I Send, Therefore I Leak: Information Leakage in Low-Power Wide Area Networks

Patrick Leu, ETH Zurich, Switzerland

Ivan Puddu, ETH Zurich, Switzerland

Aanjhan Ranganathan, Northeastern University, USA

Srdjan Čapkun, ETH Zurich, Switzerland

LPWANs address communication needs of IoT

- LPWANs provide communications to cheap, widely distributed end-devices
- Requirements on communications:
  - Cheap, easy, large-scale deployment
  - Long battery life
  - Long communication range, O(km)
  - Usually low data rate (periodic sensor readings, binary states, …)
  - No complex medium access protocol, avoid channel sensing
- LoRa, SigFox, NB-IOT, Weightless, ...

LoRaWAN™

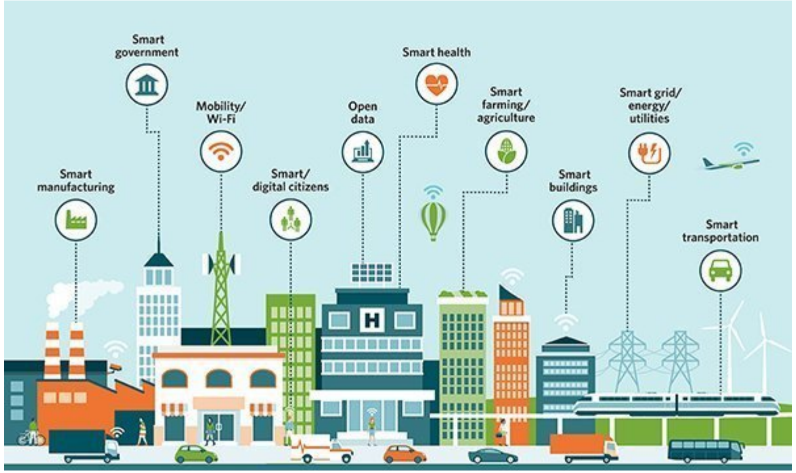

Problem: The mere existence of a transmission can leak sensitive information

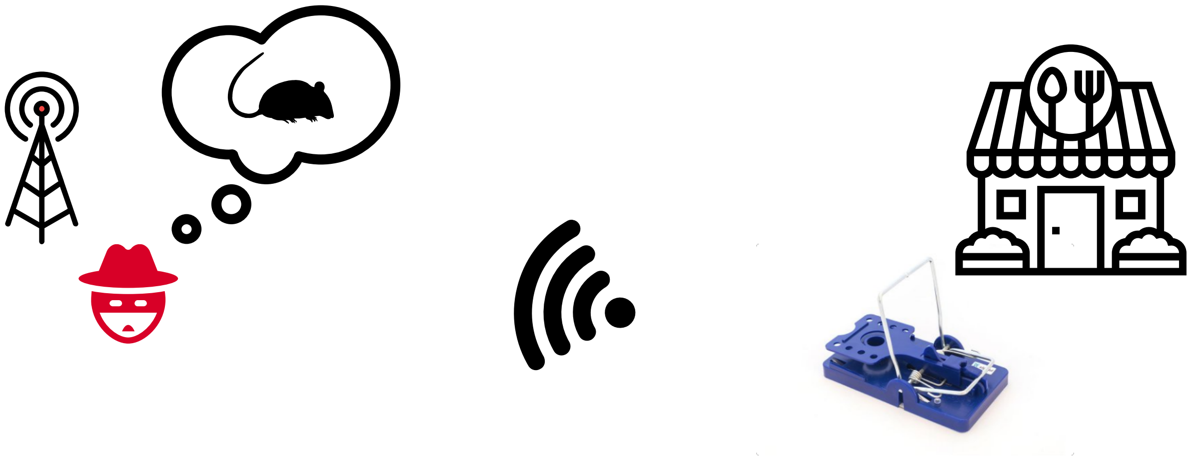

This is a fundamental difference to other wireless technologies, such as cellular or WiFi.

Event-driven communication leaks information

- Event-driven communication
  - Devices send upon sensing a real-world event: Push button, IR sensor, humidity sensor, ...
- Eavesdropping is easy, inexpensive and can be done from a distance
  - Robust encoding helps the eavesdropper.
  - LoRa PHY has been reverse engineered. SDRs can be used.
- Existing work in LoRa/LPWAN security:
  - Replay attack
  - Acknowledgement spoofing
  - Physical key extraction
  - Device fingerprinting
  - Reactive jamming
- **Privacy implications not studied so far.**

Contributions

- We show that event-driven communication in LPWANs inherently leaks information.
- We identify two classes of leakage.
- We show that full leakage prevention is very difficult as it involves high amounts of excess power.

