Analyzing TETRA Location Privacy and Network Availability

Martin Pfeiffer
Secure Mobile Networking Lab
TU Darmstadt, Germany
mpfeiffer@seemoo.tu-darmstadt.de

Jan-Pascal Kwiotek
Secure Mobile Networking Lab
TU Darmstadt, Germany
jkwiotek@seemoo.tu-darmstadt.de

Jiska Classen
Secure Mobile Networking Lab
TU Darmstadt, Germany
jclassen@seemoo.tu-darmstadt.de

Robin Klose
Secure Mobile Networking Lab
TU Darmstadt, Germany
rklose@seemoo.tu-darmstadt.de

Matthias Hollick
Secure Mobile Networking Lab
TU Darmstadt, Germany
mhollick@seemoo.tu-darmstadt.de

ABSTRACT

Terrestrial Trunked Radio (TETRA) is a digital communication standard taking over from analog communication in various emergency services and governmental agencies in Europe since the late 1990s. TETRA has to meet stringent requirements for a dependable communication infrastructure as it is used in the public safety sector by professional users and first responders. In fact, TETRA is claimed to enhance resilience, security, and availability compared to former analog communication standards. In this paper, we demonstrate that TETRA's location privacy and dependability can easily be undermined and weakened. In particular, TETRA devices can be localized by means of antenna arrays and direction finding techniques on the physical layer. Further, we practically evaluate the impact of various jamming and fuzzing attacks on two commonly used devices. The digital standard even raises new attack vectors, since regular beaconing enables localization of idle devices, and forged frames can provably crash devices.

1. INTRODUCTION

Digital communication systems offer promising features compared to analog systems, including better voice quality, data transmission capabilities, higher spectral efficiency, group calls, encryption, and authentication. Emergency services and governmental agencies recently migrated most communication infrastructure from analog technologies to TETRA or TETRAPOL in European countries, as shown in Figure 1. TETRA provides additional flexibility, as it allows on the one hand a clear separation of institutions such as police, medical services, and firefighters, but on the other hand supports cooperation between them, as required in larger operations. In addition, TETRA increases the network's capacity enabling various simultaneous calls. Security and privacy protection is part of the standard, since it is important to control which information is disclosed to whom: e.g., criminals should not know where the police is, and the press should not know about ongoing operations.

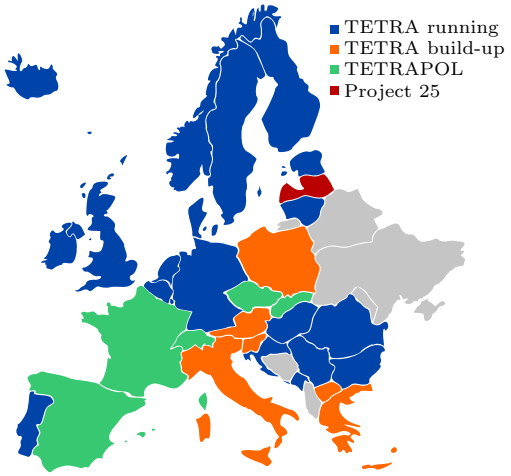

Figure 1: Current state of TETRA rollout, based on [4].

Yet, the transition from analog to digital communication equipment introduces new vulnerabilities that must be addressed. In particular, the location privacy might be broken at the physical layer by estimating the angle-of-arrival (AoA) of radio waves. Localization becomes even more effective in digital modes as a base station (BS) transmits continuously and a mobile station (MS) regularly performs signaling to the BS, even if they are not used actively. Knowledge of BS and MS positions could for instance enable criminals to shut down the network or to gain valuable information about police actions. Moreover, digital modes raise new attack vectors due to higher system complexity, as has been demonstrated for the Global System for Mobile Communications (GSM), where forged text message transmissions were shown to crash recent phones [7].

Our research is limited to TETRA unencrypted direct mode operation (DMO) due to legal issues. Nevertheless,

SPSM'16, October 24 2016, Vienna, Austria

DOI: http://dx.doi.org/10.1145/2994459.2994463

we suspect that other modes are also affected. DMO can be used on amateur radio frequencies, such that experiments do not harm any emergency services. Critical experiments are performed in an electromagnetically isolated tent. We make the following contributions based on Universal Software Radio Peripheral (USRP) N210 and GNURadio:

- We implement a localization attack of TETRA equipment based on AoA estimation of TETRA signals,
- we implement a TETRA fuzzing framework for unencrypted DMO transmissions, and
- we demonstrate vulnerabilities that cause MS devices to crash, reboot or mute.

