Reducing the Complexity of Fingerprinting-Based Positioning using Locality-Sensitive Hashing

Conference Paper**Author(s):**

Tang, Larry; Ghods, Ramina; [Studer, Christoph](#)

Publication date:

2019-11

Permanent link:

<https://doi.org/10.3929/ethz-b-000461317>

Rights / license:

[In Copyright - Non-Commercial Use Permitted](#)

Originally published in:

<https://doi.org/10.1109/IEEECONF44664.2019.9048657>

Reducing the Complexity of Fingerprinting-Based Positioning using Locality-Sensitive Hashing

Larry Tang, Ramina Ghods, and Christoph Studer

*School of Electrical and Computer Engineering, Cornell University, Ithaca, NY
email: {lwt29, rg548, studer}@cornell.edu; website: vip.ece.cornell.edu*

Abstract—Localization of wireless transmitters based on channel state information (CSI) fingerprinting finds widespread use in indoor as well as outdoor scenarios. Fingerprinting localization first builds a database containing CSI with measured location information. One then searches for the most similar CSI in this database to approximate the position of wireless transmitters. In this paper, we investigate the efficacy of locality-sensitive hashing (LSH) to reduce the complexity of the nearest neighbor search (NNS) required by conventional fingerprinting localization systems. More specifically, we propose a low-complexity and memory efficient LSH function based on the sum-to-one (STOne) transform and use approximate hash matches. We evaluate the accuracy and complexity (in terms of the number of searches and storage requirements) of our approach for line-of-sight (LoS) and non-LoS channels, and we show that LSH enables low-complexity fingerprinting localization with comparable accuracy to methods relying on exact NNS or deep neural networks.

I. INTRODUCTION

Localization of wireless transmitters in indoor and outdoor scenarios has received growing interest over the decade. Conventional methods for outdoor localization are mainly based on triangulation or trilateration methods, which map time-of-flight (ToF), angle-of-arrival, or received signal strength (RSS) features to a specific location using geometric models [1]–[4]. All of these approaches rely on line-of-sight (LoS) connectivity to multiple basestations, access points, or satellites. A popular instance of such methods are global navigation-satellite systems (GNSS) that combine ToF measurements with geometric and physical models. In situations that lack LoS connectivity (which is the case for indoor scenarios) or systems where GNSS is unavailable (which is the case for ultra-low-power sensors), alternative positioning methods are required.

A. Localization using CSI-Fingerprinting

Location fingerprinting is a prominent method to enable wireless positioning in challenging propagation environments, such as non-LoS scenarios or channels with complex multipath propagation; see, e.g., [2], [5]–[8] and the references therein. The principle of location fingerprinting is as follows. In a first phase, one tabulates a large number of measured

channel-state information (CSI), which includes RSS [6], [7] or other measured channel features in the time or frequency domain [9]–[11] or at multiple receive antennas [12], with associated location information in a given area—the resulting CSI-fingerprint/location tuples are stored in a database. In a second phase, a new CSI-fingerprint is extracted from a wireless transmitter and the most similar CSI fingerprints in the database are retrieved. The stored locations associated with the nearest fingerprints are then used to generate an estimate of the transmitter’s location. While a nearest neighbor search (NNS) in the CSI-fingerprint database can provide a simple (often accurate) estimate of the transmitter’s location¹, the complexity of NNS in large fingerprinting databases can quickly become a bottleneck. More recently, localization approaches based on deep neural networks have been proposed in [15]–[19]. Such methods avoid a NNS and directly map measured CSI-fingerprints to location—learning these networks, however, typically requires a large number of CSI measurements at fine resolution in space (often of the order of a few wavelengths). Furthermore, the large number of network parameters can easily be of the same order as the size of the fingerprint database.

