a racing drone sitting on its launch pad, or a tree branch waving in front of a static monitor sensor should also be considered. These cases, and in fact, most use-cases benefit from the ability to adaptively select a compression mode on a per-scan basis.

The Jiffy encoder could employ brute force to select a compression mode. Because I- and P-scan compression modes are extremely fast, it might be reasonable to encode each scan using both methods and then select the better result. Although brute force trial compression ensures the most effective mode is used, it also reduces throughput by 50%. As explained in subsection IV-E, a trial compression of a few scanlines in each scan (excluding mask compression) proves to be successful in selecting the most effective compression mode for the full scan.

## IV. EXPERIMENTS

### A. Software and System

All code, including benchmarks, compatible datasets, implementations of competing codecs, and unit tests, are freely available<sup>3</sup>. All codecs and benchmarks are implemented in python3. Performance critical code relies on python3 bindings to fast C/C++ libraries such as numpy, Zstd, SIMD-PFOR, zlib, lz4, etc. Benchmarks were performed on a 12th Gen Intel i7-1250U CPU and a NVidia Jetson Xavier NX ARM CPU (Table III), both running at 1.9GHz. The benchmark process is pinned to CPU 0 (a performance core) and configured to use the performance CPU governor. Sensors used are summarized in Table I.

### B. Datasets

Benchmarks for several LiDAR use cases are employed: autonomous driving, infrastructure monitoring, handheld mapping, and drone flying. To simplify evaluation, all datasets were constructed using Ouster OS0, OS1, or OS2 LiDAR sensors. The autonomous driving and infrastructure monitoring datasets were supplied by the sensor manufacturer[25]. The handheld mapping dataset is extracted from the park sequence of the multi-camera Newer College dataset[26], and the drone datasets are derived from the nya\_02 sequence of the NTU VIRAL dataset[27]. The NTU drone dataset includes simultaneous scans from one vertically-mounted and one horizontally-mounted OS0-16 LiDAR sensor. Because their compression characteristics differ significantly, the horizontal- and vertical-mounted drone LiDAR scans are treated as separate use cases.

The autonomous driving datasets were collected using the short-range, medium-range, and long-range sensors, detailed in Table I. Results for the codec comparisons in subsection IV-C are averaged across all three sequences. All datasets contain range scans with single-precision floating point samples. The autonomous driving datasets additionally contain signal, reflectivity, near-infrared, and second-return scan types. The second-return range2 scan type contains 32-bit ranges. The signal scan type and its second-return equivalent, signal2, record the 16-bit unsigned integer intensity of the LiDAR return pulse. The reflectivity and reflectivity2 scan types record an 8-bit unsigned integer estimate of surface reflectivity for every sample in a scan. The near-infrared scan type is a 16-bit unsigned integer measurement of ambient light. Dataset features are summarized in Table II.

TABLE I  
A SUMMARY OF BENCHMARKED LiDAR SENSORS

Sensor	Range	Resolution	Accuracy	Precision
Ouster OS0	0.3 - 50 m	1 mm	± 3 cm	± 1-5 cm
Ouster OS1	0.3 - 120 m	1 mm	± 3 cm	± 0.7-5 cm
Ouster OS2	1.0 - 240 m	1 mm	± 3 cm	± 2.5-8 cm

TABLE II  
A SUMMARY OF BENCHMARKED DATASETS

Dataset Name	Sensor	Scan Shape	# Scans	Attributes
Driving-0	OS0	128 × 1024	1432	✓
Driving-1	OS1	128 × 1024	719	✓
Driving-2	OS2	128 × 1024	971	✓
Infrastructure	OS1	128 × 1024	690	✗
Handheld	OS0	128 × 1024	1947	✗
Drone (Horizontal)	OS0	16 × 1024	4187	✗
Drone (Vertical)	OS0	16 × 1024	4186	✗