This paper is structured as follows. Section 2 briefly describes the TETRA network hierarchy and underlying protocol. In Section 3, we demonstrate the localization attack. In Section 4, we perform a fuzzing analysis and discuss its impact. We conclude our work in Section 5.

2. BACKGROUND

TETRA supports an infrastructure mode called trunked mode operation (TMO) requiring a BS, and an ad-hoc mode called direct mode operation (DMO). Devices can operate in both modes simultaneously, and even operate as repeater or gateway. In the following, we compare TMO and DMO [5].

The basic physical layer is very similar for both modes. Device localization and stateless jamming can be implemented on this layer regardless of upper layer features like encryption and authentication.

On the data link layer, TMO requires more management overhead, while DMO requires more synchronization information. In TMO, a BS automatically assigns up- and downlink channels for short time slots. DMO users need to actively select a channel different from the TMO channels; devices coordinate themselves to avoid collisions. The signaling channels (SCHs) are almost equal: both modes support SCH/F for full slots, and the uplink- and downlink half slot SCH/HU and SCH/HD in TMO correspond to the undirected SCH/H in DMO. DMO normal bursts and TMO up- and downlink bursts have similar data formats, but they differ with respect to preambles and guard intervals, and the TMO downlink burst can carry broadcasts.

For the network layer functionality, TMO provides much more modes than DMO. While TMO offers the full voice and data (V+D) service set, DMO only has a subset in the direct mode call control (DMCC). Still, DMCC is very similar to V+D, for example, the short data service (SDS) is available in both modes and equal on the network layer. Additional services offered in TMO could open further attack vectors, which we do not research.

Legal limitations restrict our research to the unencrypted DMO. We perform our experiments based on protocol parts that are very similar to TMO. However, we cannot confirm the impact in real governmental TETRA installations.

3. LOCALIZATION

Locations of MSes and BSes in digital trunked radio networks are sensitive information, as they allow to track users or to gain physical access to network infrastructure. In this section, we show that TETRA equipment can be localized by means of AoA estimation, which effectively breaks location privacy. The localization and tracking of devices in TETRA networks is additionally fostered by the fact that management data is transmitted periodically in addition to user data. While there are different physical layer techniques to localize a transmitter, we particularly use a direction finding technique based on AoA estimation, which is functional without access to the infrastructure and without the need to transmit specific signals. We provide a method for detecting, processing, and tracking TETRA signals to estimate their AoA at an antenna array, and present a practical implementation of an AoA estimator for TETRA signals based on Ettus Research N210 USRPs and GNURadio.

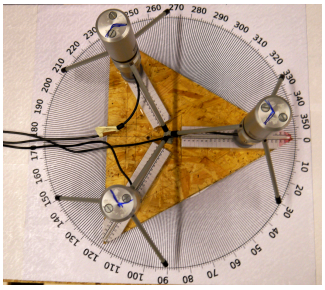

Figure 2: Antenna array.

3.1 Localization Model

AoA estimation techniques of radio signals are based on the fact that a wave front inciding on an antenna array reaches each of the individual antennas at slightly different times. The amount of this time difference effectively depends on the AoA, the antenna spacing, and the propagation speed c of the wave. Further, the relative group delay of the wave front inciding on different antennas translates to linear phase shifts in the frequency domain between the signals received at the different antennas. Hence, the AoA can be estimated by means of measured relative phase shifts of the received signals. There have been various approaches to estimate the AoA of radio signals, many of which are based on MUSIC [8]. An AoA estimation technique for narrow-band signals with analog and digital modulation is presented in [3].

In the following, we describe the AoA estimation technique of our practical implementation, which is partially based on [3]. In order to detect signals from any direction on the horizontal plane around the antenna array, we consider a uniform circular array (UCA): all antennas are placed equally spaced on a circle. In particular, we consider three antennas on an equilateral triangle as shown in Figure 2.