B. Contributions

In this paper, we reduce the complexity of the NNS step in traditional fingerprinting localization via locality-sensitive hashing (LSH) [20]–[23]. We design a computationally efficient LSH function that builds upon a randomized sum-to-one (STOne) transform [24], and we use approximate hash matches that reduce the complexity of the NNS and the size of the hash tables. We evaluate the accuracy and complexity (in terms of the number of searches and storage requirements) of our method for LoS and non-LoS channels in a massive multiuser (MU) multiple-input multiple-output (MIMO) wireless system. Finally, we compare the accuracy, complexity, and storage requirements of our approach to that of recent neural network-based localization methods proposed in [15]–[18].

II. FINGERPRINTING LOCALIZATION VIA LSH

We now introduce the principles of fingerprinting localization and summarize the basics of LSH. We then detail our approach for low-complexity LSH with approximate hash lookups.

¹More sophisticated methods, such as neural networks [13] or Bayesian methods [14] can be used to improve the nearest neighbor search estimate.

The work of LT was supported in part by the Cornell University Engineering Learning Initiatives (ELI). The of RG and CS was supported in part by the semiconductor research corporation (SRC) nCORE task 2758.004, by Xilinx Inc., and by the US National Science Foundation under grants ECCS-1408006, CCF-1535897, CCF-1652065, CNS-1717559, and ECCS-1824379. The authors thank E. Lei and O. Castañeda for their help with CSI generation.

A. Basics of Fingerprinting Localization

Fingerprinting localization proceeds in two phases. In the first phase, CSI-fingerprints $\{\mathbf{f}_n\}_{n=1}^N = \mathcal{F}$ and associated positions $\{\mathbf{x}_n\}_{n=1}^N = \mathcal{X}$ are measured in a given area and stored in a database. Here, the vector $\mathbf{f}_n \in \mathbb{R}^D$ corresponds to a D -dimensional CSI-fingerprint at position $\mathbf{x}_n \in \mathbb{R}^d$, where D is typically high-dimensional and d is either two or three dimensions. Depending on the application, CSI-fingerprints can represent RSSs acquired at multiple receivers, power-delay profiles, angle-of-arrival, and many others; see [2] for a survey.

In the second phase, an estimate for the location $\mathbf{x}_{n'}$ of a new transmitting device with index n' is generated. To this end, a CSI-fingerprint $\mathbf{f}_{n'}$ is extracted. Then, the indices associated with the K most similar fingerprints in $\{\mathbf{f}_n\}_{n=1}^N$ are identified

$$\mathcal{N}_K = \{n : \|\mathbf{f}_n - \mathbf{f}_{n'}\| < r, \mathbf{f}_n \in \mathcal{F}, n = 1, \dots, N\},$$

where $r > 0$ ensures that $|\mathcal{N}_K| = K$. One then approximates the location of transmitter n' from the set of similar locations $\mathcal{X}_K = \{\mathbf{x}_n\}_{n \in \mathcal{N}_K}$. If $K = 1$, then one can simply pick the location associated with the nearest CSI fingerprint; more sophisticated approaches are discussed in [13], [14].

B. Locality-Sensitive Hashing (LSH)

Finding the K nearest neighbors in a large dataset containing high-dimensional vectors suffers from the curse of dimensionality; this implies that carrying out an exhaustive search inevitably results in high complexity [21]. Locality-sensitive hashing (LSH) is a powerful method to perform an approximate nearest neighbor search in large datasets [20]–[23]. The principle behind LSH is to construct locality-sensitive hash functions for which similar datapoints have matching hash values and dissimilar datapoints have mismatched hash values. Mathematically, LSH relies on functions $h : \mathbb{R}^D \rightarrow \mathcal{S}$ for which the hash collision probability $\Pr[h(\mathbf{p}) = h(\mathbf{q})] = P_1$ of any two datapoints $\mathbf{p}, \mathbf{q} \in \mathbb{R}^D$ is large if $\|\mathbf{p} - \mathbf{q}\| \leq R$ and $\Pr[h(\mathbf{p}) = h(\mathbf{q})] = P_2$ is small if $\|\mathbf{p} - \mathbf{q}\| \geq cR$ with R being an application-dependent radius and $c > 1$ a constant; for LSH, one is particularly interested in the case $P_1 \gg P_2$. The finite-cardinality set $\mathcal{S}$ contains so-called buckets, which are nothing but unique hash values.