LPWAN applications

Industrial applications

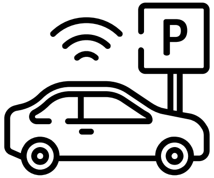

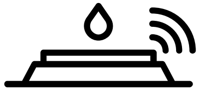

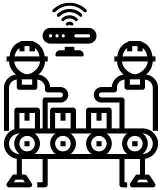

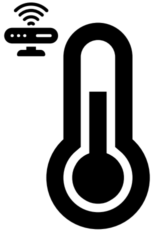

Smart homes and cities

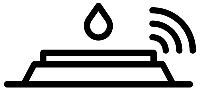

LPWAN architecture

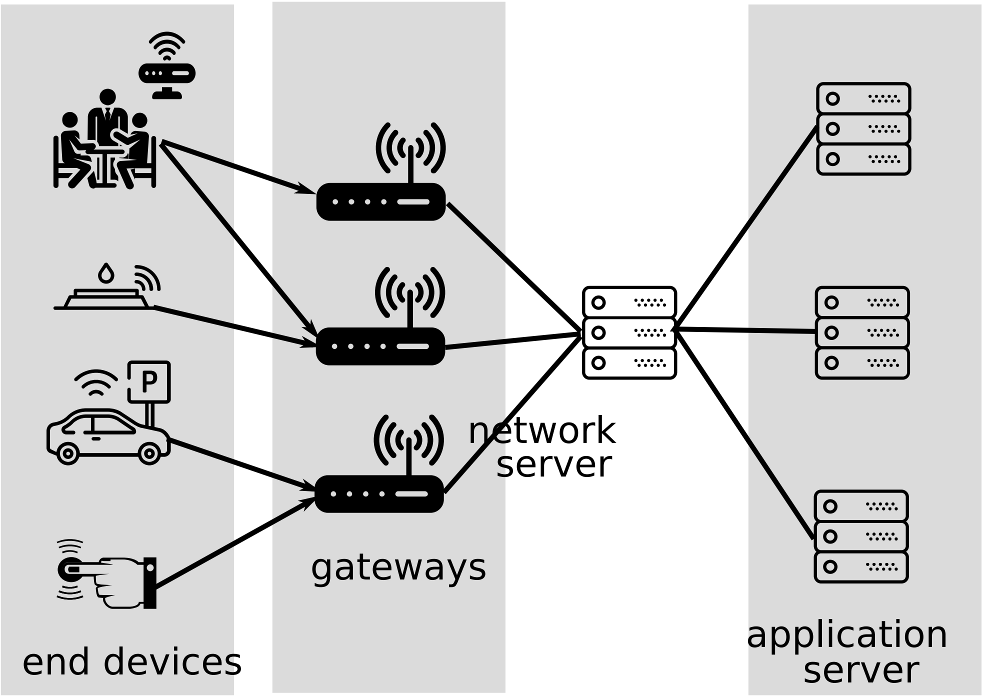

LPWAN architecture

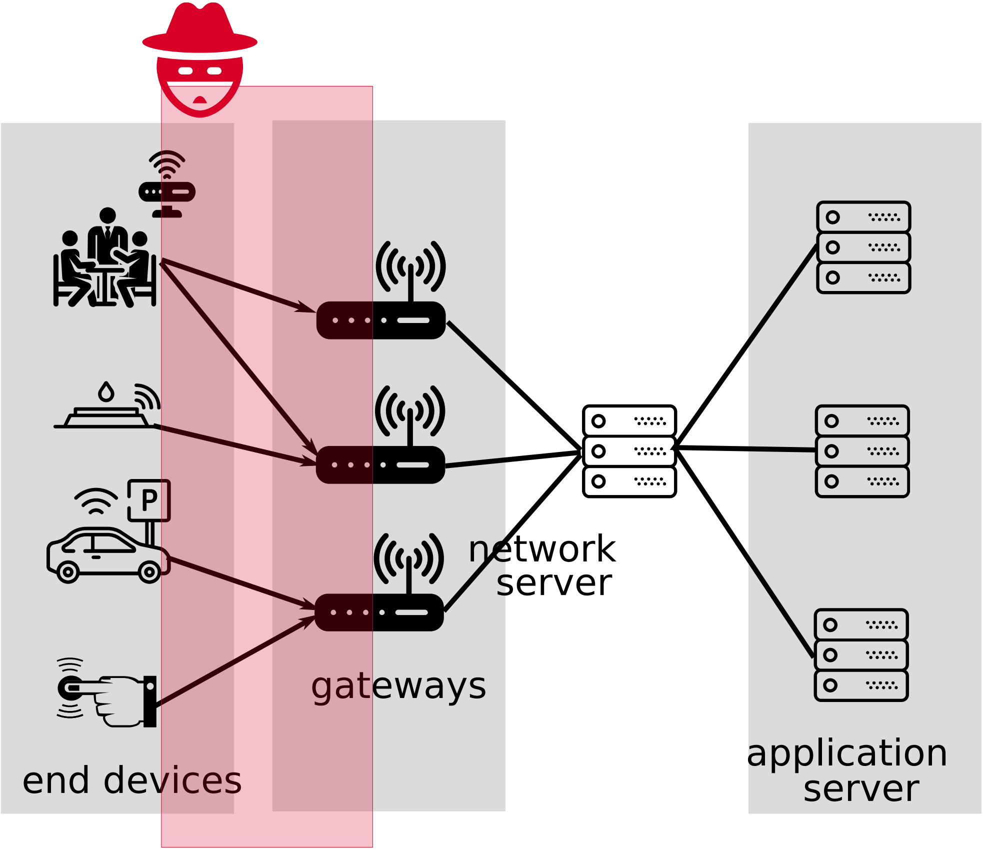

Eavesdropping attacker

- *Attacker's intent*: Obtain sensitive information which is associated with **real-world events** that trigger transmissions
  - Equipment failure, emergency situations, presence/absence of personnel, ...
  - Irrespective of application-level encryption
- *Attacker's approach*: Inspect per-application message timings
  - Can separate applications by frame header, device fingerprinting or based on location

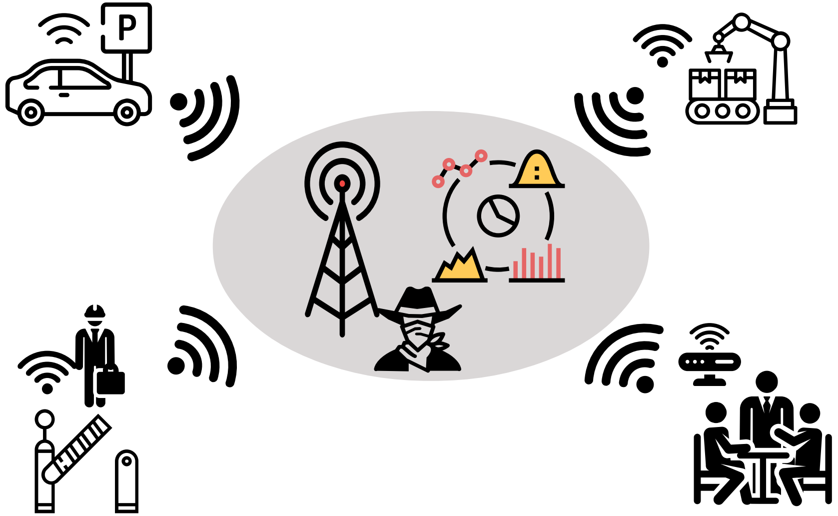

Example: Company parking space
How can real-world events show?

