TELECOMMUNICATIONS AND RADIO ENGINEERING

UDC 621.396:004.056(045)
DOI:10.18372/1990-5548.87.20885

Pavlo Chernikov

URBAN LORAWAN RESILIENCE UNDER EW INTERFERENCE: AN OPERATIONAL MODEL FOR MUNICIPAL NETWORKS

Specialized Municipal Enterprise "Kyivteleservis", Kyiv, Ukraine
E-mail: pavlo.chernikov@gmail.com ORCID 0009-0006-7390-9698

Abstract*—This paper reports field experience from a municipal LoRaWAN network that periodically faces electronic-warfare interference. We connect the observed service failures to a simple SINR-based reception model and translate that model into practical actions an operator can take – without expensive instrumentation and without heavy changes on end devices. The result is an operations-oriented playbook: which radio frequency symptoms to watch for in gateway / network-server data, how to recognize likely interference patterns, and what mitigations help keep core city services running under degraded radio conditions.*

Keywords—LoRaWAN; LPWAN; electronic warfare; jamming; interference; smart city; municipal services; operational resilience.

I. INTRODUCTION

Low-Power Wide-Area Networks (LPWAN) are increasingly used for municipal monitoring and control tasks due to their wide coverage and low device power consumption. However, in real deployments the radio layer is exposed to interference that can be accidental (co-channel emitters, industrial noise) or intentional (electronic warfare jamming). This paper consolidates engineering observations and a practical operational model for maintaining LoRaWAN service continuity under such conditions.

LoRaWAN-based low-power wide-area networks (LPWANs) are widely adopted in smart-city telemetry because they offer long range, low device complexity, and multi-year battery lifetime. Municipal deployments support heterogeneous sensing tasks (metering, environmental monitoring, infrastructure status, and safety-related telemetry). Reliability is often measured in routine conditions, yet the most critical requirement appears during infrastructure disruptions: when services degrade, the city depends on telemetry for situational awareness and resource planning.

A specific vulnerability of LPWAN telemetry is susceptibility to electromagnetic interference. In addition to accidental interference sources (industrial emitters, faulty power equipment, or local radio frequency (RF) congestion), intentional interference can occur in contested environments. Under severe interference, packet delivery can drop abruptly, and standard 'increase retransmissions' tactics may backfire by consuming airtime and draining batteries. The practical problem addressed in this paper is how a city-owned LoRaWAN network can sustain acceptable packet reception during intermittent high-power interference while preserving energy efficiency and operational cost constraints.

II. LITERATURE ANALYSIS

LoRa physical-layer properties (chirp spread spectrum, processing gain through spreading factor) have been studied extensively. Gateway diversity and spatial redundancy are recognized as major determinants of reliability in urban deployments. Network mechanisms such as adaptive data rate (ADR), confirmed uplinks, and channel hopping influence airtime and decoding probability. Recent work also investigates intentional jamming and reactive interference against LoRa PHY, including sweep jamming and protocol-aware interference, and proposes countermeasures such as frequency diversity, time diversity, and detection-based response.

However, many published results assume controlled laboratory setups or access to extensive instrumentation (spectrum analyzers, synchronized gateways, or device-side RF sampling). Municipal operators frequently lack the ability to instrument every area and must act with limited visibility. They also face constraints that are rarely modeled explicitly: (i) limited flexibility in spectrum use due to regional regulations and existing channel plans; (ii) budgetary limits on rapid gateway densification;

(iii) service continuity requirements for legacy devices and long replacement cycles; and (iv) strict energy budgets for battery-powered endpoints. These realities motivate an operational model that bridges the RF layer and actionable network-operations steps.

III. Problem Statement

Purpose and objectives of the study. The purpose of this study is to provide an engineering case-study perspective and an operational model for improving LoRaWAN resilience under intermittent EW interference. The objectives are to: (1) describe the generalized architecture and constraints of a municipal LoRaWAN deployment; (2) formalize a reception feasibility condition via signal-to-interference-plus-noise ratio (SINR) to reason about failure modes; (3) propose mitigation actions that require minimal changes to end devices; and (4) outline monitoring and response workflows that support 'degraded mode' operation of municipal services during interference episodes.