Approximate NNS via LSH is carried out in two phases. In the first phase, one computes hash values for all points in the dataset, i.e., $\{h(\mathbf{f}_n)\}_{n=1}^N$. In the second phase, for a given query point $\mathbf{f}_{n'}$, one computes $h(\mathbf{f}_{n'})$ and compares the resulting hash value to those in the dataset. One can then compare the true, high-dimensional CSI feature distance $\|\mathbf{f}_m - \mathbf{f}_{n'}\|$ associated to *only* those indices $m \in \{n : h(\mathbf{f}_n) = h(\mathbf{f}_{n'}), n = 1, \dots, N\}$ for which there was a collision. By using a set of T distinct and carefully crafted LSH functions (instead of just one LSH function), one can improve the likelihood to find at least one nearest neighbor while often significantly reducing the complexity of approximate NNS, even for very large datasets.

C. Fast LSH via the Randomized STOne Transform

The literature describes a number of ways to construct LSH functions with the desired properties [21]. In what follows, we

are interested in LSH functions that can be computed at low complexity and require low memory footprint—our goal is to reduce the complexity of LSH-based location fingerprinting.

Our approach builds upon prominent LSH functions for normalized datapoints [25] and uses ideas from [24], [26] to lower the complexity. In [25], the authors introduce a family of hash functions called cross polytope LSH for applications in which data points are distinguished by angular distance. For cross polytope LSH, one hashes a point $\mathbf{x} \in \mathbb{R}^D$ that lies on the unit sphere by computing $\mathbf{y} = \mathbf{A}\mathbf{x}/\|\mathbf{A}\mathbf{x}\|$, where $\mathbf{A} \in \mathbb{R}^{D \times D}$ is a random matrix with i.i.d. Gaussian entries. The closest standard basis vector of $\mathbb{R}^D$ to $\mathbf{y}$ is then used as the hash of $\mathbf{x}$. To make cross polytope hash functions more practical, reference [25] suggests the use of pseudo-random rotations. Concretely, instead of multiplying an input vector $\mathbf{x}$ by a random Gaussian matrix, one can “fake multiply” with a Gaussian matrix [26] using **HD**, where **H** is a Hadamard matrix (which has a fast transform) and **D** is a diagonal matrix with a fixed but random sequence of antipodal entries.

In our application, we select a diagonal matrix $\mathbf{D} \in \mathbb{R}^{D \times D}$ in which the diagonal entries are pseudo-random ± 1 values. We also select a pseudo-random index set $\Omega \subset \{1, 2, \dots, D\}$, where $|\Omega| = L$ will be the length (dimension) of the hash values. For **H**, we use the sum-to-one (STOne) transform matrix $\mathbf{H} \in \mathbb{R}^{D \times D}$, which is a Hadamard transform that has (i) $D \log(D)$ complexity and (ii) multi-scale properties [24]. The STOne transform matrix can be constructed using a 4×4 stencil matrix defined as [24]

$$S_4 = \frac{1}{2} \begin{bmatrix} -1 & 1 & 1 & 1 \\ 1 & -1 & 1 & 1 \\ 1 & 1 & -1 & 1 \\ 1 & 1 & 1 & -1 \end{bmatrix}. \quad (1)$$

One then forms a larger STOne transform matrix by computing the Kronecker product of two (or more) stencil matrices, e.g., $S_{16} = S_4 \otimes S_4$ where $\otimes$ is the Kronecker product.