Attacker observable	Real-world event
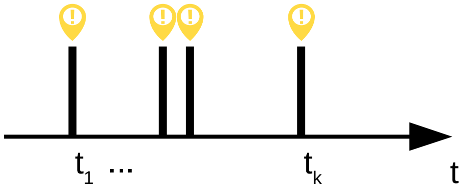	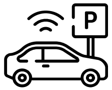 Single car joining or leaving
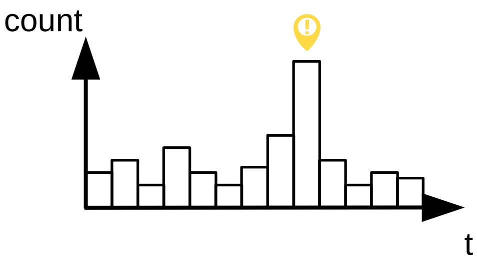	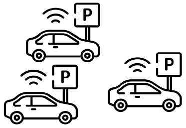 Higher-than-usual number of cars joining or leaving

Two types of information leakage

- **Leakage**: Eavesdropper learns about the occurrence of a **real-world event** by observing message timings and aggregates thereof.
- **Existential Leakage**
  - Transmission of a single message leaks occurrence of an event
- **Statistical Leakage**
  - Statistics of message counts over time leaks information
  - Attacker is interested in observing anomalies. These are likely to represent real-world events.

Example: Company parking space
How can real-world events show?

	Attacker observable	Real-world event
Existential Leakage	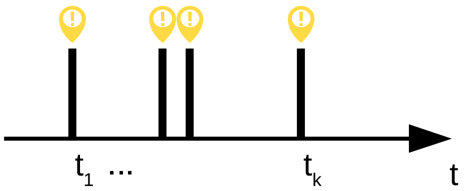	 Single car joining or leaving
Statistical Leakage	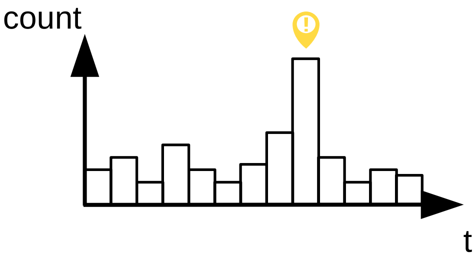	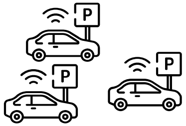 Higher-than-usual number of cars joining or leaving

Can leakage be prevented?

- Assumptions
  - Delay-intolerant messaging
  - In particular: No aggregation
  - Power budget for obfuscation max. identical
- Approach: Dummy messages
- Can we prevent leakage of event information with dummies? At what cost?

Preventing existential leakage

- Messages cannot be removed
- Messages can only be added
- Leakage prevention: Add **dummies** with identical temporal distribution as real packets.
- Each dummy message also represents a fake event ().
- For an anonymity set of size k, increase power by factor of k

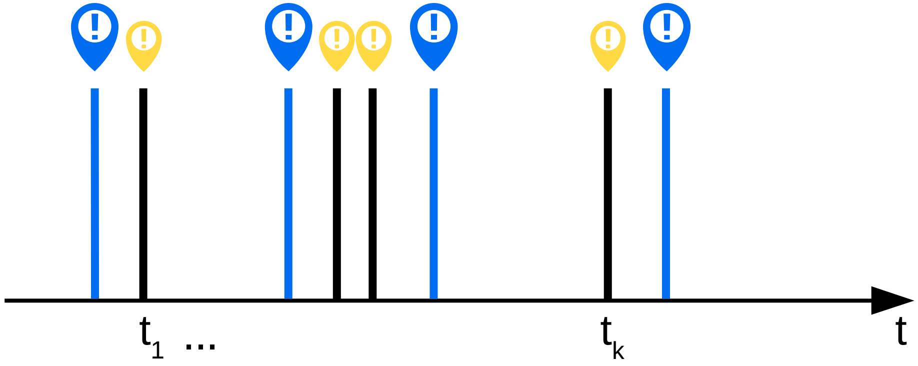

Preventing statistical leakage

Waterfilling

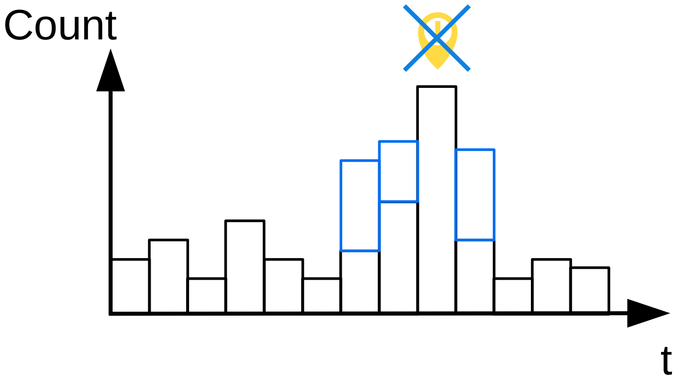

Fake anomaly

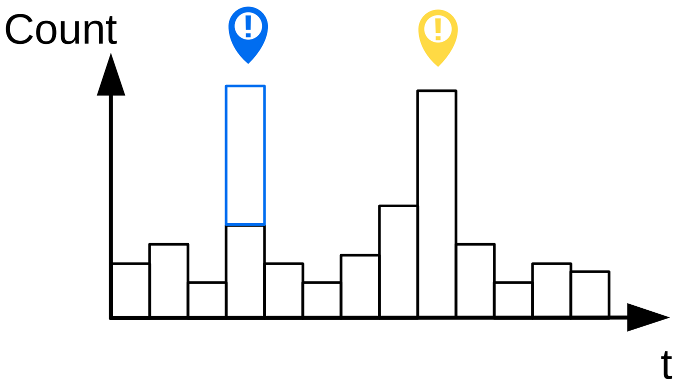

Can we protect from statistical leakage while keeping power consumption within reasonable bounds?