IV. Elements of the Methodological Framework

The work follows a practitioner's case-study approach. To avoid disclosure of sensitive operational parameters, the architecture and results are presented in generalized form. The considered deployment follows the canonical LoRaWAN stack: battery-powered end devices transmit uplinks in an ISM band to one or more gateways; gateways forward frames over IP backhaul to a network server; the server performs deduplication, ADR policy, device session management, and routes data to applications. The municipal operator maintains the network server and gateways and integrates telemetry with service workflows as shown in Fig. 1.

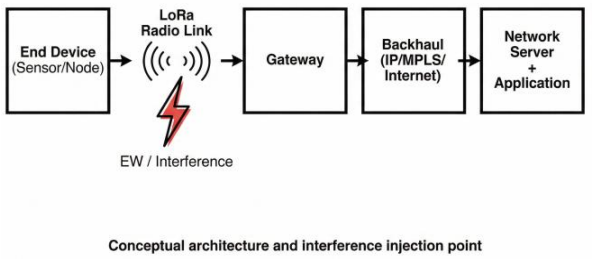

Fig. 1. High-level LoRaWAN architecture and interference injection point (conceptual)

LoRa uses chirp spread spectrum. For a given bandwidth (BW) and spreading factor (SF), the symbol duration is:

$$T_{\text{sym}} \frac{2^{\text{SF}}}{\text{BW}}. \quad (1)$$

Longer symbols (higher SF) increase processing gain and improve robustness to noise and interference at the cost of longer time-on-air. In urban municipal deployments, a common operational tradeoff is between robustness (higher SF) and network capacity (duty-cycle constraints, collision probability, and battery usage). In contested RF conditions, the operator may temporarily favor robustness for a subset of critical devices while keeping the overall network load within regulatory limits.

Reception feasibility under interference can be expressed via the SINR:

$$\text{SINR} = \frac{P_s}{(P_i + N_0)} \geq \gamma_{\min}, \quad (2)$$

where P_s is the received signal power; P_i is aggregate interference power; N_0 is noise power, and $\gamma_{\min}$ is the minimum SINR threshold required for reliable reception (dependent on SF, BW, coding rate, and receiver implementation). While $\gamma_{\min}$ values are typically provided by receiver sensitivity curves, the operational value of (2) is its directional guidance: actions that increase P_s at least one gateway (coverage improvements and spatial diversity), reduce effective P_i (avoid impaired channels, exploit frequency diversity), or lower the required threshold by choosing more robust settings can restore reception.

In practice, municipal operators rarely observe P_s and P_i directly. Instead, they observe proxy variables reported by gateways and the network server: per-gateway RSSI and SNR estimates, the number of gateways receiving each uplink (redundancy), frame counters, and application-level delivery rates. Therefore, we define an operational monitoring vector $M(t)$ that includes: (a) uplink success rate by device class; (b) distribution of gateway-reported SNR; (c) deduplication count per uplink; (d) confirmed-uplink retry counts (when used); and (e) per-channel load / packet error indicators available from the network server. An interference episode is detected when $M(t)$ deviates beyond predefined thresholds for a sustained interval.

Interference modes relevant to urban LoRaWAN can be grouped into three categories. Category A is wideband high-power noise that reduces SINR across many channels, typically resulting in a broad drop of uplink delivery and lower reported SNR values. Category B is partial-band or swept interference that intermittently affects subsets of channels, manifesting as channel-dependent loss and oscillating performance. Category C is localized interference that impacts specific neighborhoods due to geometry and line-of-sight, visible as spatially

clustered loss and decreased gateway diversity for affected devices. Although the RF environment is complex, these categories are sufficient to guide practical mitigations while avoiding disclosure of sensitive parameters, as shown in Fig. 2.