With these ingredients, we can design the following LSH function: $h(\mathbf{f}) = [\text{sign}(\mathbf{H}\mathbf{D}\mathbf{f})]_\Omega$. Here, the operator $[\cdot]_\Omega$ extracts the vector whose entries are associated to the indices in Ω ; $\text{sign}(\cdot)$ operates element-wise on vectors. For this LSH function, each hash value (also called bucket) is in the set $\mathcal{S} = \{-1, +1\}^L$; this means that our hash values are of dimension L . Since $\tilde{\mathbf{f}} = \mathbf{D}\mathbf{f}$ can be computed in linear time and the STOne transform $\mathbf{H}\tilde{\mathbf{f}}$ in $D \log(D)$ time [24], computing hash values $h(\mathbf{f})$ is efficient and requires a minimum amount of storage (only for diagonal entries of **D** and the subset Ω).

To design multiple hash functions and tables, we simply generate T pseudo-random subsets $\Omega_t \subset \{1, \dots, D\}$, $t = 1, \dots, T$. We then compute T hash functions as $h_t(\mathbf{f}) = [\text{sign}(\mathbf{H}\mathbf{D}\mathbf{f})]_{\Omega_t}$, $t = 1, \dots, T$, which requires the computation of $\text{sign}(\mathbf{H}\mathbf{D}\mathbf{f})$ only once (instead of computing T Hadamard transforms).

D. LSH with Approximate Matches

Our final ingredient addresses the storage requirements of LSH-based location fingerprinting. To this end, we declare a match whenever the Hamming distance between $h(\mathbf{p})$ and $h(\mathbf{q})$ is within a given threshold $\delta \geq 0$ (instead of declaring

Fig. 1. Accuracy of LSH-based location fingerprinting for LoS and non-LoS massive MU-MIMO scenarios. "NN in feature space" are the nearest neighbors found through an exhaustive search and "NN in real space" are the true nearest neighbors in real space.

a match only if the two hash values are equal). Note that for $\delta = 0$, we perform classical LSH. The key advantage of approximate matches (i.e., $\delta > 0$) is that fewer hash tables are sufficient to achieve the same performance (in terms of finding the K nearest neighbors) as conventional LSH; see the next section for more details. Furthermore, recent advances in processing-in-memory (PIM) enables one to quickly identify approximate hash matches in a hardware efficient manner [27].

III. RESULTS AND COMPARISONS

We now show results of the proposed low-complexity LSH-based fingerprinting localization method. We provide a detailed accuracy/complexity trade-off analysis and a comparison to recent neural-network localization approaches [15]–[17]. Our main goals are as follows: (i) reduce the complexity of searching similar CSI fingerprints in large databases and (ii) reduce the complexity at minimal storage overhead.

A. Simulated Scenario

To evaluate the efficacy of our approach, we consider a massive MU-MIMO-OFDM localization scenario in LoS and non-LoS scenarios with a single basestation containing 32 antennas operating at 2.68 GHz with a bandwidth of 20 MHz and localizing 2000 transmitters distributed uniformly at random in an area of 40,000 m²; the noisy channel vectors are generated using channel models from [28]. The CSI features are $D = 256$ dimensional (32 antennas and 8 maximally-spaced subcarriers) and correspond to the absolute value of

beamspace/delay domain channel vectors as in [16], [19]. We use the $K = 2$ nearest neighbors and average their location; we also use approximate matches with a Hamming distance of $\delta = 1$. Both of these algorithm choices yield consistently good results. Below, we provide simulation results that confirm this claim. To assess the complexity and storage requirements, we use the following metrics: (i) the fraction of high-dimensional vectors that have been compared relative to an exhaustive search; (ii) the memory area indicating the size of the precomputed LSH tables (the total memory depth times the length L of each hash value); and (iii) the hardware cost, which is the fraction of comparisons times the memory area.

