CoRa

A Collision-Resistant LoRa Symbol Detector of Low Complexity

José I. Álamos[1], Thomas C. Schmidt[1], Matthias Wählisch[2].
[1]Hamburg University of Applied Sciences, Germany
[2]TU Dresden & Barkhausen Institut, Germany

INFOCOM 2025
London, UK

LoRa Wireless Modulation Technique

Long range (up to 15 km)

High time on air (up to seconds)

Low power consumption (mJ)

Motivation

- LoRaWAN uses uncoordinated communication (ALOHA)

Draft IG LPWA Report (IEEE P802.15-17-0528-00-lpwa)

Motivation

- LoRaWAN uses uncoordinated communication (ALOHA)
- Long range increases aggregation at LoRaWAN gateways

Region	**Population ($\frac{1}{km^2}$)**	**10km-Radius Mean Arrival ($\frac{1}{s}$)**
Paris	21000	18325
London	5518	4815
Berlin	4000	3490

Draft IG LPWA Report (IEEE P802.15-17-0528-00-lpwa)

Motivation

- LoRaWAN uses uncoordinated communication (ALOHA)
- Long range increases aggregation at LoRaWAN gateways
- **High collision rate** due to long time on air

Region	Population ($\frac{1}{km^2}$)	10km-Radius Mean Arrival ($\frac{1}{s}$)
Paris	21000	18325
London	5518	4815
Berlin	4000	3490

Draft IG LPWA Report (IEEE P802.15-17-0528-00-lpwa)

Existing LoRa Receivers

- Baseline fails in multi-packet collisions

Existing LoRa Receivers

- Baseline fails in multi-packet collisions
- Prior work trades off **throughput**, **complexity** or needs **symbol boundary information**

Method	**Throughput**	**Complexity**	**Requires Symbol Boundary Info.**
AlignTrack (*ICNP. '21*)	Moderate	High	Yes
CIC (*SIGCOMM '21*)	High	High	Yes
TnB (*CoNEXT '22*)	High	Low	Yes
CoRa	High	Low	No

Scope

- Focus exclusively on the **demodulation** process

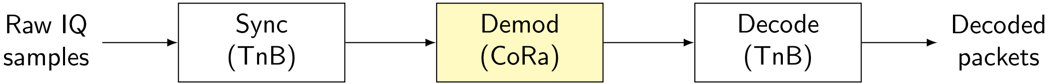

Scope

- Focus exclusively on the **demodulation** process
- No reliance on successful detection of colliding frames

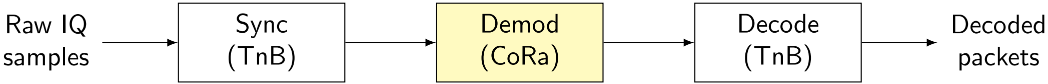

Agenda

Background

Core Idea

Evaluation

Conclusion

LoRa Modulation

Chirp Spread Spectrum

$$Y_0(t) = e^{j(\pi B t^2)}$$

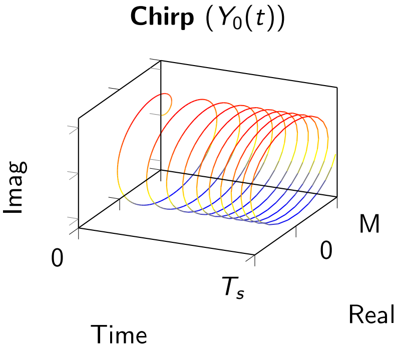

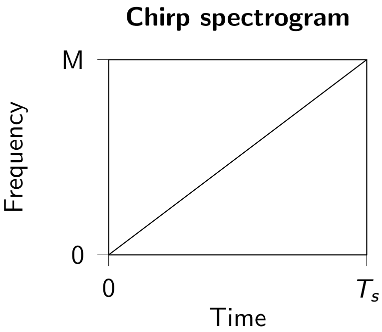

Dechirp Procedure

LoRa symbol ($X(t)$)

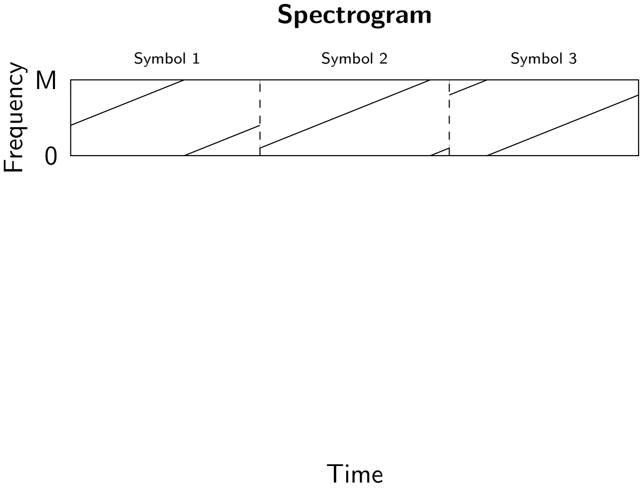

Dechirp Procedure

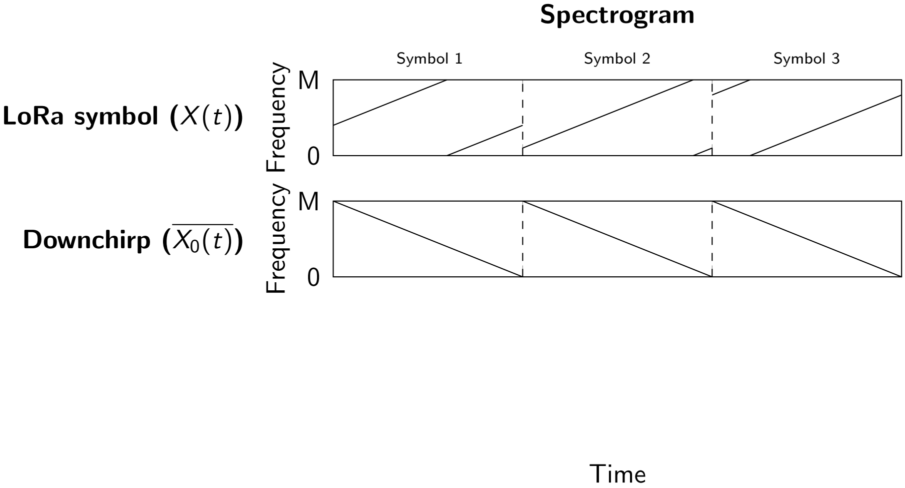

Dechirp Procedure

LoRa symbol ($X(t)$)