### C. Compression Results vs. Existing Lossless Algorithms

A comparison of encoding speed and compression ratio for the Jiffy codec and competing lossless methods was performed using the autonomous vehicle data sets[25]. The mean speed and mean compression ratio over all range scans are presented in Fig. 3. Lz4 is consistently fastest, but its compression ratio ranks among the worst. The Jiffy codec is the clear winner in both speed and compression ratio for the first-return 'range' scan types and the second-return 'range2' and 'signal2' scan types. These scan types exhibit the most sparsity and the highest dynamic range, so they benefit most from the zero bitmasking and SIMD-PFOR algorithms. Horizontal delta encoding provides compression gains due to the low spatial frequency content of this scan type, and adaptive coding of I- and P-scans also contributes to higher compression ratios when the vehicle comes to a stop.

Reflectivity (first- and second- returns) were the most challenging for Jiffy, as these scan types had the smallest dynamic range (8 bits) coupled with the highest spatial frequency content. The high frequency content counters the benefit of delta encoding, and coupling that deficit with the 8-bit dynamic range provides little opportunity for the SIMD-PFOR codec to effectively pack the data. The Zlib codec was a close competitor for these scan types. The Jiffy compression ratio for the nearIR 16-bit scan type was very competitive with the Bzip2 and Lzma codecs at a 7x to 21x higher compression speed.

### D. Compression Results by Use Case

Compression performance was tested for five unique usage scenarios: autonomous driving, with short- medium- and long-range sensors, infrastructure monitoring with a stationary long range sensor, a short range walking scan using a handheld sensor without motion stabilization, and a drone equipped with two short range sensors, one scanning horizontally and one vertically (Fig. 4).