B. Accuracy vs. Complexity vs. Memory Trade-offs

Figure 1 shows results for a LoS and a non-LoS scenario. As a reference, we include the performance of an exhaustive search (denoted by "NN in feat. space") and that of finding the true nearest neighbors in real space (denoted by "NN in real space"). Each curve is parametrized by the number of hash functions T . We see in Fig. 1(a) that increasing the hash length L can reduce the number of comparisons by 10 $\times$, while achieving the same average distance error as an exhaustive search (indicated by the dashed "NN in feat. space" lines). As shown in Fig. 1(b), we see that increasing L increases the storage requirements for the hash values. This trade-off between reduced complexity but increased memory area is a result of the increase in the number of required hash tables but more sparsely populated hash buckets. Similarly, we see that

Fig. 2. Hardware cost vs. average distance error of LSH-based location fingerprinting with approximate Hamming matches $\delta \in \{0, 1, 2\}$ for a LoS scenario. A threshold of $\delta = 1$ yields consistently good results.

increasing the number of hash functions T reduces the average distance error but increases the number of comparisons and the memory area. The hardware cost is defined as fraction of comparisons times the memory area and is shown in Fig. 1(c), which reveals that the best trade-off in terms of fraction of comparisons and memory area is achieved by relatively small hash values, i.e., with $L = 12$ bits, and a few hash tables, i.e., with $T = 4$ tables. Figures 1(d,e,f) show results for a non-LoS scenario. The trends are similar to that of the LoS case, which confirms that the proposed method is able to reduce the NNS complexity of LSH-based fingerprinting localization under different propagation conditions.

Figure 2 shows results for the LoS scenario where we compare different Hamming-distance thresholds δ to find exact or approximate matches. We see that increasing δ from 0 to 1 reduces the average distance error for the same hardware costs as for $\delta = 0$. Increasing δ to 2 yields only a marginal improvement in terms of the average distance error, confirming that $\delta = 1$ is a reasonable choice. We would like to reiterate that approximate hash matches can be found efficiently using specialized PIM hardware, such as the one in [27].

C. Comparison to Deep Neural Networks

We now compare the performance of our LSH-based approach to that of a deep neural network in terms of complexity and memory area. We use fully connected neural networks (FCNNs) with different numbers of layers (1 to 6) to compare accuracy, complexity, and memory area. The 6-layer FCNN has been used in [18] and consists of 512, 256, 128, 64, 32, and 2 activations per layer; each layer uses rectified linear unit (ReLU) except the last layer, which uses a linear activation. To obtain the FCNNs with 5 to 1 layers, we successively remove the layer with the largest number of nodes.

To evaluate the complexity of LSH-based localization and FCNNs, we compute the total number of multiplications and additions required to find the approximate position of a single test point. In our LSH-based approach, we compute the complexity using the average number of comparisons needed for a given test point. The memory area for both approaches is the total number of stored values and parameters.

Figures 3(a,b,c,d) show results for LoS and non-LoS scenarios. As a reference, we show the performance of an exhaustive search and that of finding true nearest neighbors in real space. For the LoS scenario, Figs. 3(a,b) show that the LSH-based approach can achieve lower average distance error with lower complexity than the neural network based approach. In terms of memory area, we observe that the FCNN consistently performs better since it only needs to store weights and bias values, whereas the LSH-based approach requires storage for the entire fingerprint/location database so that comparisons can be made to points in CSI feature space. For the non-LoS scenario, Figs. 3(c,d) show that the LSH-based approach no longer outperforms the neural network in terms of search complexity, i.e., FCNNs achieve slightly better average distance error with about the same complexity as the LSH-based approach. We see similar trends in memory area as we did in the LoS scenario.

IV. CONCLUSIONS

We have demonstrated that the complexity of fingerprinting-based localization can be reduced significantly by using locality-sensitive hashing (LSH). To this end, we have proposed computationally efficient LSH functions that build upon the sum-to-one (STOne) transform and approximate hash matches. Our results have shown that LSH-based fingerprinting is able to achieve the same accuracy as an exhaustive search over the CSI fingerprinting database for line-of-sight (LoS) and non-LoS scenarios, but at orders-of-magnitude lower complexity and storage requirements. We have also demonstrated that LSH-based positioning is able to outperform neural network-based approaches under LoS propagation conditions. However, we have not observed any advantage of our approach for non-LoS propagation conditions. These observations imply that (i) the best approach for fingerprinting positioning depends strongly on the propagation conditions and (ii) neural networks are not necessarily the best approach for localization in terms of complexity and positioning accuracy.