Downchirp ($\overline{X_0(t)}$)

Dechirped symbol ($X(t) \cdot \overline{X_0(t)}$)

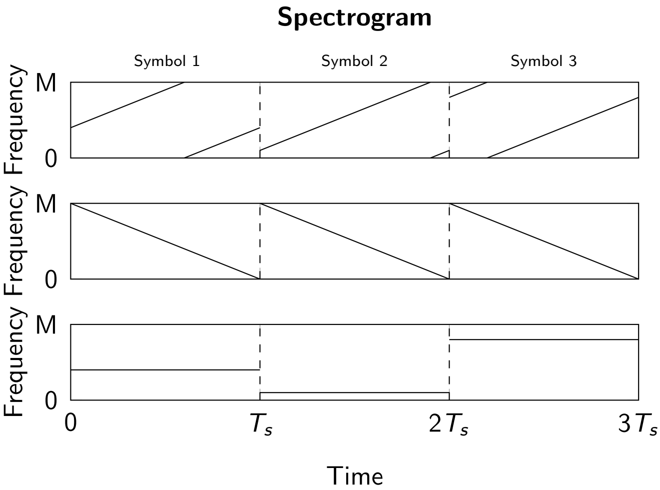

Extracting Symbol Value from Dechirped Signal

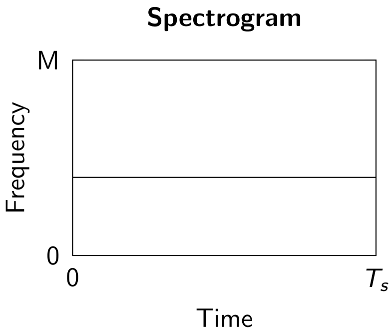

Extracting Symbol Value from Dechirped Signal

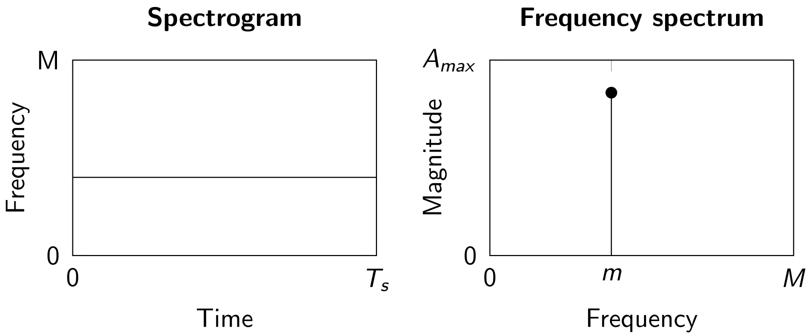

Extracting Symbol Value from Dechirped Signal

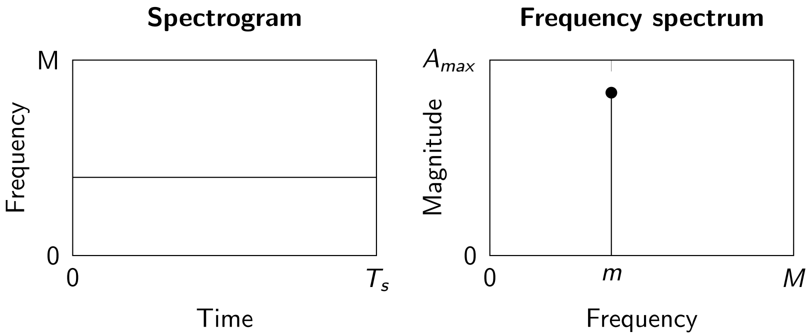

Select frequency bin with the highest magnitude

LoRa Collisions

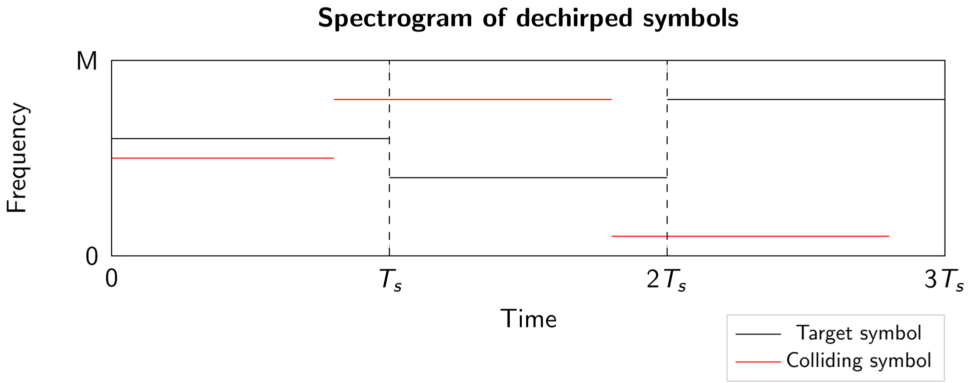

DFT Artifacts from Collisions

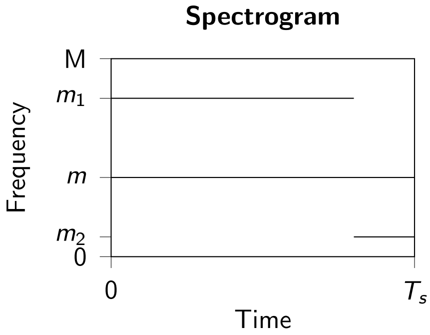

DFT Artifacts from Collisions

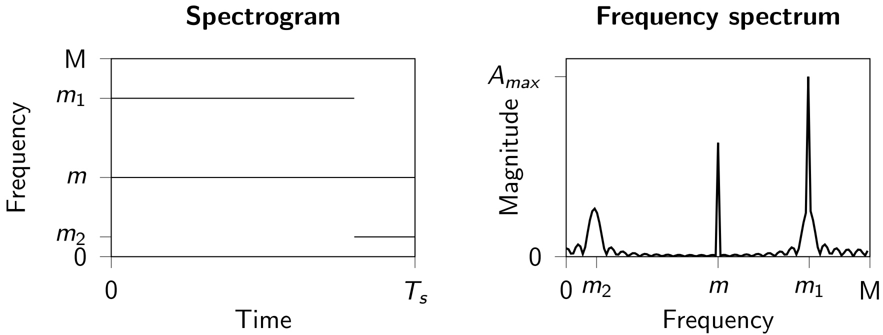

DFT Artifacts from Collisions

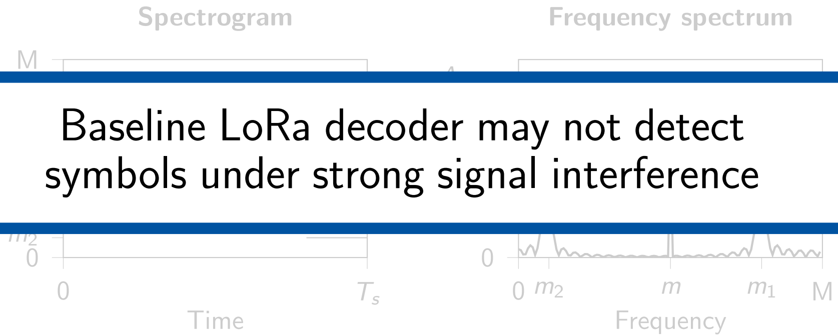

Baseline LoRa decoder may not detect symbols under strong signal interference

Goal

Detect the **complete** waveform in the dechirped signal within the demodulation window

Decomposition of a dechirped LoRa signal under collision