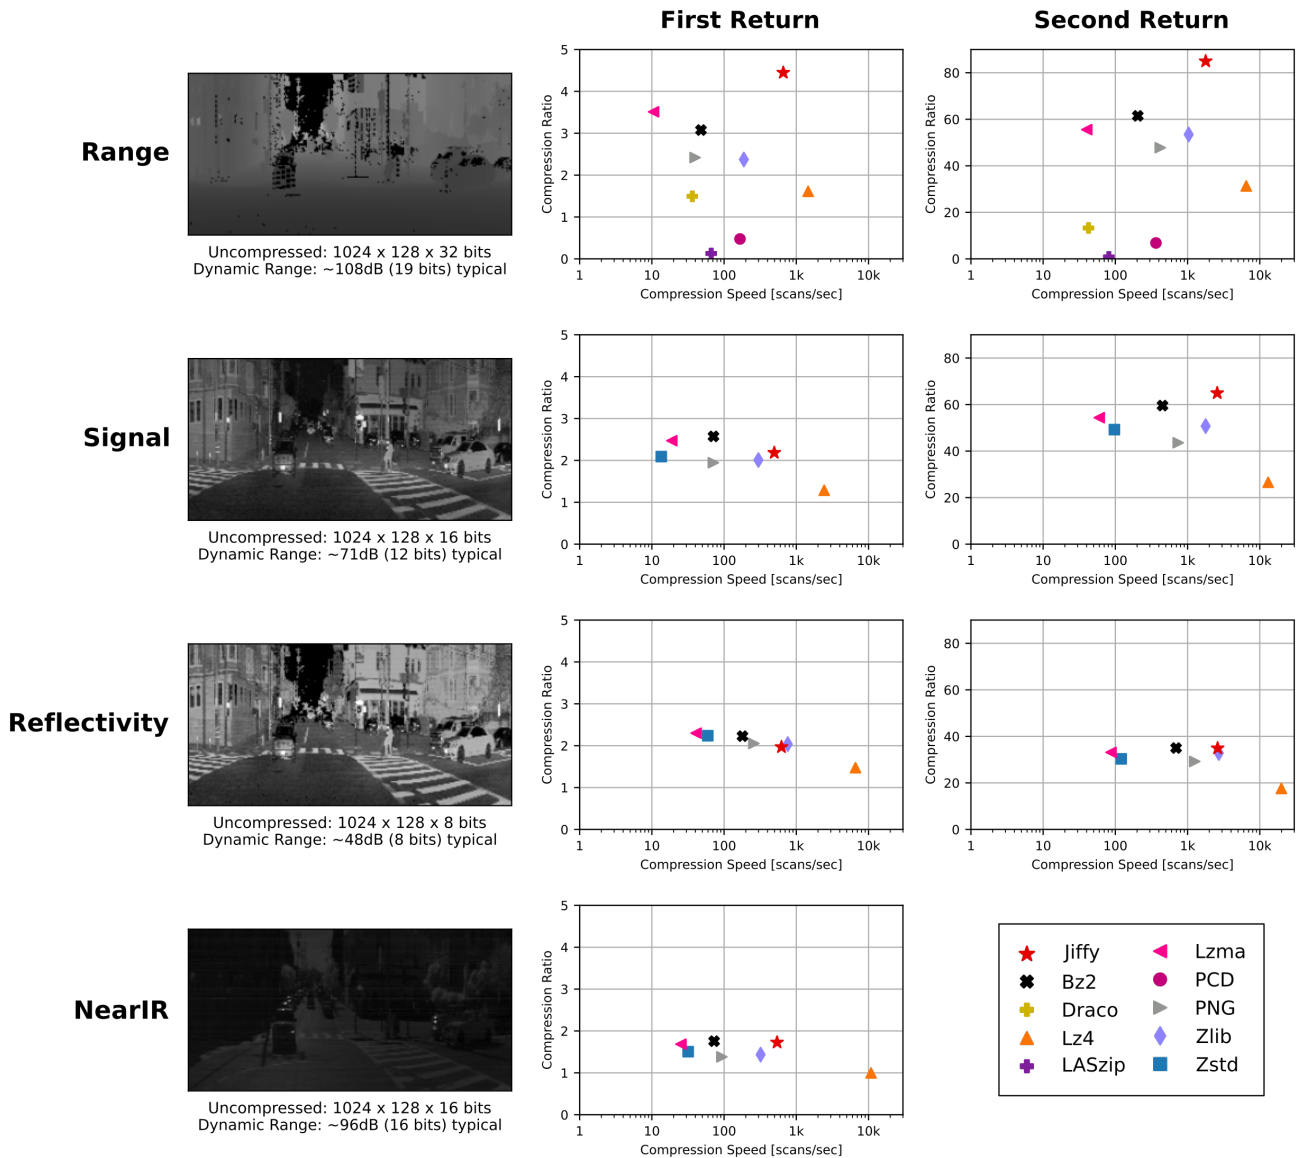


Fig. 3. The Jiffy codec compresses LiDAR range scans by a factor of 4.4 at a rate of 660 scans per second on a single core of a 12th Gen Intel i7-1250U processor. The same codec is pareto-optimal for compression of 16-bit per pixel intensity (Signal) and ambient infrared (NearIR) scan data and is near-optimal for compression of 8-bit per pixel reflectivity data. For sparse second-return data, the Jiffy codec operates at 1000 scans per second and compresses at 35-85x compression ratios. Results presented here are averaged across three autonomous vehicle datasets collected using Ouster OS0-128, OS1-128, and OS2-128 LiDAR sensors.

For the autonomous driving scenario at the default 1 mm precision, the short range sensor had the best compression ratio (4x), and the long range sensor had the lowest (2.6x). This result clearly demonstrates the efficient bit packing of the SIMD-PFOR algorithm, because the short range sensor data uses fewer of the original 32 bits per sample than the medium- and long-range sensors.

The drone scenario clearly demonstrates the effect of the zero bitmask compression. The horizontally- and vertically-mounted sensors have the same dynamic range, but the vertically mounted sensor data has a higher degree of sparsity relative to the horizontally mounted sensor, so it achieves a higher compression ratio (3.3x horizontal vs. 4.0 vertical).