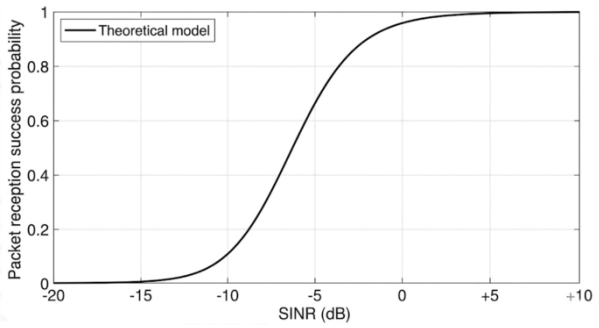

Fig. 2. Illustrative relationship between SINR and packet reception success probability

Operational monitoring focus (gateway logs).

In day-to-day operations, the primary diagnostic source was gateway logging rather than application-layer counters. This choice was practical: gateway logs provide near-real-time visibility into radio-layer conditions (RSSI/SNR, channel, spreading factor, and frame integrity indicators such as cyclic redundancy check (CRC) results) even when upstream services are partially degraded.

We prioritized a small set of gateway-derived indicators that could be tracked continuously and compared against a stable baseline: (i) RSSI and SNR distributions over time, (ii) the share of packets with CRC/validation failures, (iii) per-channel reception counts (to identify frequency-localized loss), and (iv) gateway diversity for repeated uplinks. Tracking these indicators supported rapid differentiation between interference-like signatures and non-radio causes such as backhaul outages or server-side processing delays.

For reporting and post-event analysis, we used aggregated, anonymized summaries (e.g., median / percentiles and ratios to baseline) instead of raw logs. This approach keeps the discussion reproducible for engineering readers while avoiding disclosure of sensitive deployment specifics.

Based on the above monitoring vector and interference categories, we propose an operational model with three layers: detection, classification, and response. Detection uses server – and gateway-level telemetry to flag abnormal degradation. Classification assigns the episode to Category A/B/C based on observable patterns: breadth of impact across channels and geography, changes in gateway diversity, and the stability of SNR distributions. Response selects a mitigation bundle matched to the category and to service priorities.

For Category A (broadband degradation), the primary lever is spatial diversity: increase the chance that at least one gateway receives a decodable frame by improving coverage in critical zones and ensuring multiple gateways have line-of-sight alternatives. Operationally, the response is to switch critical device groups to more robust settings where feasible (e.g., limiting Adaptive Data Rate (ADR) downshifts and allowing higher spreading factors), while reducing non-critical traffic to preserve capacity.

For Category B (partial-band/swept), frequency diversity is central: redistribute traffic across channels, enable or encourage hopping strategies, and identify persistently degraded channels for temporary avoidance. For Category C (localized), the response is targeted: use gateway placement and directional antennas (when applicable) to improve reception in the affected area, and rely on redundancy and buffering for services until RF conditions normalize.

At the service level, municipal applications can implement ‘degraded mode’ logic. Examples include: (i) accepting delayed reporting and increasing buffer sizes for non-critical telemetry; (ii) using multi-sensor redundancy or alternative communications for high-criticality assets; (iii) escalating alerts only when multiple independent indicators fail, to avoid false positives during temporary RF degradation; and (iv) switching service workflows to conservative defaults when telemetry is unavailable (e.g., scheduled manual checks for critical infrastructure). These steps reduce the operational impact of telemetry gaps without requiring immediate RF-layer changes.

Finally, post-episode analysis is essential. The operator should record episode timelines, affected channels, gateway diversity patterns, and recovery behavior. This enables iterative improvement of channel plans, prioritization of gateway upgrades, and validation of detection thresholds. While the current study does not publish sensitive quantitative metrics, the proposed monitoring vector provides a reproducible framework for future measurement campaigns where disclosure is acceptable.

V. Prospective Issues and Practical Implications

Technical context for interference-resilient LoRaWAN operation. To make the case-study discussion transferable, it is useful to formalize the physical-layer trade-offs that dominate LoRaWAN behavior under intentional or unintentional interference. LoRa uses chirp spread spectrum

(CSS), where increasing the spreading factor (SF) increases processing gain and receiver sensitivity, but also increases time-on-air and reduces network capacity. A commonly used approximation of processing gain is:

$$\mathrm{PG} \approx 10 \cdot \log_{10}(2^{\mathrm{SF}}), \quad (3)$$

where PG is the processing gain in dB (illustrative). Higher SF therefore improves robustness to wideband noise and some forms of interference, but it also increases the collision window and duty-cycle consumption for both uplink and downlink.