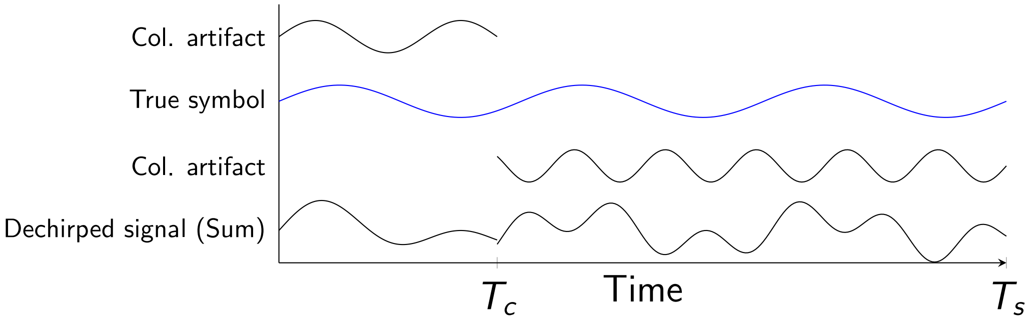

Detecting Complete Waveforms

Proposal: Perform cross-correlation with a complex waveform template, inverted in the second half

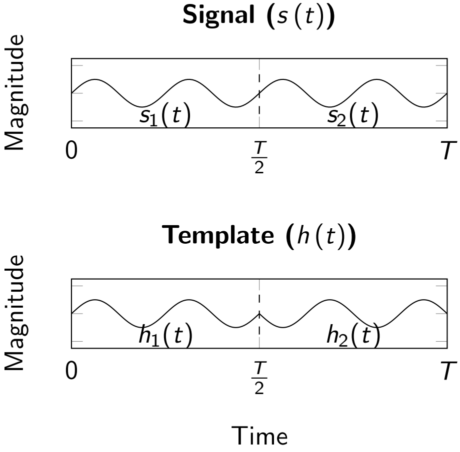

Detecting Complete Waveforms

Proposal: Perform cross-correlation with a complex waveform template, inverted in the second half

- **1. Linearity:**

$$\langle s, h\rangle = \langle s_1, h_1\rangle + \langle s_2, h_2\rangle$$

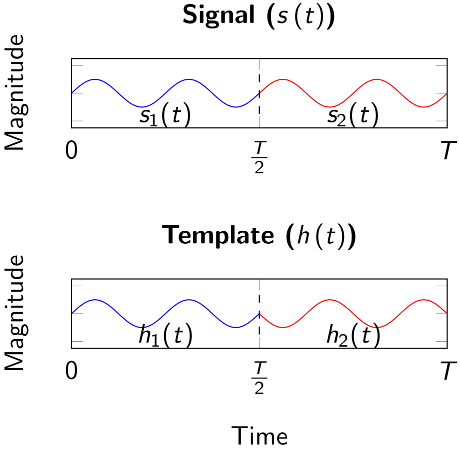

Detecting Complete Waveforms

Proposal: Perform cross-correlation with a complex waveform template, inverted in the second half

- **1. Linearity:**

$$\langle s, h \rangle = \langle s_1, h_1 \rangle + \langle s_2, h_2 \rangle$$

- **2. Phase-Shifting:**

$$s_1(t) = k \cdot h_1(t) \quad , \quad s_2(t) = -k \cdot h_2(t)$$

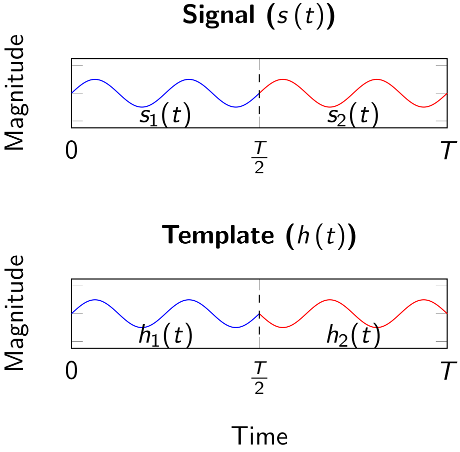

Detecting Complete Waveforms

Proposal: Perform cross-correlation with a complex waveform template, inverted in the second half