As signal processing takes place in the baseband, we start with a discussion on the effect of a radio frequency (RF) signal's delay on its corresponding baseband signal after downconversion. Equation 1 describes the signal $f_s(t)$ transmitted after upconverting the baseband signal $s(t)$ to carrier frequency f_c:

$$f_s(t) = s(t) \cdot e^{j2\pi f_c t} \tag{1}$$

If the transmitted signal reaches a receive antenna with a delay τ on an ideal channel, the receiver's frontend will see:

$$f_r(t) = f_s(t-\tau) = s(t-\tau) \cdot e^{j2\pi f_c (t-\tau)} \tag{2}$$

After downconversion, the received baseband signal is:

$$r(t) = f_r(t) \cdot e^{-j2\pi f_c t} = s(t-\tau) \cdot e^{-j2\pi f_c \tau} \tag{3}$$

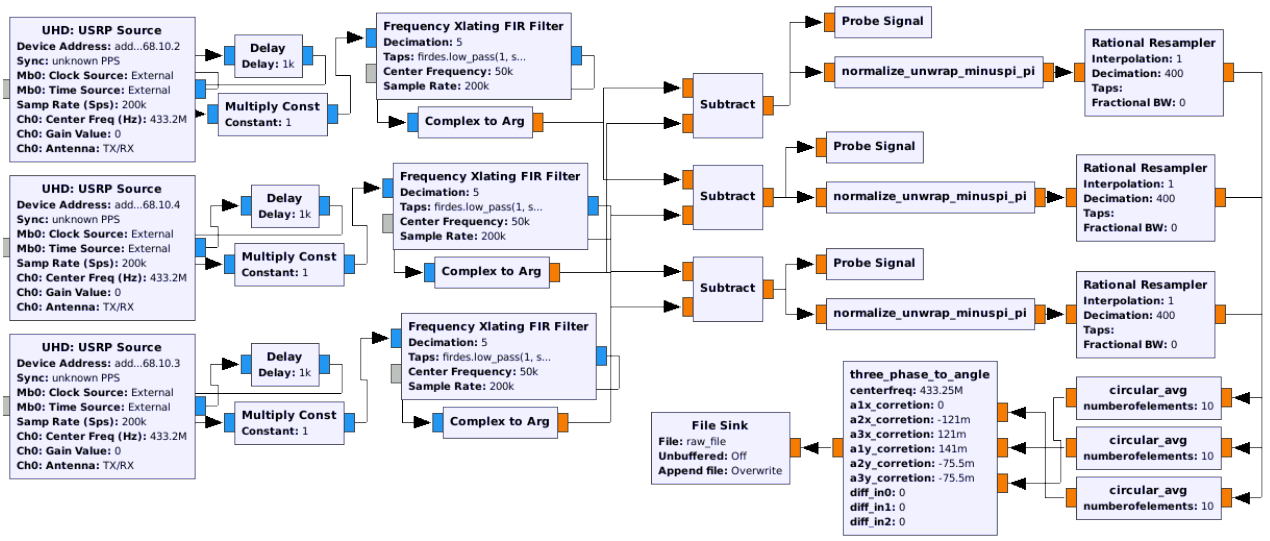

Figure 3: GNURadio localization program.

If a second antenna receives the same signal with a slight additional delay $\triangle$, it will see the following baseband signal:

$$r(t) = s(t - \tau - \triangle) \cdot e^{-j2\pi f_c(\tau+\triangle)} \tag{4}$$

The relative delay $\triangle$ between both versions of the received baseband signal is marginal and might introduce a measurable relative linear phase shift only for very high baseband frequencies, such that it can be neglected for narrow-band signals, especially if f_c is close to the actual signal's radio frequency. More importantly, both versions of the received baseband signal experience a constant relative phase shift $e^{-j2\pi f_c \triangle}$, which can be used to estimate the relative group delay of the wave front of a narrow-band signal from a measured relative phase shift Φ:

$$\hat{\triangle} = \frac{\Phi}{2\pi f_c} \tag{5}$$

By means of the estimated relative group delay $\hat{\triangle}$ and the propagation speed c_{air}, one can estimate the distance $\hat{z}$ travelled by the wave front between the two antennas:

$$\hat{z} = \hat{\triangle} * c_{\text{air}} \tag{6}$$

Note that $\hat{z}$ is bounded by the actual distance between both antennas and may become 0 if the wave front incides on both antennas sideways at the same time. For an antenna array with three antennas A_0, A_1, and A_2 at positions x_{A_0}, x_{A_1}, and x_{A_2}, respectively, one can measure three phase differences and therefore estimate $\hat{z}_{A_{0,1}}$, $\hat{z}_{A_{0,2}}$, and $\hat{z}_{A_{1,2}}$. Let the normalized propagation vector $\vec{S}$ of the wave front be defined in a 2-dimensional space, where x_0 and x_s denote the center of the antenna array and the position of the transmitter, respectively, and $||.||$ is the L2 norm:

$$\vec{S} = \frac{x_0 - x_s}{||x_0 - x_s||} \tag{7}$$

Further, define vectors between antenna pairs:

$$\vec{x_{A_{0,1}}} = x_{A_1} - x_{A_0} \tag{8}$$

$$\vec{x_{A_{0,2}}} = x_{A_2} - x_{A_0} \tag{9}$$

$$\vec{x_{A_{1,2}}} = x_{A_2} - x_{A_1} \tag{10}$$

The signal's propagation vector $\vec{S}$ can be estimated by solving the following linear system or by approximating a solution in the case of practical measurements:

$$\begin{pmatrix} x_{A_{0,1}} \\ x_{A_{0,2}} \\ x_{A_{1,2}} \end{pmatrix} \cdot \vec{S} = \begin{pmatrix} \hat{z}_{A_{0,1}} \\ \hat{z}_{A_{0,2}} \\ \hat{z}_{A_{1,2}} \end{pmatrix} \tag{11}$$

Finally, the AoA γ at the antenna array can be estimated from the components of $\vec{S}$:

$$\hat{\gamma} = \arctan \frac{S_y}{S_x} \tag{12}$$

3.2 Implementation

We use USRP N210 software-defined radios (SDRs) with SBX daughter boards for our implementation due to the following reasons. First, USRPs can easily be operated in the frequency bands of TETRA networks. Second, the available bandwidth is large enough to capture up- and downlinks of different TETRA networks simultaneously, which allows to flexibly detect TETRA transmissions and potentially track the AoAs of multiple devices simultaneously. Third, USRPs allow for clock synchronization, which is required for the phase shift estimation of the carrier wave. We have also tested various antennas to meet the following requirements:

- a homogenous radiation pattern to receive signals from all directions,
- a high gain in the range of 380 MHz up to 480 MHz since TETRA networks are operated in this band, and
- a small form factor since the antennas must be integrated in an antenna array.

Due to their good characteristics, we use SIRIO GP430LP groundplane antennas.

Synchronization

In order to get precise signal phase shift estimates between antennas, all receivers must be synchronized on a scale of nanoseconds and their clocks must be phase-locked to each other. According to an application note by Ettus Research

[6], the synchronization of USRPs N210 SDR platform requires a common 10 MHz reference signal to match the oscillators, a common 1 pps signal to adjust the clocks and the sampling times, and a calibration signal to initialize the phase offset between the USRPs. The setup of our implementation is depicted in Figure 4. The USRP in the bottom is equipped with a global positioning system disciplined oscillator (GPSDO) to synchronize the remaining USRPs. Other than supposed, a single USRP hardware driver (UHD) source does not reliably phase-lock more than two USRPs in our setup. Therefore, we use three separate UHD sources and customized code to synchronize the clocks using the 1 pps signal. Sampling clocks are finally started on all devices at a common point in time.

Calibration

Before being able to estimate phase differences between the received signals, we use a calibration procedure to initialize the relative phases of the different receiver chains. In doing so, we transmit a calibration signal from a fourth USRP to the inputs of the receiver USRPs as shown in Figure 4. When the calibration has finished, the receivers switch to their second inputs, respectively, that are connected to the antennas. We have validated the proper functioning of this procedure in experiments that are out of scope of this paper.

Signal Selection

There is a limited number of frequencies used in TETRA and each of them corresponds to an absolute radio frequency channel number (ARFCN). Our system automatically detects TETRA signals by continuously tracking the received signal strenght indication (RSSI) values at the frequency bins corresponding to ARFCNs and feeds them to our localization program shown in Figure 3.

Experiment Setup

To validate the AoA estimator, we set up a sender transmitting a TETRA compliant signal with 25 kHz bandwith and $\pi/4$ differential quadrature phase shift keying (DQPSK) modulation in the ham radio frequency band. The building blocks for this program are taken from the osmocomTETRA project [2].