The infrastructure scenario clearly demonstrates the coding gain of the temporal (P-Scan) compression. This scenario is using the long-range sensor, yet it exhibits a compression ratio that rivals the best-compressed scenarios using the short range sensor, drone(vertical), walking, and driving.

#### E. Validation of Spatiotemporal Encoder Selection Heuristic

Testing the efficacy and error rate of the *IP* scan selection heuristic was performed by comparing the *predicted* optimal compression method (I or P), to the *actual* optimum compression method for each of 22,877 range scans in all test datasets. The *actual* optimal method was determined by comparing the relative results of compressing the scans using an I-scan only compressor and a P-scan only compressor.

Using just 4 test lines in the heuristic test, the algorithm correctly predicted the optimal compression method for 96% of the 22,877 scans tested, with 0.5% sub-optimal I-scans and 3.5% sub-optimal P-scans.

### F. Compression / Decompression Speed

Jiffy encode/decode speed was tested using three autonomous driving datasets. All three datasets use high data rate sensors that produce  $128 \times 1024$  scans at 10 Hz. Benchmark results for each scan type are presented in Table III. On a single CPU core, the Jiffy codec sustains a mean encoding(decoding) rate of more than 1300(2100) scans/sec over all scan types. On a single thread, Jiffy can encode(decode) all seven scan types at a rate of 120 Hz, twelve times faster than the sensor's 10 Hz output rate.

TABLE III  
JIFFY COMPRESSION/DECOMPRESSION RATES BY SCAN TYPE

Scan Type	Intel Encode / Decode [scans/sec]	ARM Encode / Decode [scans/sec]
Range	661 / 881	595 / 782
Range2	1782 / 3335	1602 / 2749
Signal	496 / 579	572 / 743
Signal2	2574 / 4541	1777 / 2953
Reflectivity	624 / 799	548 / 631
Reflectivity2	2608 / 4454	1767 / 2972
Near Infrared	542 / 579	558 / 613

### G. Effect of Quantization on Compression Ratio

Jiffy quantizes LiDAR range data with a precision of one millimeter by default. For less precise sensors, or in circumstances where quantization errors are acceptable, more aggressive quantization can be used to improve compression. Increased quantization leads to a predictable improvement in compression performance with bounded absolute errors (Fig. 4). Doubling the quantization precision removes one bit of information from each range measurement. This enables Jiffy to compress each range measurement using at least one less bit and produces a corresponding improvement in

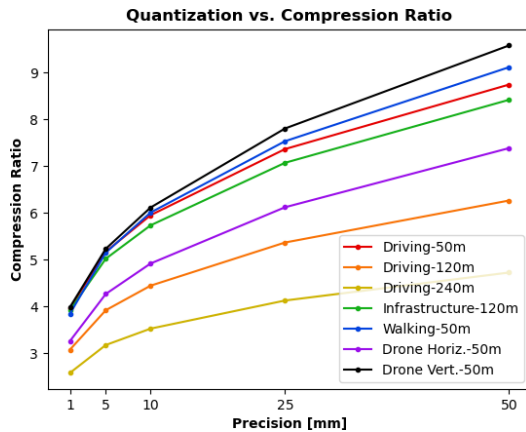


Fig. 4. Jiffy quantizes LiDAR range measurements prior to lossless compression. For applications where some error is acceptable, aggressive quantization can be traded for improved compression. Doubling the quantization precision  $P_q$  decreases the compressed output by approx. one bit per measurement. Max. roundtrip error is upper bounded by  $0.5P_q$ .

compression ratio. Because Jiffy compression of quantized scan data is lossless, the decoded error will not exceed half of the quantization precision setting.