REFERENCES

- [1] F. Gustafsson and F. Gunnarsson, "Mobile positioning using wireless networks: Possibilities and fundamental limitations based on available wireless network measurements," *IEEE Signal Process. Mag.*, vol. 22, no. 4, pp. 41–53, Jul. 2005.
- [2] H. Liu, H. Darabi, P. Banerjee, and J. Liu, "Survey of wireless indoor positioning techniques and systems," *IEEE Trans. on Systems, Man, and Cybernetics, Part C (Applications and Reviews)*, vol. 37, no. 6, pp. 1067–1080, Nov. 2007.
- [3] S. Kumar, S. Gil, D. Katabi, and D. Rus, "Accurate indoor localization with zero start-up cost," in *Int. Conf. Mobile Comput.*, Sep. 2014, pp. 483–494.
- [4] N. Garcia, H. Wymeersch, E. G. Larsson, A. M. Haimovich, and M. Coulon, "Direct localization for massive MIMO," *IEEE Trans. on Signal Process.*, vol. 65, no. 10, pp. 2475–2487, May 2017.
- [5] P. Bahl, V. N. Padmanabhan, V. Bahl, and V. Padmanabhan, "RADAR: an in-building RF-based user location and tracking system," in *IEEE Conf. on Computer Commun. (INFOCOM)*, Mar. 2000.
- [6] K. Kaemarungsi and P. Krishnamurthy, "Modeling of indoor positioning systems based on location fingerprinting," in *IEEE Infocom*, vol. 2, Mar. 2004, pp. 1012–1022.
- [7] —, "Properties of indoor received signal strength for WLAN location fingerprinting," in *Intl. Conf. on Mobile and Ubiquitous Syst.: Networking and Services (MOBIQUITOUS)*, Aug. 2004, pp. 14–23.

Fig. 3. Comparison between LSH-based fingerprinting localization and neural network-based positioning. For LoS propagation conditions, LSH-based positioning reduces complexity while achieving superior accuracy; the memory requirements for neural network-based positioning is generally lower. For non-LoS propagation conditions, neural network-based positioning slightly outperforms LSH-based positioning in terms of accuracy, complexity, and memory.