Our antenna array consists of three antennas placed on an equilateral triangle that can be flexibly rotated as depicted in Figure 2. The antenna spacing should be equal between all three antenna pairs and be chosen according to the carrier wavelength and the maximum expected estimation error:

$$d_{\max} = \lambda/2 - \epsilon \tag{13}$$

Subtracting the estimation error avoids phase jumps in the estimates at the expense of a little precision. The transmitter's antenna and the antenna array are placed 5 m apart of each other at the same height.

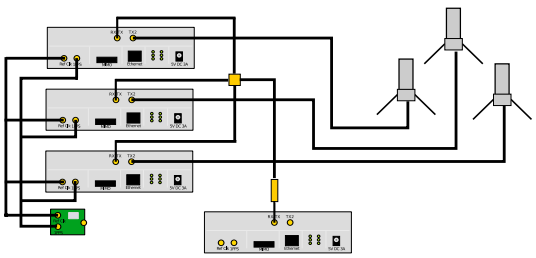

Figure 4: Localization setup.

3.3 Evaluation

The antenna array is rotated in steps of 10°, while in each position at least 500 AoA estimates are taken. The transmitter using ham radio frequencies and DMO is located at approximately 630 m distance. Figure 5 shows the estimated AoA as a function of the actual angle after averaging. The estimator achieves an accuracy of at least ±10°. The slight oscillation of the estimates suggests a systematic symmetric error. In fact, this allows to enhance the estimation accuracy by rotating the antenna array by 360° and averaging the estimated angle while compensating the rotation angle of the antenna array. In doing so, the estimation error is reduced to ±4°.

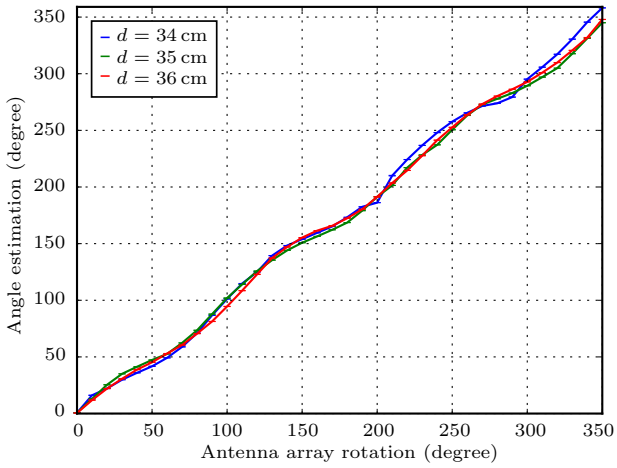

Figure 5: Estimated AoAs in all angular directions with different antenna distances $d \leq d_{\max}$ of the antenna array.

Consequently, transmitting devices of TETRA networks can be localized in case of line-of-sight through the measurement of at least two AoAs at different known positions. Note that the localization error highly depends on the device location relative to the AoA measurement positions: the localization error is lower if the device is closer; and a small intersection area also requires that the AoA estimates are taken from preferably orthogonal directions.

The signaling interval between a MS and its BS depends on the device configuration and is typically smaller than 2 min in idle mode, enabling a location tracking threat. Unlike older analog emergency service communication systems, TETRA requires this regular signaling since MSes need to register at BSes, authenticate, and exchange management data. Hence, localization of stations and devices is a severe threat introduced by TETRA by design.

4. FUZZING

In this section, we describe our fuzzing analysis setup as well as successful jamming and fuzzing attacks.

4.1 Analysis Setup

We implement a fuzzing framework based on osmocomTETRA [2] and Dizzy [1] for the DMO over all layers and especially the SDS text messages within DMCC on the network