- **1. Linearity:**

$$\langle s, h\rangle = \langle s_1, h_1\rangle + \langle s_2, h_2\rangle$$

- **2. Phase-Shifting:**

$$s_1(t) = k \cdot h_1(t) \quad , \quad s_2(t) = -k \cdot h_2(t)$$

- **Result:**

$$\langle s_2, h_2\rangle = -\langle s_2, h_1\rangle \quad \Rightarrow \quad \langle s, h\rangle = 0$$

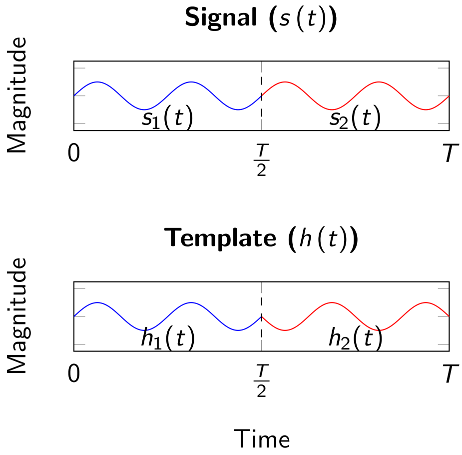

Efficient Cross-Correlation with DFT Tricks

Multiply the dechirped symbol by a phase-shifted mask $m[n]$:

$$\tilde{x}[n] = x[n] \cdot m[n]$$

Then compute:

$$\tilde{X}_k = \text{DFT}\{\tilde{x}[n]\}$$

where

$$m[n] = \begin{cases} +1, & n < N/2 \\ -1, & n \geq N/2 \end{cases}$$

⇒ *Gives correlation with phase-shifted waveform at each k*

Derivation of Half-Period Discriminator (HPD)

- Phase-shifted mask introduces spectral leakage and therefore cannot be used directly as a feature

Phase-shifted DFT ($\tilde{X}_k$)

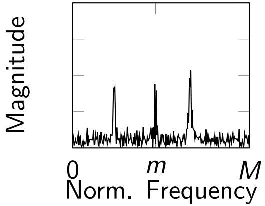

Derivation of Half-Period Discriminator (HPD)

- Phase-shifted mask introduces spectral leakage and therefore cannot be used directly as a feature
- **Solution:** Pointwise minimum between X_k and $\tilde{X}_k$

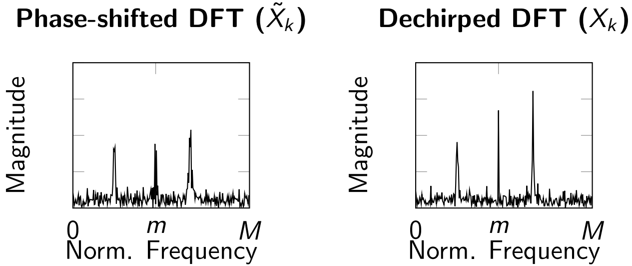

Derivation of Half-Period Discriminator (HPD)

- Phase-shifted mask introduces spectral leakage and therefore cannot be used directly as a feature
- **Solution:** Pointwise minimum between X_k and $\tilde{X}_k$

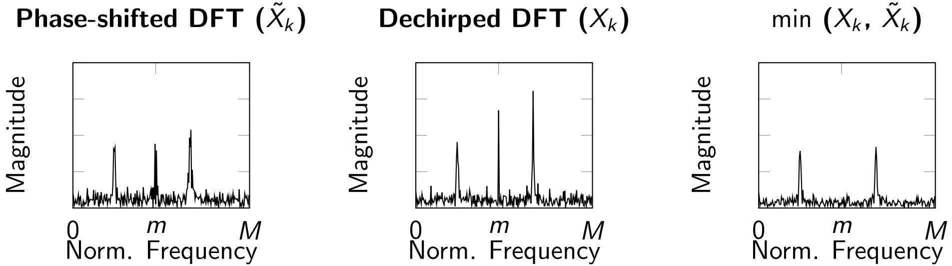

Derivation of Half-Period Discriminator (HPD)

- Phase-shifted mask introduces spectral leakage and therefore cannot be used directly as a feature
- **Solution:** Pointwise minimum between X_k and $\tilde{X}_k$
- HPD: ratio of minimum-filtered to original DFT

$$\mathbf{h}_k = \frac{\min\left(|X_k|, |\tilde{X}_k|\right)}{|X_k|}$$

Phase-shifted DFT ($\tilde{X}_k$)

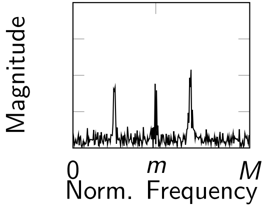

Dechirped DFT (X_k)

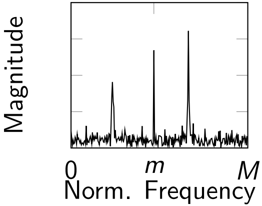

min (X_k, $\tilde{X}_k$)