Simulation model: traffic model

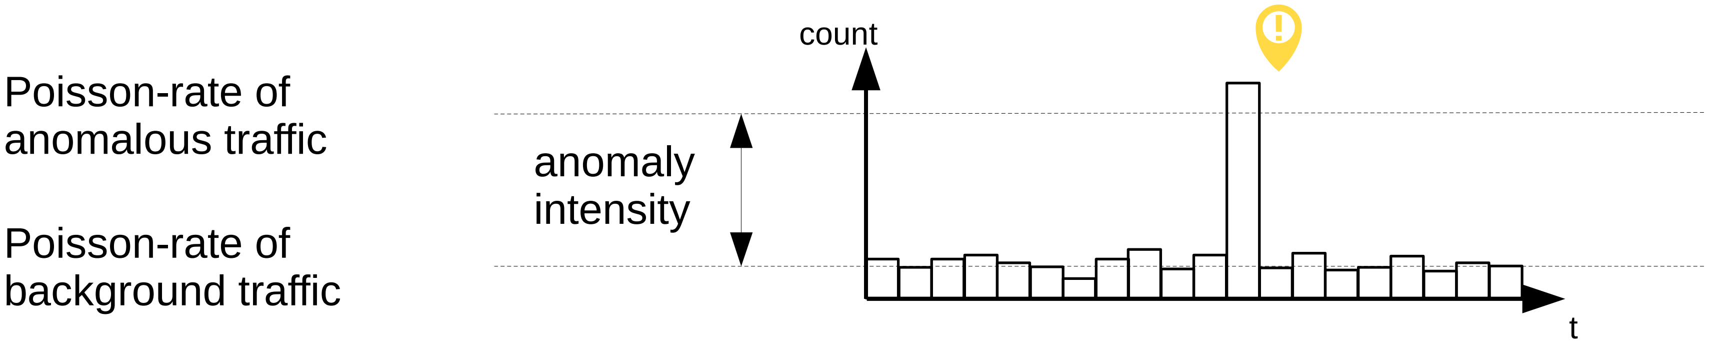

Simulation model: time discretization

Poisson-rate of anomalous traffic

Poisson-rate of background traffic

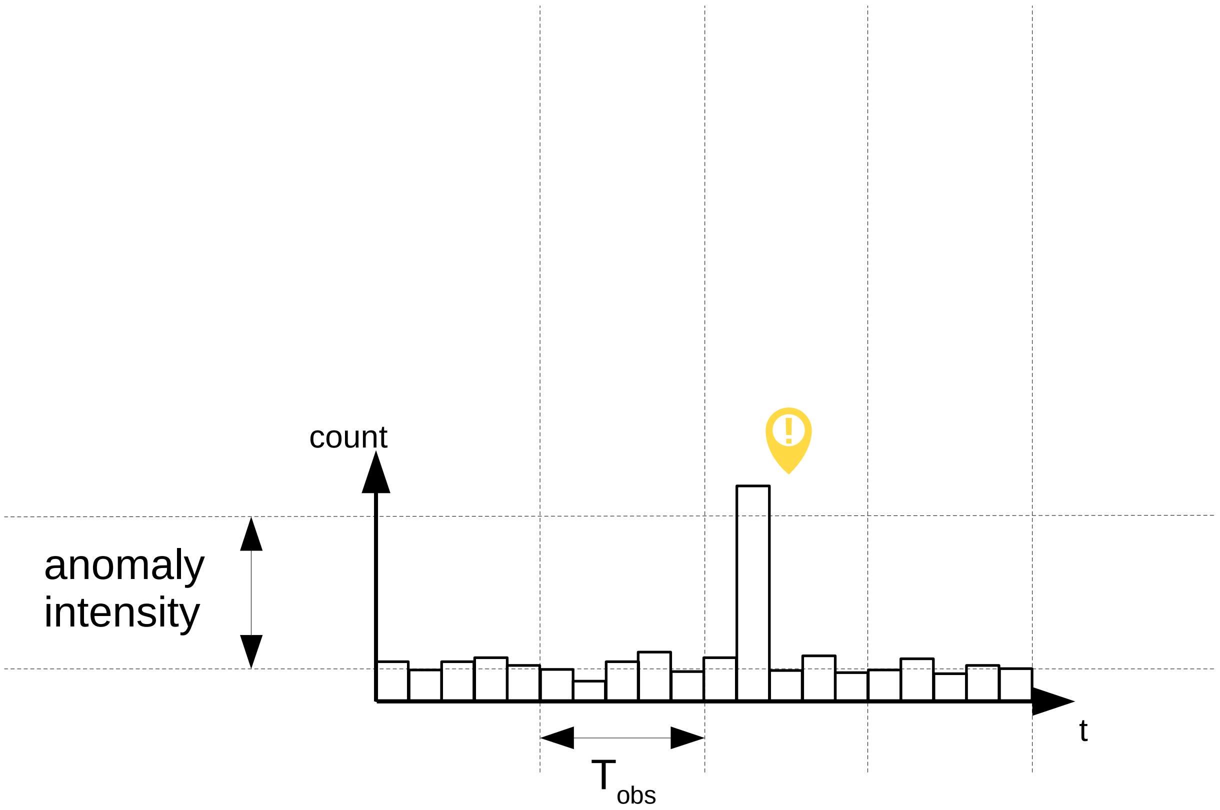

Simulation model: attacker observable

We choose the anomaly indicator as index of dispersion of the count within T_{obs}: $\frac{\sigma^2}{\mu}$