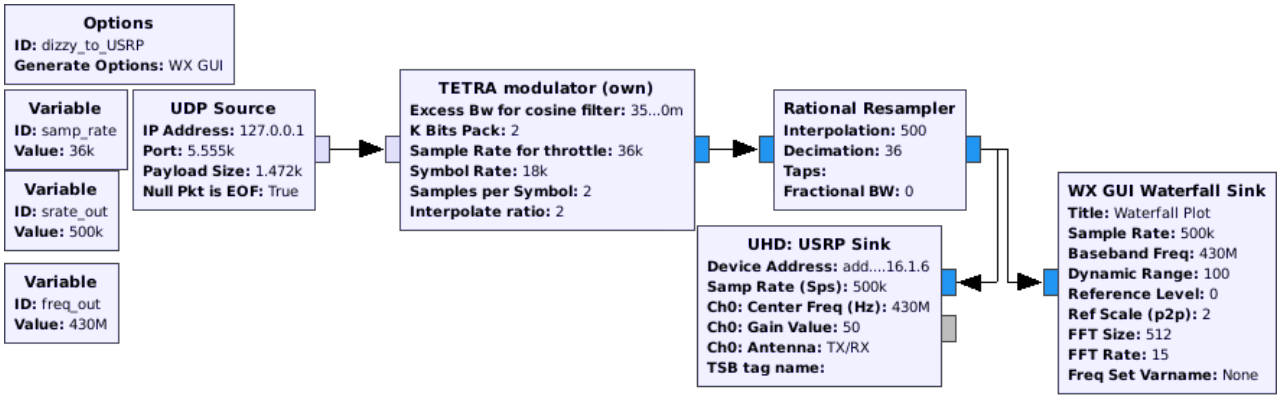

Figure 6: GNURadio DMO transmission program without modulation details.

layer. Our setup is shown in Figure 7. All experiments are performed with real TETRA MSes and the SDR platform USRP N210. Both MSes are from different popular vendors, which are also used in emergency services. We call them device A and device B throughout the paper.

4.2 Implementation

Our TETRA fuzzing framework can replay recorded physical layer signals and generate signals from upper layer TETRA bursts. For fuzzing, we give Dizzy [1] a representation of the network layer DMCC fields and data link layer fields as well as synchronization bursts. A socket connecting Dizzy to GNURadio waits for DMCC and synchronization bursts, and adds the data link layer block encoding and training sequences. The physical layer is implemented in GNURadio based on osmocomTETRA [2]—several modifications are required to add DMO reception and transmission, since only TMO reception was supported. Our DMO transmission GNURadio block is shown in Figure 6. The fuzzing framework is also used for replay attacks but without field processing, to get clean, noise-free signals.

4.3 Stateless Jamming

We implement a straight forward jamming-based Denial of Service (DoS) by replaying a simple synchronization burst 2 times per second. Receivers of the burst are jammed successfully. During the jamming attack, both MSes are neither able to start voice calls nor to send text messages, so they are **temporarily muted**. An error message occurs indicating the channel is busy, which is also shown by continuously flashing light emitting diodes (LEDs). Up to this point, the devices behave correctly.

Figure 7: Fuzzing setup.

Yet, after the jamming attack, when no more synchronization bursts are transmitted, both MSes sometimes are **continuously muted**. We observe that the devices stop the voice functionality and the push-to-talk button shows no action. In addition, the button sounds are disabled. Sometimes, the devices continue working after a few minutes, but often they require a reboot to fix the problem. The continuous muting is more likely to appear the higher the jamming rate and duration are.

4.4 Reactive Jamming

TETRA is vulnerable to reactive jamming by design, because the protocol is designed to have long synchronization periods and the sender's identity is transmitted. Hence, jammers have quite a long time between receiving information and selectively jamming stations. Long burst durations enable jamming even with our comparatively slow python-based software stack. This is in contrast to analog transmissions, where voice is transmitted immediately and the origin can only be estimated with sophisticated physical layer methods such as device on-off-transition characteristics or localization.

In our setup, two regular MSes are communicating to each other while the osmocomTETRA USRP reacts to the destination address of the target. We confirm that our jammer is fast enough by setting the jamming frequency slightly off the original signal's frequency. We record the overlapping signals with another USRP and confirm the temporal alignment of both transmissions in the spectrum. When jamming specific frames, we observe **unreliable communication**, but no crashes or other effects on the device firmware.

4.5 Arbitrary Message Types

We modify the SCH/H part of a synchronization burst. Fields are altered according to the standard, such as increasing the frame counter and slot number. Yet, we mark the message type field for full fuzzing. Our device A reacts similarly as to the stateless jamming and is **selectively and continuously muted**: reception and transmission are no longer possible, and the top LED keeps flashing, but

the menu is still working. To restore functionality, device A needs to reboot. This only affects the addressed device, making this attack more powerful than stateless jamming.

The misbehavior could be caused by one of the following issues. First, the false message type could cause the device to wait for more data that actually never arrives. Second, implementations of other message types might not check boundaries and values that are not expected, which can cause a buffer overflow or endless waiting.