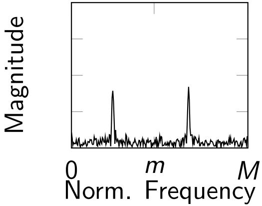

HPD Alone Is Insufficient

Noise-induce spurious peaks yield false positives

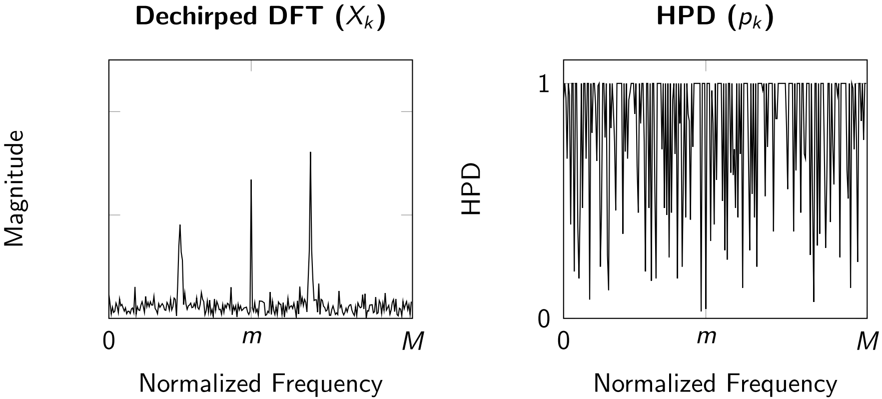

Extending Feature Space: Peak Magnitude Deviation (PMD)

- Select peaks in the dechirped spectrum based on how closely their magnitude matched the expected peak magnitude (measured from preamble).

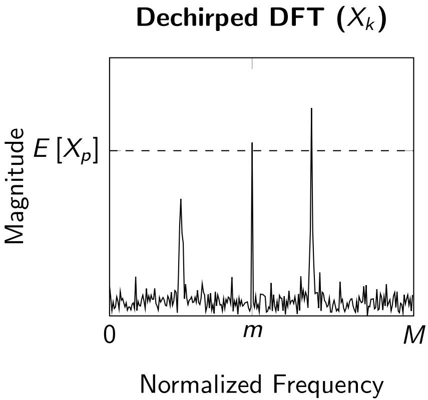

Extending Feature Space: Peak Magnitude Deviation (PMD)

- Select peaks in the dechirped spectrum based on how closely their magnitude matched the expected peak magnitude (measured from preamble).
- We define the Peak Magnitude Deviation as:

$$\mathbf{p}_k = \min\left(\frac{||X_k| - E[X_p]|}{E[X_p]}, 1\right)$$

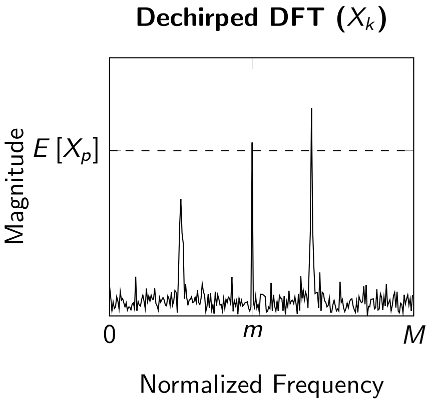

Extending Feature Space: Peak Magnitude Deviation (PMD)

- Select peaks in the dechirped spectrum based on how closely their magnitude matched the expected peak magnitude (measured from preamble).
- We define the Peak Magnitude Deviation as:

$$\mathbf{p}_k = \min\left(\frac{||X_k| - E\left[X_p\right]|}{E\left[X_p\right]}, 1\right)$$

- The true symbol peak yields a PMD value near zero

Dechirped DFT (X_k)

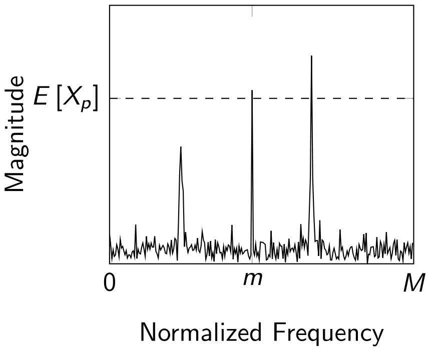

PMD Alone Is Insufficient

Collision peaks may match expected magnitude

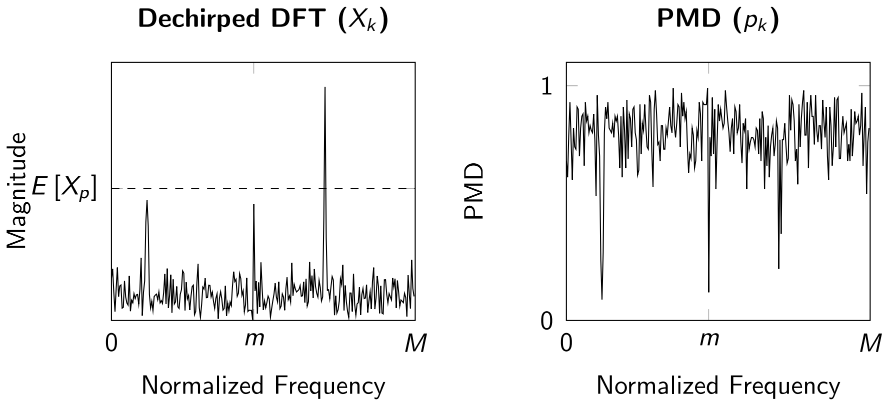

Combining Features: Symbol Classifier

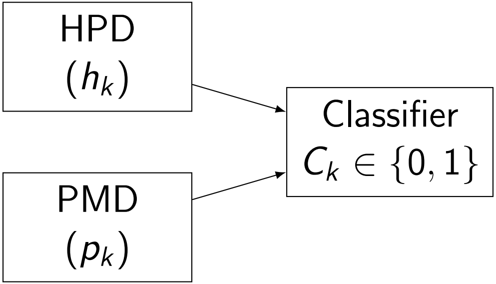

$C_k = 1 \rightarrow$ valid symbol candidate $\quad C_k = 0 \rightarrow$ invalid (e.g., noise or collision)