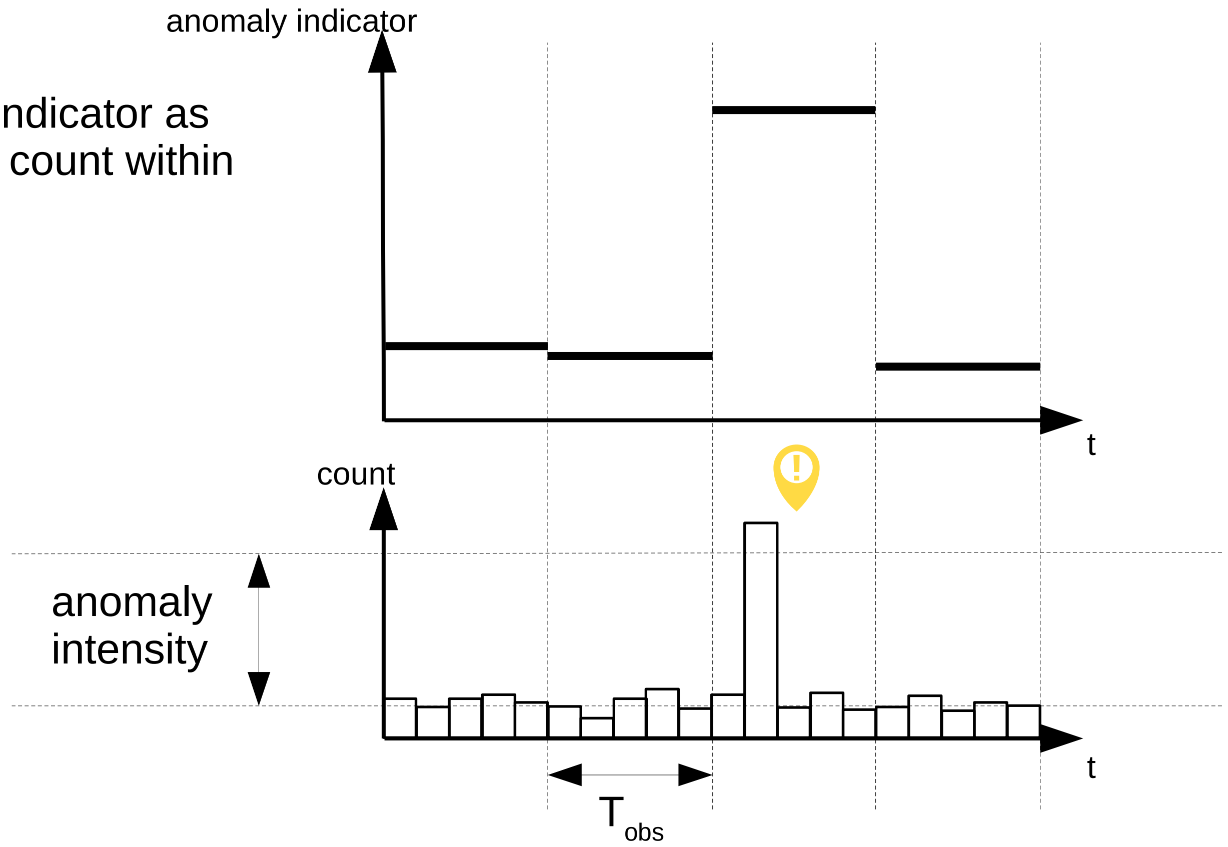

Poisson-rate of anomalous traffic

Poisson-rate of background traffic

Attacker performs binary classification

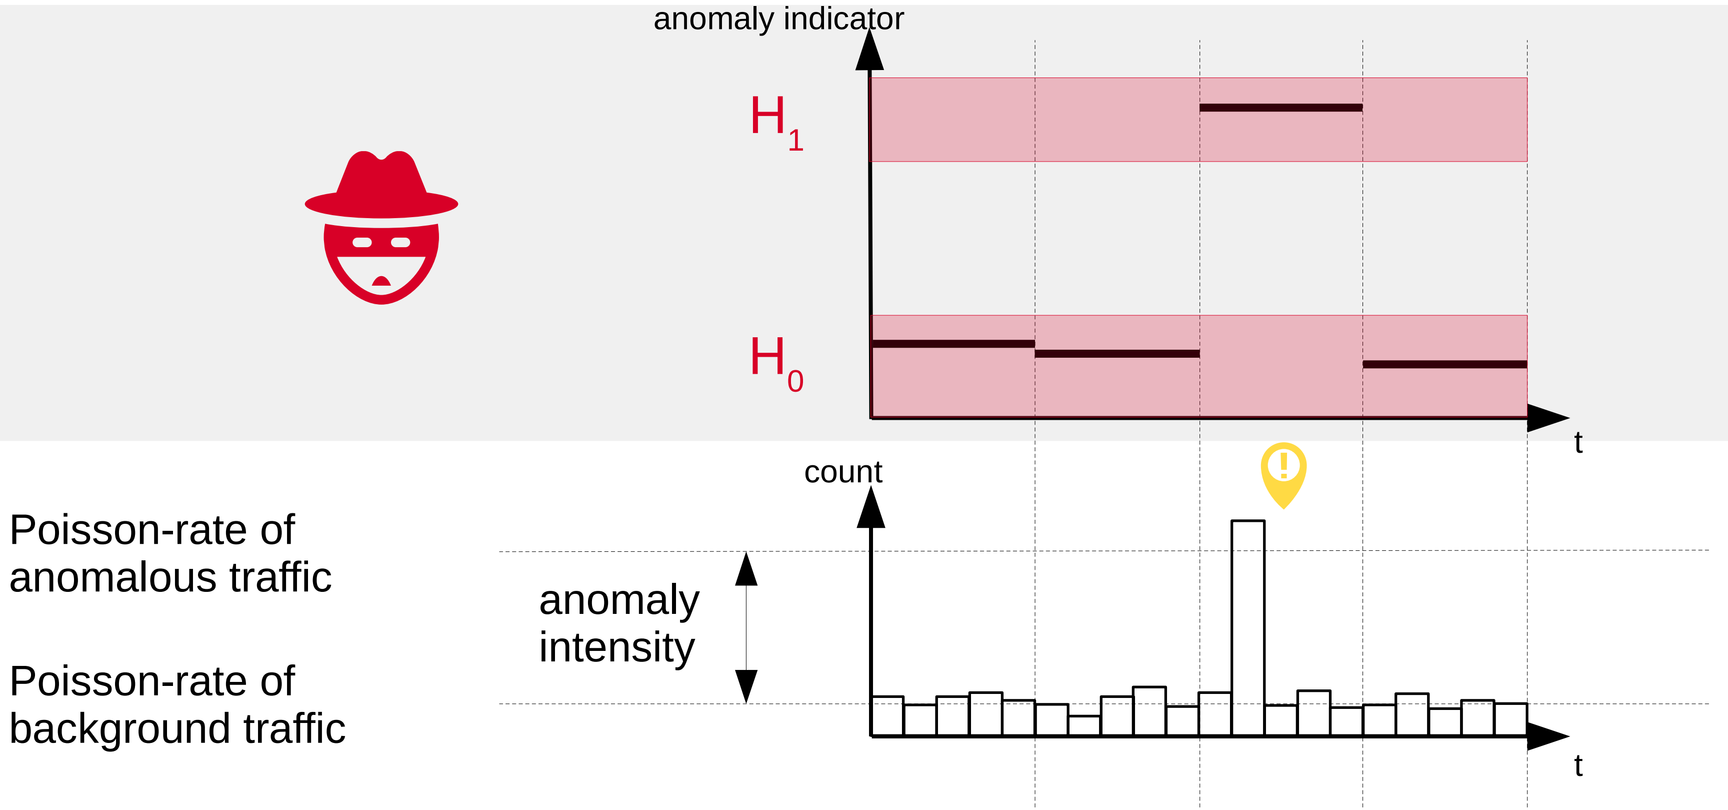

Obfuscation strategy

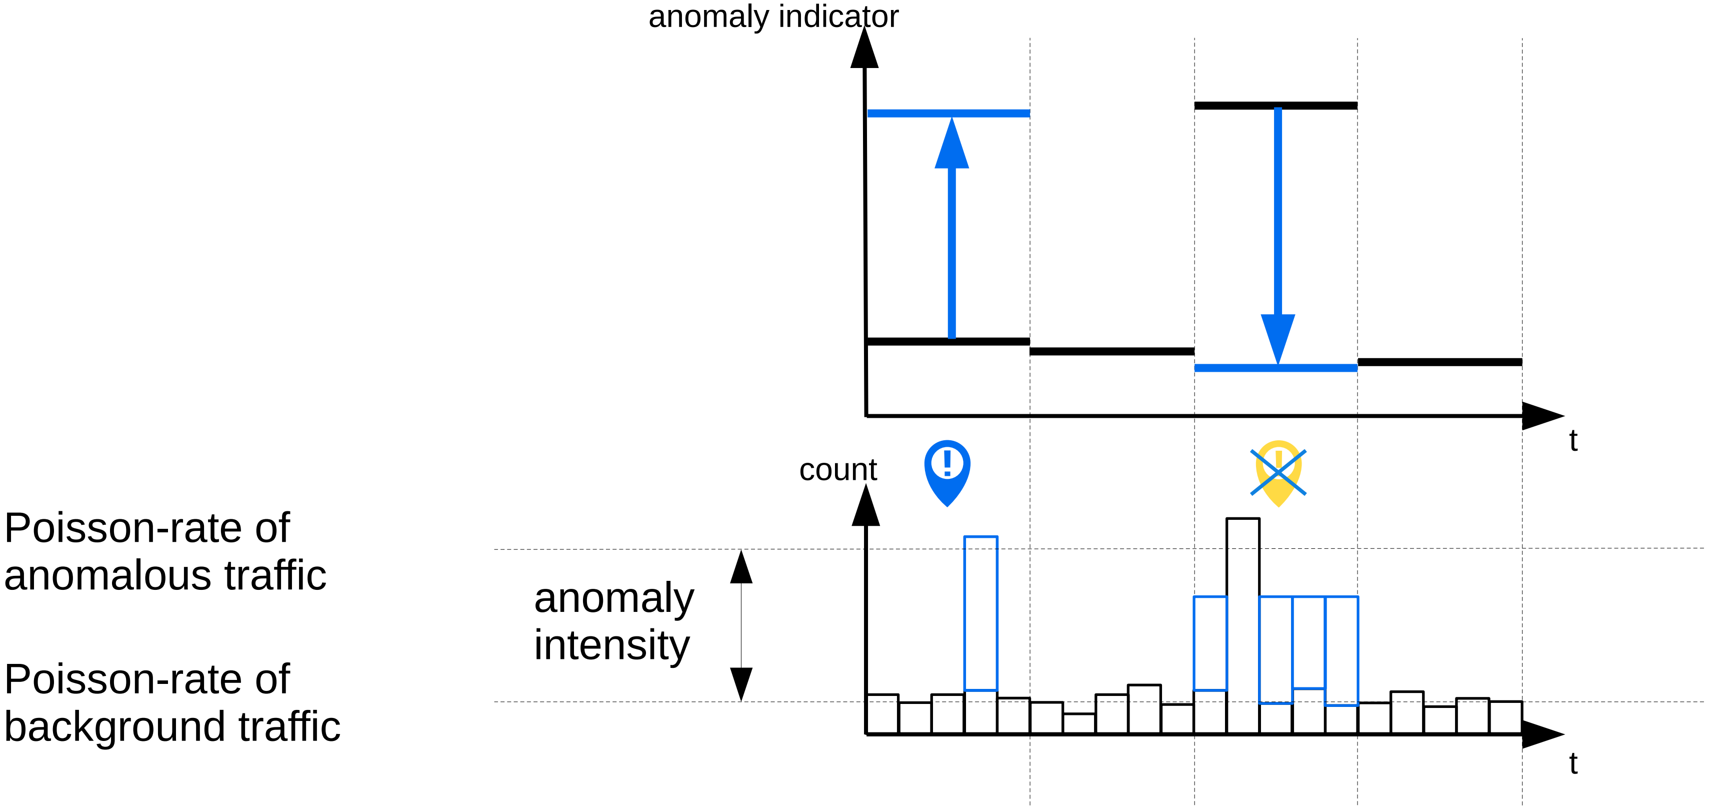

Obfuscation cost depends on strategy

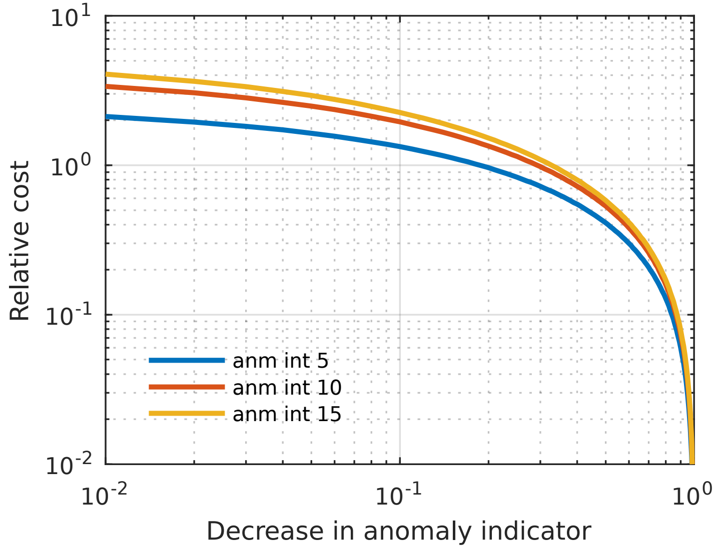

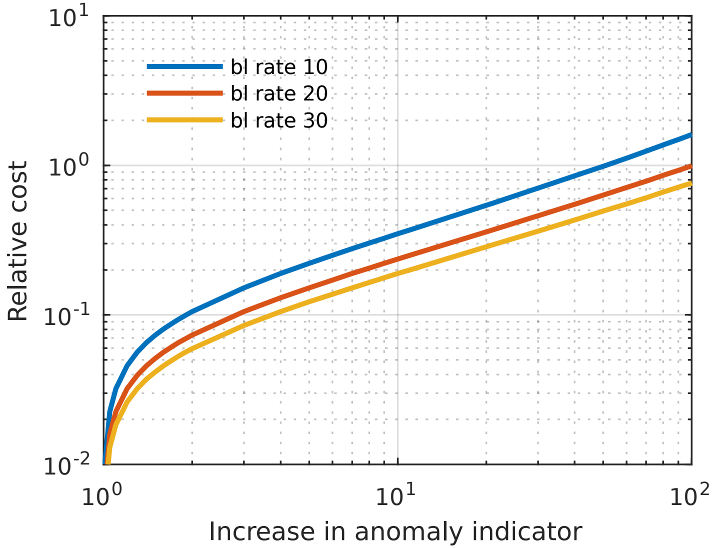

Obfuscation strategy