We assume that this also affects devices of type A in TMO, where the message type field also exists. Even though device B did not react to the attack, further devices and even BSes might be vulnerable.

4.6 High Amount of Text Messages

When sending a high amount of text messages using SDS to device A, they do not just fill up the inbox, but after a while the **device reboots**. During the reboot, the device will be unavailable for a short period of time. After the reboot, normal functionality is restored.

We expect that results related to SDS also apply to TMO since they are similar on the network layer. Moreover, the way how messages are stored is technically equal for both modes, since they share one inbox.

4.7 Evaluation

An overview of the attack success is shown in Table 1.

	Device A	Device B
Stateless jamming	muted	muted
Reactive jamming	unreliable	unreliable
Message type	muted	—
Many text messages	reboot	—

Table 1: Successful attacks found by our analysis.

Stateless and reactive jamming are a device unspecific protocol attack—even if the firmware would not go to a continuously muted state, the communication would at least be temporarily muted. Nevertheless, jamming should not continuously mute a device. In a large-scale emergency scenario with many MSes, it would have severe consequences, if devices would crash and continuously be muted, even though the communication channel was actually free again.

The injection of arbitrary message types might also happen randomly due to transmission failures, but it only leads to misbehavior either if an implementation parses an erroneous message type before the checksum or if enough disturbed bits randomly generate a valid checksum, with the latter being very unlikely. Typically, false message types are caused by an attacker and should be filtered out by the firmware. This could have a severe impact and affect further firmware versions, including TMO and BSes.

In contrast, a full message inbox is never a problem for a BS, since it does not store messages. The only exception might be internal message caches, but to fill these, a much higher message transmission rate would be required over a long time, depending on memory and firmware limitations. However, it is a problem for all MSes: they might reject new messages or delete old messages, both causing information loss. The reboot vulnerability we found is even worse, but very device specific.

5. CONCLUSION

Due to their more complicated nature, digital communication systems are prone to be vulnerable. We demonstrate a practical localization and fuzzing framework for TETRA. Since the localization technique is based on the physical layer, it applies to all TETRA systems, independent of communication mode and security features. In contrast, fuzzing vulnerabilities are device specific and firmware updates or TETRA protocol changes could fix them. Future work includes expanding the research to more devices and a wider range of message types.

Emergency services will profit from TETRA security research. The knowledge of the localization problem due to the frequent BS signaling is important for all secret operations and many safety related tasks. During such operations, TETRA devices should better be switched off. In contrast, just not exchanging voice messages would have been sufficient for the old analog devices. Errors that are found by fuzzing or triggered by simple jamming can also appear naturally due to distortions and transmission failures without any attacker present. Fixing these errors in the firmware will make devices more stable and fault-tolerant. We highly encourage device vendors to also perform fuzzing tests on their devices.

Acknowledgments

This work has been funded by the DFG within CROSSING and SFB 1053 MAKI. We also thank Hendrik Schmidt (ERNW GMBH), who co-supervised the work of Jan-Pascal Kwiotek.

6. REFERENCES

[1] Dizzy. https://www.ernw.de/research/dizzy/index.html.

[2] OsmocomTETRA. https://osmocom.org/projects/tetra/wiki/OsmocomTETRA.

[3] L. Balogh and I. Kollar. Angle of Arrival Estimation Based on Interferometer Principle. In *IEEE International Symposium on Intelligent Signal Processing*, pages 219–223, 2003.

[4] Bundesanstalt für den Digitalfunk der Behörden und Organisationen mit Sicherheitsaufgaben. Digitalfunk in Europa. http://www.bdbos.bund.de/DE/Digitalfunk_BOS/Digitalfunk_in_Europa/digitalfunk_in_europa.html.

[5] J. Dunlop, D. Girma, and J. Irvine. *Digital mobile communications and the TETRA system*. John Wiley & Sons, 2013.

[6] Ettus Research. Application Note Synchronization and MIMO Capability with USRP Devices. http://www.ettus.com/content/files/kb/mimo_and_sync_with_usrp_updated.pdf.

[7] C. Mulliner, N. Golde, and J.-P. Seifert. SMS of Death: From Analyzing to Attacking Mobile Phones on a Large Scale. In *USENIX Security Symposium*, 2011.

[8] R. O. Schmidt. Multiple Emitter Location and Signal Parameter Estimation. *IEEE Transactions on Antennas and Propagation*, 34(3):276–280, 1986.