Bayesian Classification

- **Goal:** Estimate $P(C_k \mid p_k, h_k)$

Bayesian Classification

- **Goal:** Estimate $P(C_k \mid p_k, h_k)$
- Bayes theorem:

$$P(C_k \mid p_k, h_k) = \frac{P(p_k, h_k \mid C_k)\, P(C_k)}{P(p_k, h_k)}$$

Bayesian Classification

- **Goal:** Estimate $P(C_k \mid p_k, h_k)$
- Bayes theorem:

$$P(C_k \mid p_k, h_k) = \frac{P(p_k, h_k \mid C_k)\, P(C_k)}{P(p_k, h_k)}$$

- Likelihoods from simulated collisions

Bayesian Classification

- **Goal:** Estimate $P(C_k \mid p_k, h_k)$
- Bayes theorem:

$$P(C_k \mid p_k, h_k) = \frac{P(p_k, h_k \mid C_k)\, P(C_k)}{P(p_k, h_k)}$$

- Likelihoods from simulated collisions
- Priors from class frequencies

Likelihood of True Symbol

- For each candidate bin k, we compute the posterior probability:

$$P(C_k = 1 \mid p_k, h_k)$$

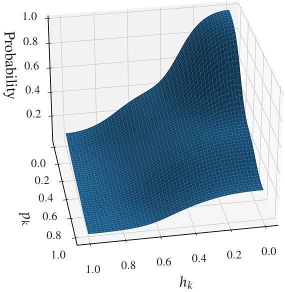

Likelihood of True Symbol

- For each candidate bin k, we compute the posterior probability:

$$P(C_k = 1 \mid p_k, h_k)$$

- True symbols cluster in regions with low PMD and low HPD.

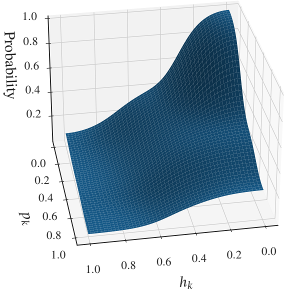

Output of Features and Bayes Classifier

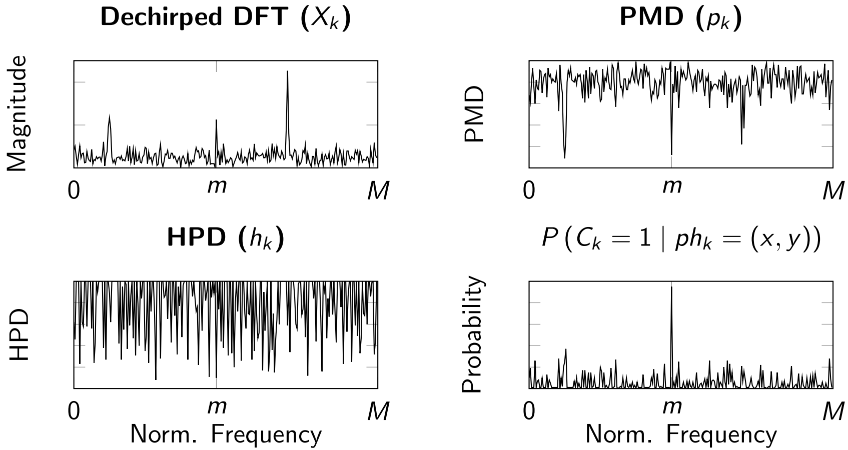

Preamble Collision

- Preamble peaks exhibit identical frequency across consecutive symbols

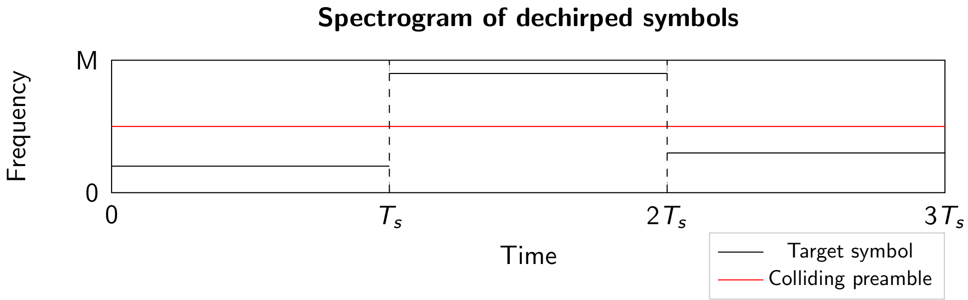

Preamble Collision

- Preamble peaks exhibit identical frequency across consecutive symbols
- Interference preamble peaks resemble valid symbols (full oscillation, correct magnitude).

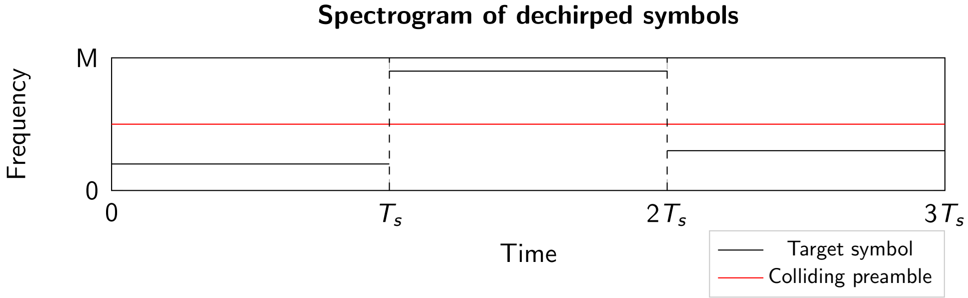

Preamble Filtering

- Peaks stable across consecutive symbols are likely preamble artifacts.