### H. Ablation Study of the Jiffy Algorithm

An ablation study of the Jiffy encoder using an infrastructure monitoring LiDAR dataset[25] demonstrates the compression achieved by each component of the Jiffy algorithm (Table IV). The entire dataset is encoded using bare SIMD-PFOR compression, resulting in a 3.2x compression ratio. Adding horizontal delta-encoding along the scanlines and compressing the residuals using SIMD-PFOR predictably lowers the ratio to 2.2x, due to the effect of negative integers on SIMD-PFOR. The ZigZag transform resolves this problem, boosting the ratio to 3.8x. The out-of-range mask encodes zero values in the range scan using much less than one bit per sample. The bitmask also allows the removal of zero values from the delta-encoded measurement stream. This effectively removes large residuals at the boundaries between out-of-range and in-range values. With masking, delta encoding, the ZigZag transform, and SIMD-PFOR compression, the dataset is compressed by a factor of 5.4x. Adding the final algorithmic improvement, P-scan prediction, completes the Jiffy encoding algorithm, achieving a final dataset compression ratio of 6.2x.

TABLE IV  
PERFORMANCE OF ABLATED JIFFY CODECS APPLIED TO THE INFRASTRUCTURE MONITORING LiDAR DATASET[25]

Ablated Codec Description	Compression Ratio
SIMD-PFOR	3.2
Delta→SIMD-PFOR	2.2
Delta→ZigZag→SIMD-PFOR	3.8
Mask→Delta→ZigZag→SIMD-PFOR	5.4
Mask→Predict→Delta→ZigZag→SIMD-PFOR (Jiffy)	6.2

## V. CONCLUSIONS

The Jiffy codec successfully addresses the need for a fast, lossless LiDAR compression algorithm. In multiple-attribute benchmarks from a variety of use cases, Jiffy achieved best-in-class lossless compression at a rate of hundreds of scans per second on a single CPU core. This is more than 10x faster than the uncompressed sensor output rate for all of the sensors tested. The ability to control the quantization of range data is an effective and intuitive way to influence the compression ratio with a predictable error ceiling.

The techniques introduced here could be easily adapted to compress high dynamic range depth images and videos produced by stereo cameras or solid-state LiDAR sensors. Further research could investigate compression of per-pixel velocity measurements produced by continuous-wave LiDAR sensors or of derived perception data such as per-pixel segmentation labels. Lossy compression is also of considerable interest, but research into lossy algorithms should also quantify the downstream impact of compression artifacts on perception and localization algorithms.