Noise floor and practical SINR interpretation. In field operation, engineers rarely have direct measurements of interference power P_i at the receiver input. However, gateway-reported RSSI and SNR can be interpreted against the expected thermal noise floor. For a receiver bandwidth BW, a standard engineering estimate of noise power in dBm is:

$$N_0(\mathrm{dBm}) \approx -174 + 10 \cdot \log_{10} \mathrm{BW} + \mathrm{NF}, \quad (4)$$

where 174 dBm/Hz is the thermal noise density at room temperature, BW is in Hz, and NF is the receiver noise figure (in dB). When measured RSSI rises while SNR decreases across many channels and gateways, the most plausible explanation is an elevated effective noise floor (broadband interference) rather than a single narrowband emitter.

Link budget view of "how much interference is too much". A link-budget form is helpful because it separates geometry and hardware from interference. The received power can be expressed as:

$$P_r(\mathrm{dBm}) = P_t(\mathrm{dBm}) + G_t + G_r - L_p - L_m, \quad (5)$$

where P_t is transmit power; G_t and G_r are antenna gains; L_p is path loss; and L_m captures implementation losses (cabling, connectors, polarization mismatch). Substituting P_r into the SINR inequality highlights that there are only a few practical levers during an interference episode: (i) increase effective P_r via antenna placement / diversity, (ii) reduce required $\gamma_{\min}$ by increasing SF or improving coding, or (iii) reduce P_i by frequency planning, filtering, and physical measures.

Network-level capacity under constraints. Even when a single link can be made robust, municipal deployments are limited by network capacity. LoRaWAN uplink access is typically ALOHA-like; with many nodes, collisions become a dominant loss mechanism especially at higher SF (longer packets). A classical throughput approximation for pure ALOHA is:

$$S = G \cdot e^{-2G}, \quad (6)$$

where G is the offered load and S is the successful throughput (both in normalized packet-per-slot units, illustrative). This is not a precise LoRaWAN model, but it explains why "increase SF everywhere" can solve interference for a subset of devices while simultaneously degrading overall delivery due to longer airtime and higher collision probability. For municipal services, resilience therefore often means selective hardening: raising SF only for critical sensors or only during constrained periods, while keeping the rest of the fleet on capacity-efficient settings.

What to monitor during suspected interference.

From a practical engineering perspective, an interference-aware monitoring vector should include: per-packet RSSI and SNR distributions; packet error / CRC failure counts; successful frame counters; gateway diversity (how many gateways receive the same uplink); channel utilization by frequency; and temporal correlation with known grid-power or backhaul disruptions. These observables allow classification without disclosing sensitive deployment parameters.

Interference fingerprints (engineering interpretation). Broadband noise / jamming typically manifests as a simultaneous RSSI increase and SNR decrease across multiple channels and multiple gateways, often with a sharp rise in CRC failures. Narrowband interference may produce a localized degradation confined to specific frequencies; in that case, RSSI may not rise dramatically, but demodulation / CRC failures concentrate on the affected channels. Pulsed or sweeping interference can create bursty loss patterns: periods of normal reception interleaved with short intervals of near-total packet loss. These signatures can be exploited operationally: broadband events favor gateway diversity and SF escalation for critical traffic, while narrowband events favor channel re-mapping and filtering.

Mitigation levers and their trade-offs. Mitigations fall into (a) radio / physical measures and (b) protocol / operations measures. Radio / physical measures include: increasing gateway density or improving placement to increase diversity; using higher-gain or directional antennas for gateways in problematic corridors; improving feeder/cable losses; adding band-pass filtering and better grounding to reduce desensitization; and, where

feasible, separating critical gateways from strong local emitters. Protocol / operations measures include controlled ADR policies, prioritization of critical device classes, downlink budgeting (because acknowledgements and MAC commands also consume airtime), and application-layer buffering to tolerate delayed delivery. The key is to apply mitigations selectively so that the capacity penalty associated with higher SF does not dominate overall network performance, as summarized in Table I.