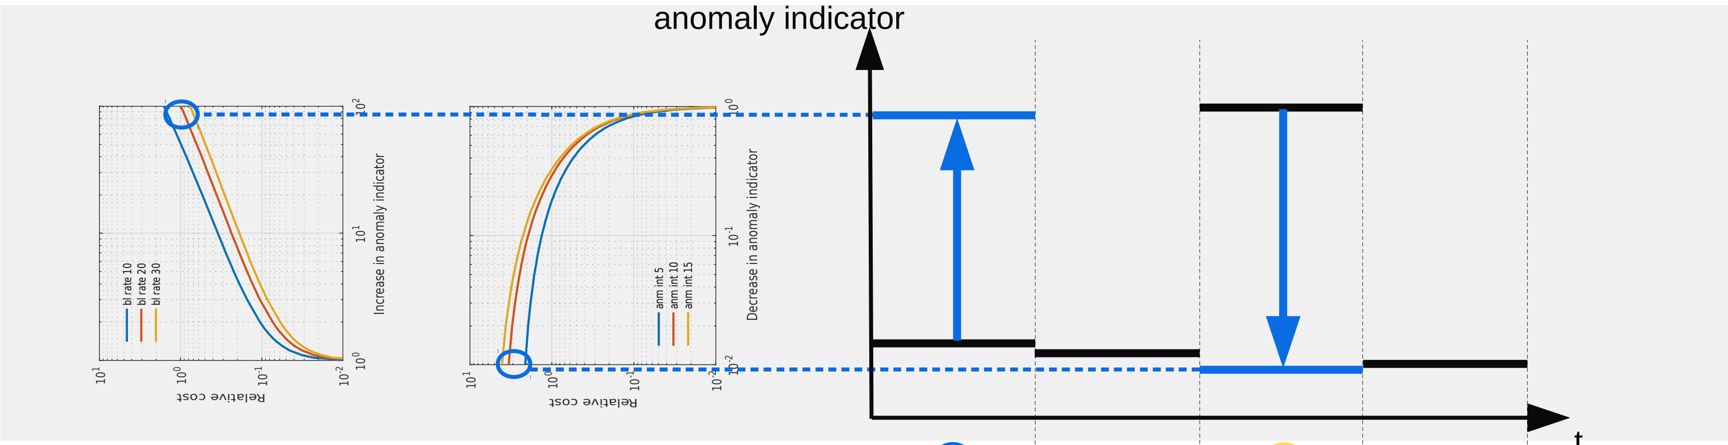

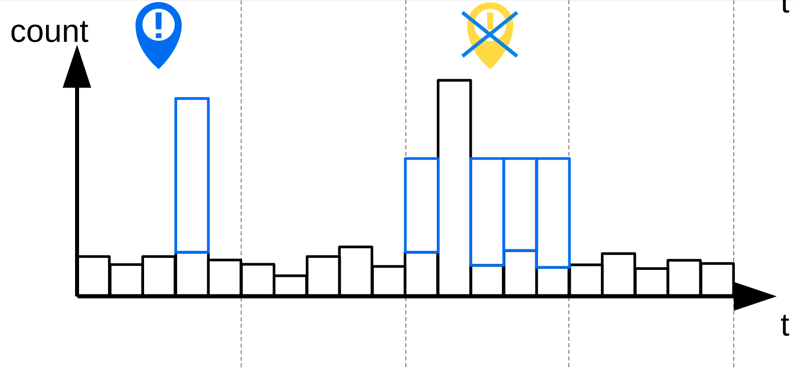

- *Obfuscator's goal*: Optimize obfuscation subject to an average power constraint
- Results in a **strategy**:
  - Probability of waterfilling an anomaly
  - *and* Probability of introducing a fake anomaly

Results

- We consider the performance of a guessing attacker
  - Observes the anomaly indicator per interval
  - Knows the rate of anomalies
  - Knows the obfuscation strategy
  - *Attacker's goal*: correctly assign anomalies to intervals.
- Obfuscation cost limited to the power of real transmissions.
- Average error
  - Which fraction of anomalies was correctly assigned by the guessing attacker?
- Conditional entropy
  - Entropy in the system after the attacker seeing the observable.