Preamble Filtering

- Peaks stable across consecutive symbols are likely preamble artifacts.
- We refine symbol demodulation by penalizing candidates present in both the current symbol T_m and the previous T_{m-1}.

$$P_{\text{final}}(k) = P_{T_m}(C_k = 1 \mid p_k, h_k) \cdot \left(1 - P_{T_{m-1}}(C_k = 1 \mid p_k, h_k)\right)$$

Preamble Filtering

- Peaks stable across consecutive symbols are likely preamble artifacts.
- We refine symbol demodulation by penalizing candidates present in both the current symbol T_m and the previous T_{m-1}.

$$P_{\text{final}}(k) = P_{T_m}(C_k = 1 \mid p_k, h_k) \cdot \left(1 - P_{T_{m-1}}(C_k = 1 \mid p_k, h_k)\right)$$

- Select k maximizing $P_{\text{final}}(k)$

Evaluation platform

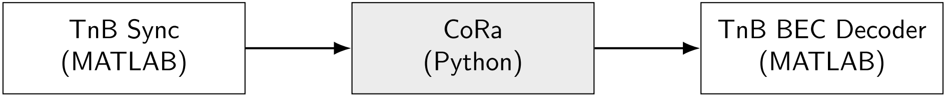

Target platform:
Intel(R) Core(TM) i7-7500U CPU @ 2.70GHz

Evaluation scenario

- Real-world SDR captures under varying SNR conditions

	CIC Dataset (*SIGCOMM '21*)				TnB Dataset (*CoNEXT '22*)		
Parameter	**D1**	**D2**	**D3**	**D4**	**Indoor**	**Outdoor1**	**Outdoor2**
Spreading Factor	8				8, 10		
Bandwidth [kHz]	250				125		
Coding Rate	4/5				4/5, 4/6, 4/7, 4/8		
Sampling Rate [MHz]	2				1		
SNR (SF8) [dB]	30 to 42	30 to 42	10 to 30	-17 to 5	-5 to 17	-10 to 10	-8 to 14
SNR (SF10) [dB]	-				-9 to 23	-17 to 12	-13 to 12
Payload size [B]	12				12		
Deployment area [m^2]	15 × 10	100 × 60	170 × 100	1200 × 1600	100 × 120	400 × 240	220 × 200
TX Rate [pkt/s]	5 to 100				20, 25		
Capture Time [s]	60				30		

Evaluation scenario

- Real-world SDR captures under varying SNR conditions
- **CIC dataset:** fixed SF and coding rate, multiple transmission rates, short time on air

	CIC Dataset (*SIGCOMM '21*)				**TnB Dataset** (*CoNEXT '22*)		
Parameter	**D1**	**D2**	**D3**	**D4**	**Indoor**	**Outdoor1**	**Outdoor2**
Spreading Factor	8				8, 10		
Bandwidth [kHz]	250				125		
Coding Rate	4/5				4/5, 4/6, 4/7, 4/8		
Sampling Rate [MHz]	2				1		
SNR (SF8) [dB]	30 to 42	30 to 42	10 to 30	-17 to 5	-5 to 17	-10 to 10	-8 to 14
SNR (SF10) [dB]	-				-9 to 23	-17 to 12	-13 to 12
Payload size [B]	12				12		
Deployment area [m^2]	15×10	100×60	170×100	1200×1600	100×120	400×240	220×200
TX Rate [pkt/s]	5 to 100				20, 25		
Capture Time [s]	60				30		

Evaluation scenario

- Real-world SDR captures under varying SNR conditions
- **CIC dataset:** fixed SF and coding rate, multiple transmission rates, short time on air
- **TnB dataset:** fixed rate, varying SF and coding rate

	CIC Dataset (*SIGCOMM '21*)				**TnB Dataset** (*CoNEXT '22*)		
Parameter	**D1**	**D2**	**D3**	**D4**	**Indoor**	**Outdoor1**	**Outdoor2**
Spreading Factor	8				8, 10		
Bandwidth [kHz]	250				125		
Coding Rate	4/5				4/5, 4/6, 4/7, 4/8		
Sampling Rate [MHz]	2				1		
SNR (SF8) [dB]	30 to 42	30 to 42	10 to 30	-17 to 5	-5 to 17	-10 to 10	-8 to 14
SNR (SF10) [dB]	-				-9 to 23	-17 to 12	-13 to 12
Payload size [B]	12				12		
Deployment area [m^2]	15×10	100×60	170×100	1200×1600	100×120	400×240	220×200
TX Rate [pkt/s]	5 to 100				20, 25		
Capture Time [s]	60				30		

Throughput evaluation: CIC Dataset

- Up to 60 pkt/s under high SNR, and up to 40 pkt/s below the noise floor

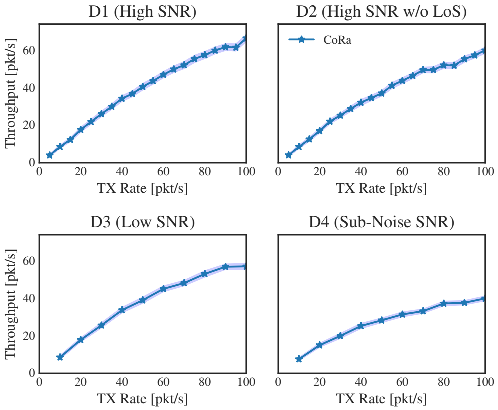

Throughput evaluation: CIC Dataset

- Up to 60 pkt/s under high SNR, and up to 40 pkt/s below the noise floor
- Outperforms TnB, CIC and the baseline