- [8] M. Bshara, U. Orguner, F. Gustafsson, and L. Van Biesen, "Fingerprinting localization in wireless networks based on received-signal-strength measurements: A case study on WiMAX networks," *IEEE Trans. on Veh. Tech.*, vol. 59, no. 1, pp. 283–294, Jan. 2010.
- [9] M. Ibrahim and M. Youssef, "CellSense: An accurate energy-efficient GSM positioning system," *IEEE Trans. on Veh. Tech.*, vol. 61, no. 1, pp. 286–296, Jan. 2012.
- [10] X. Wang, L. Gao, S. Mao, and S. Pandey, "DeepFi: Deep learning for indoor fingerprinting using channel state information," in *IEEE Wireless Commun. and Networking Conf. (WCNC)*, Mar. 2015, pp. 1666–1671.
- [11] K. Wu, J. Xiao, Y. Yi, D. Chen, X. Luo, and L. M. Ni, "CSI-based indoor localization," *IEEE Trans. on Parallel and Distributed Syst.*, vol. 24, no. 7, pp. 1300–1309, 2013.
- [12] Y. Chapre, A. Ignjatovic, A. Seneviratne, and S. Jha, "CSI-MIMO: An efficient Wi-Fi fingerprinting using channel state information with MIMO," *Elsevier Pervasive and Mobile Computing*, vol. 23, pp. 89–103, Oct. 2015.
- [13] A. Smailagic, J. Small, and D. P. Siewiorek, "Determining user location for context aware computing through the use of a wireless LAN infrastructure," *Institute for Complex Engineered Systems Carnegie Mellon University, Pittsburgh, PA*, vol. 15213, Dec. 2000.
- [14] P. Castro, P. Chiu, T. Kremenek, and R. Muntz, "A probabilistic room location service for wireless networked environments," in *Int'l Conf. Ubiquitous Computing*. Springer, Dec. 2001, pp. 18–34.
- [15] X. Wang, L. Gao, and S. Mao, "CSI phase fingerprinting for indoor localization with a deep learning approach," *IEEE Internet of Things Journal*, vol. 3, no. 6, pp. 1113–1123, Dec. 2016.
- [16] J. Vieira, E. Leitingner, M. Sarajlic, X. Li, and F. Tufvesson, "Deep convolutional neural networks for massive MIMO fingerprint-based positioning," in *IEEE Intl. Symp. Personal, Indoor, and Mobile Radio Commun. (PIMRC)*, Oct. 2017, pp. 1–6.
- [17] M. Arnold, S. Dörner, S. Cammerer, and S. Ten Brink, "On deep learning-based massive MIMO indoor user localization," in *IEEE Int'l Workshop on Signal Process. Advances in Wireless Commun. (SPAWC)*, Jun. 2018, pp. 1–5.
- [18] E. Lei, O. Castañeda, O. Tirkkonen, T. Goldstein, and C. Studer, "Siamese neural networks for wireless positioning and channel charting," in *57th Annual Allerton Conf. Commun., Control, and Comput.*, Sep. 2019.
- [19] C. Studer, S. Medjkouh, E. Gönültaş, T. Goldstein, and O. Tirkkonen, "Channel charting: Locating users within the radio environment using channel state information," *IEEE Access*, vol. 6, pp. 47 682–47 698, Aug. 2018.
- [20] A. Gionis, P. Indyk, and R. Motwani, "Similarity search in high dimensions via hashing," in *Very Large Data Base (VLDB)*, vol. 99, no. 6, Sep. 1999, pp. 518–529.
- [21] A. Andoni and P. Indyk, "Near-optimal hashing algorithms for approximate nearest neighbor in high dimensions," *Communications of the ACM*, vol. 51, no. 1, p. 117, 2008.
- [22] M. S. Charikar, "Similarity estimation techniques from rounding algorithms," in *Proc. of the 34th Ann. ACM Symp. on Th. of Comp.*, May 2002, pp. 380–388.
- [23] S. Har-Peled, P. Indyk, and R. Motwani, "Approximate nearest neighbor: Towards removing the curse of dimensionality," *Theory of computing*, vol. 8, no. 1, pp. 321–350, 2012.
- [24] T. Goldstein, L. Xu, K. F. Kelly, and R. Baraniuk, "The STOne transform: Multi-resolution image enhancement and compressive video," *IEEE Trans. on Image Process.*, vol. 24, no. 12, pp. 5581–5593, Dec. 2015.
- [25] A. Andoni, P. Indyk, T. Laarhoven, I. Razenshteyn, and L. Schmidt, "Practical and optimal LSH for angular distance," in *Advances in Neural Inf. Process. Syst. (NIPS)*, Dec. 2015, pp. 1225–1233.
- [26] M. Kapralov, V. Potluru, and D. Woodruff, "How to fake multiply by a Gaussian matrix," in *Intl. Conf. on Machine Learning (ICML)*, June 2016, pp. 2101–2110.
- [27] O. Castañeda, M. Bobbitt, A. Gallyas-Sanhueza, and C. Studer, "PPAC: A versatile in-memory accelerator for matrix-vector product-like tasks," in *IEEE Intl. Conf. on Application-Specific Systems, Architectures, and Processors (ASAP)*, Jul. 2019.
- [28] S. Jaeckel, L. Raschkowski, K. Börner, L. Thiele, F. Burkhardt, and E. Eberlein, "QuaDRiGa - Quasi Deterministic Radio Channel Generator User Manual and Documentation," Fraunhofer Heinrich Hertz Institute, Tech. Rep. v2.0.0, 2017.