Attacker's guessing performance

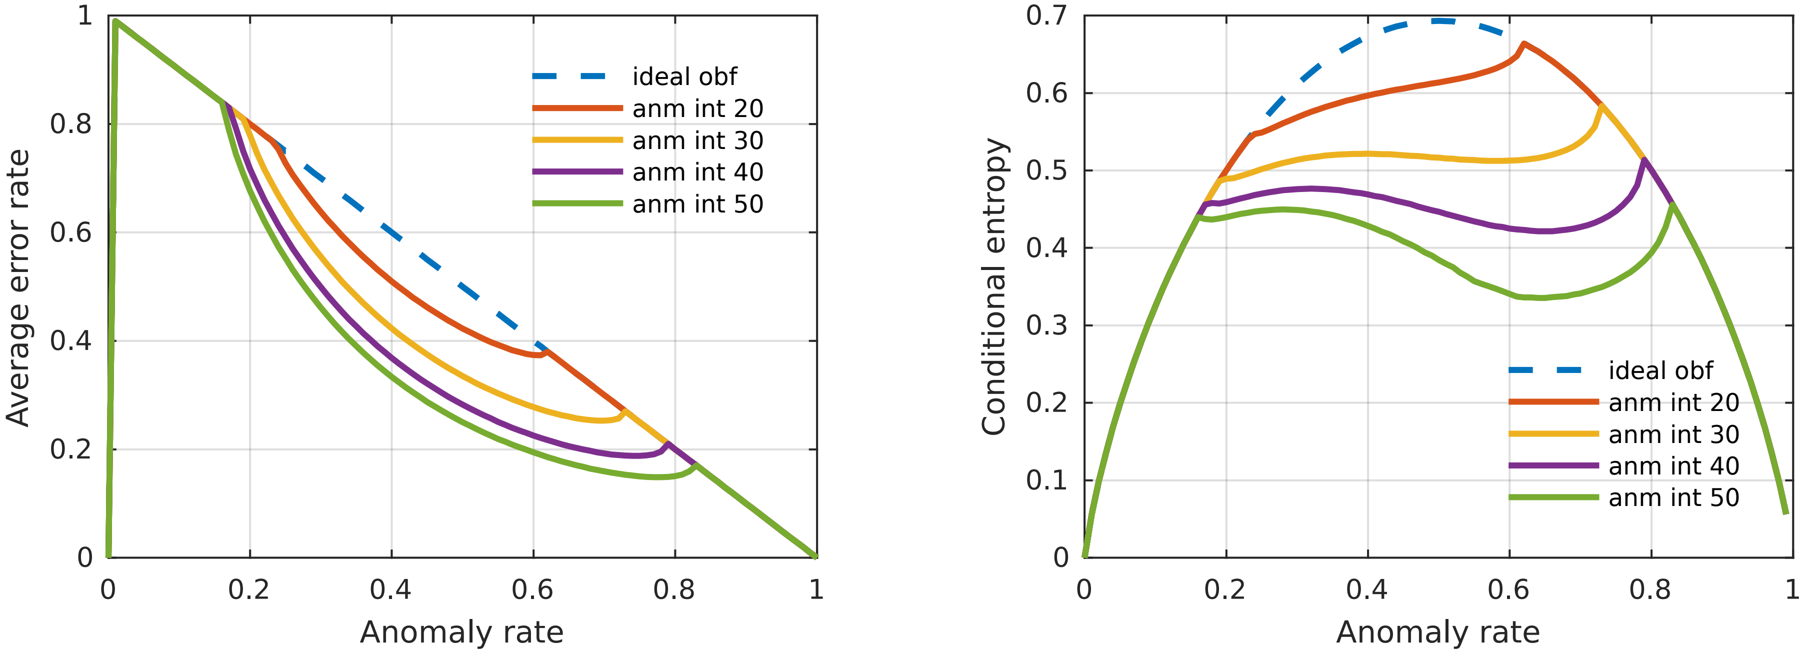

Assumption: Obfuscator has **optimal knowledge** about the occurrence of anomalies.

Attacker’s guessing performance

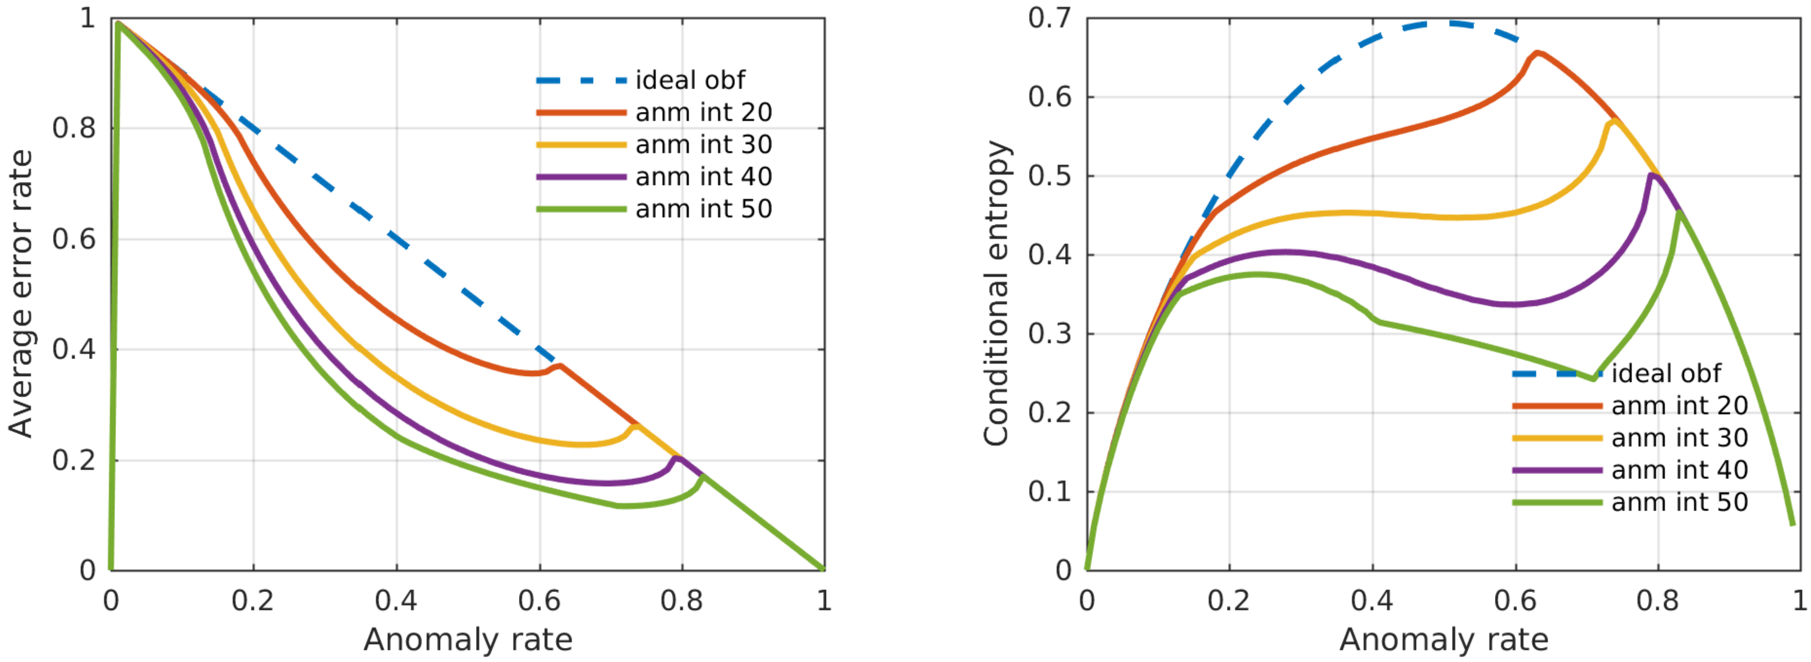

Assumption: Obfuscator has **limited knowledge** about anomaly occurrences (TNR 0.99, TPR 0.7).

Conclusion

- Event-driven communication in LPWANs inherently leaks information.
- The mere existence of messages can leak sensitive information, as do statistical patterns in general.
- Implementation of privacy-enhancing techniques in the LPWAN context hard, as their effect is limited without incurring significant additional energy cost.