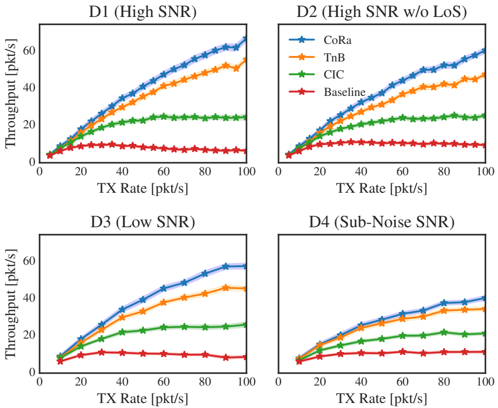

Throughput evaluation: TnB Dataset

- CoRa outperforms other alternatives at highest coding rate

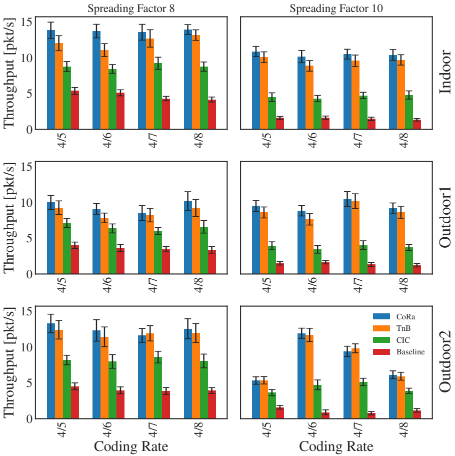

Throughput evaluation: TnB Dataset

- CoRa outperforms other alternatives at highest coding rate
- Best gain at low spreading factor and high SNR

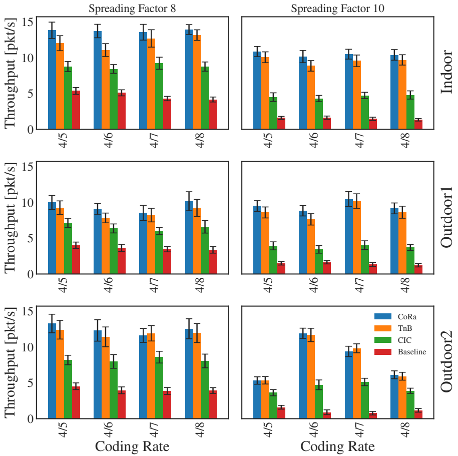

Throughput evaluation: TnB Dataset

- CoRa outperforms other alternatives at highest coding rate
- Best gain at low spreading factor and high SNR
- Matches TnB under low SNR, high spreading factors and moderate coding rate

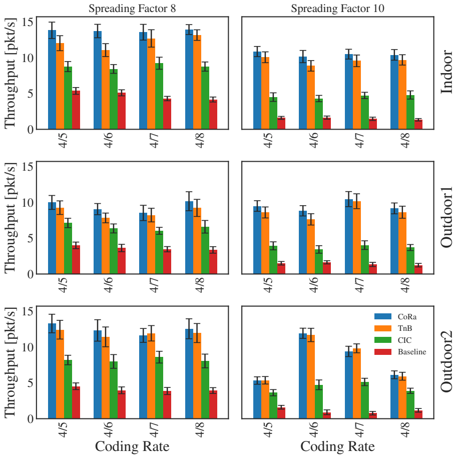

Performance Overhead

- CoRa runs in bounded time
  - $\approx 3\times$ the baseline demodulation
  - Still less than 2.5% of symbol time

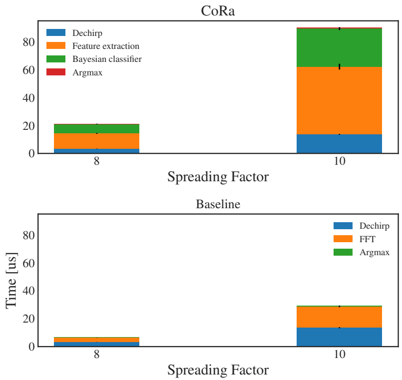

Performance Overhead

- CoRa runs in bounded time
  - $\approx 3\times$ the baseline demodulation
  - Still less than 2.5% of symbol time
- TnB and CIC are excluded due to non-comparable implementations
  - Their execution is unbounded, as it grows with the number of collisions

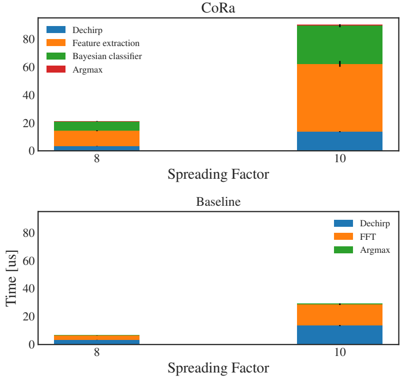

Conclusions

High throughput	Comparable or exceeding SOTA techniques
Low latency processing	Less than 2.5% of the symbol time
Symbol scoped	Robust against frame detection failures
Robust to low SNR	Reliable under weak signal conditions.

Thank you for your attention!

We fully support reproducible research and open source software

Source code & Artifacts

https://zenodo.org/records/14515243

Future work

- **Towards full packet recovery:** Explore new synchronization and decoding techniques to boost collision resolution.
- **Real-World Validation:** Implement the CoRa framework on a real SDR platform and evaluate performance under realistic deployment scenarios.

Throughput in simulated LTE ETU Channel[2]

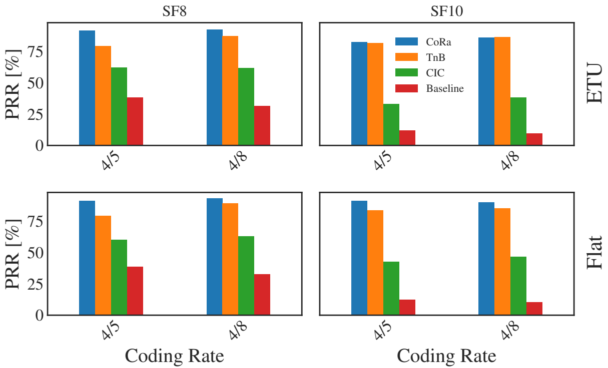

[2]: 3GPP, Technical Report 36.873 V12.7.0, 2017

CoRa

A Collision-Resistant LoRa Symbol Detector of Low Complexity

José I. Álamos[1], Thomas C. Schmidt[1], Matthias Wählisch[2].
[1]Hamburg University of Applied Sciences, Germany
[2]TU Dresden & Barkhausen Institut, Germany

INFOCOM 2025
London, UK

HAW HAMBURG TECHNISCHE UNIVERSITÄT DRESDEN barkhausen institut