## REFERENCES

- [1] Seward, J. Official Home Page for bzip2., <https://sourceware.org/bzip2/>
- [2] Gailly, J. & Adler, M. zlib 1.2.11 Manual. *Zlib 1.2.11 Manual.*, <https://www.zlib.net/manual.html>
- [3] Buchholz, F. The structure of a PKZip file. *The Structure Of A PKZIP File.*, <https://users.cs.jmu.edu/buchhofp/forensics/formats/pkzip-printable.html>
- [4] Overview of the XML file formats in Office 2010. *Microsoft Corporation.* [https://docs.microsoft.com/en-us/previous-versions/office/office-2010/cc179190\(v=office.14\)](https://docs.microsoft.com/en-us/previous-versions/office/office-2010/cc179190(v=office.14))
- [5] Portable Network Graphics (PNG) Specification (Second Edition). W3C. <https://www.w3.org/TR/2003/REC-PNG-20031110/#10Compression>
- [6] OpenEXR File Layout - OpenEXR Documentation. *Industrial Light & Magic (ILM).* <https://openexr.readthedocs.io/en/latest/OpenEXRFileLayout.html>
- [7] Collet, Y. LZ4 - Extremely Fast Compression. *LZ4.*, <https://lz4.github.io/lz4/>
- [8] ZStandard - Real-Time Data Compression Algorithm. *Meta Platforms, Inc.*, <http://facebook.github.io/zstd/>
- [9] ROS Wiki - The rosbag command-line tool. *ros.org.*, <http://wiki.ros.org/rosbag/CommandLine#compress>
- [10] Rosbag2. *ROS Index.*, <https://index.ros.org/t/rosbag2/github-ros2-rosbag2/>
- [11] Draco 3D Graphics Compression. *Google LLC.*, <https://google.github.io/draco/>
- [12] Isenburg, M. LASzip: lossless compression of LiDAR data. *Photogrammetric Engineering And Remote Sensing.* **79**, 209-217 (2013)
- [13] The PCD (Point Cloud Data) File Format. *pointclouds.org.*, [https://pointclouds.org/documentation/tutorials/pcd\\_file\\_format.html](https://pointclouds.org/documentation/tutorials/pcd_file_format.html)
- [14] Wiesmann, L., Milioto, A., Chen, X., Stachniss, C. & Behley, J. Deep Compression for Dense Point Cloud Maps. *IEEE Robotics And Automation Letters.* **6**, 2060-2067 (2021)
- [15] Tu, C., Takeuchi, E., Carballo, A. & Takeda, K. Point Cloud Compression for 3D LiDAR Sensor using Recurrent Neural Network with Residual Blocks. *2019 International Conference On Robotics And Automation (ICRA).* pp. 3274-3280 (2019)
- [16] Yu, Y., Zhang, W., Yang, F. & Li, G. Rate-Distortion Optimized Geometry Compression for Spinning LiDAR Point Cloud. *IEEE Transactions On Multimedia.* pp. 1-1 (2022)
- [17] Pereira, F., Dricot, A., Ascenso, J. & Brites, C. Point cloud coding: A privileged view driven by a classification taxonomy. *Signal Processing: Image Communication.* **85** pp. 115862 (2020), <https://www.sciencedirect.com/science/article/pii/S0923596520300771>
- [18] Nenci, F., Spinello, L. & Stachniss, C. Effective compression of range data streams for remote robot operations using H.264. *IEEE International Conference On Intelligent Robots And Systems.* pp. 3794-3799 (2014,10)
- [19] MPEG Point Cloud Compression. *MPEG.org.*, <https://mpeg-pcc.org>
- [20] Ahn, J., Lee, K., Sim, J. & Kim, C. Large-Scale 3D Point Cloud Compression Using Adaptive Radial Distance Prediction in Hybrid Coordinate Domains. *IEEE Journal Of Selected Topics In Signal Processing.* **9**, 422-434 (2015)
- [21] Tu, C., Takeuchi, E., Miyajima, C. & Takeda, K. Compressing continuous point cloud data using image compression methods. *2016 IEEE 19th International Conference On Intelligent Transportation Systems (ITSC).* pp. 1712-1719 (2016)
- [22] Goldstein J, Ramakrishnan R, Shaft U. Compressing relations and indexes. *Proceedings of the Fourteenth International Conference on Data Engineering, ICDE 98*, IEEE Computer Society: Washington, DC, USA, pp. 370-379 (1998)
- [23] Zukowski M, Heman S, Nes N, Boncz P. Super-scalar RAM-CPU cache compression. *Proceedings of the 22nd International Conference on Data Engineering, ICDE 06*, IEEE Computer Society: Washington, DC, USA; pp. 59-71 (2006)
- [24] Lemire, D. & Boytsov, L. Decoding billions of integers per second through vectorization. *Software: Practice And Experience.* **45**, 1-29 (2013,5), <https://onlinelibrary.wiley.com/doi/10.1002/spe.2203>
- [25] Ouster LiDAR Sensor Sample Data. *Ouster, Inc.*, <https://ouster.com/resources/lidar-sample-data/os-sample-data/>
- [26] Zhang, L., Camurri, M. & Fallon, M. Multi-Camera LiDAR Inertial Extension to the Newer College Dataset. (2021)
- [27] Nguyen, T., Yuan, S., Cao, M., Lyu, Y., Nguyen, T. & Xie, L. NTU VIRAL: A Visual-Inertial-Ranging-LiDAR dataset, from an aerial vehicle viewpoint. *International Journal Of Robotics Research.* (2021)