TABLE I. OBSERVABLE INDICATORS AND LIKELY INTERFERENCE CLASS (CONCEPTUAL)

Observable	Typical pattern	Likely class (hypothesis)
RSSI across channels	Rises across many channels/gateways	Broadband noise / wideband jamming
SNR distribution	Shifts downward; more negative SNR values	Broadband or mixed interference
Loss localized by frequency	Degradation concentrated on specific channels	Narrowband interference / local emitter
Loss over time	Bursty; alternating good/bad intervals	Pulsed or sweeping interference
Gateway diversity	Fewer gateways receive the same uplink	Coverage fragility / elevated interference

A. Illustrative Field Observations (Anonymized)

This subsection provides illustrative, anonymized patterns that engineers can use as a reference when diagnosing interference in LoRaWAN deployments. Values below are presented as typical ranges observed in practice and should be treated as examples (not as audited measurements for a specific site).

- During suspected interference episodes, gateway-reported SNR distributions typically shift downward (e.g., from mildly positive values toward negative values), while RSSI may rise across multiple channels, consistent with an elevated effective noise floor.
- CRC / MIC or frame validation failures commonly increase in bursts, producing short windows of near-zero successful uplinks followed by partial recovery (a pattern consistent with pulsed or sweeping interference).
- Gateway diversity often decreases uplinks that were previously received by multiple gateways may be received by only one (or none), indicating reduced margin and higher sensitivity to interference and fading.
- Loss may be either frequency-localized (narrowband) or cross-channel (wideband). Frequency-localized loss suggests a channel-specific interferer; cross-channel loss suggests broadband noise or multiple emitters.

B. Illustrative Metric Ranges (for Reporting)

If quantitative reporting is needed while keeping deployment details confidential, the following anonymized ranges can be reported as examples and replaced with project-specific values when available.

- Approximate baseline packet delivery ratio (PDR) during stable periods: 0.80–0.90 (operator-reported, aggregated).
- During interference episodes, PDR commonly degraded to 0.20–0.30, with short bursts approaching near-zero delivery in the most affected intervals (approximate).
- Gateway-reported SNR became strongly negative during affected intervals, typically decreasing by several dB relative to baseline (approximate).
- CRC/validation failures increased markedly during bursts (commonly multiple-fold vs baseline), consistent with intermittent wideband interference (illustrative).
- Gateway diversity reduction was often observed concurrently (e.g., from multi-gateway reception to single-gateway reception for the same uplink), indicating reduced link margin (illustrative).

Note: The ranges above are intentionally broad to avoid revealing operational details. For a submission that requires stronger evidence, replace them with your own aggregated statistics (even approximate) and state the data source (e.g., gateway logs / network server counters).

The case-study perspective suggests that LoRaWAN resilience in an urban municipal environment under EW interference is achievable primarily through diversity and operational discipline rather than through a single technical feature. Spatial diversity (multiple gateways), frequency diversity (channel planning and hopping), and policy diversity (service classes and degraded modes) together raise the probability that at least one valid uplink reaches the network server during interference bursts. The SINR condition (2) provides a compact way to reason about why mitigations work: they either increase P_s at some gateway (coverage improvements), reduce effective P_i by moving traffic away from impaired channels, or

improve the system's tolerance by selecting more robust modulation settings.

From an operational perspective, the key is to translate RF degradation into observable server- and gateway-level signals, and to attach those signals to predefined response bundles. This reduces ad-hoc decisions during stressful conditions and helps municipal services continue functioning with explicitly understood limitations.

Future work should include controlled measurements where feasible (e.g., non-sensitive drive tests, gateway-side spectrum snapshots, and device-side RSSI/SNR logging) to quantify the tradeoffs between reliability and energy cost, and to support more formal validation of interference classifiers. Another practical direction is integrating interference detection with incident-management processes so that city services can adapt workflows systematically during telemetry degradation.

VI. Conclusions

Making LoRaWAN resilient under interference is not only a PHY problem. It requires coordination between radio robustness, network configuration, and operational monitoring.

In practice, gateway and network-server data (RSSI/SNR trends, CRC failures, per-channel counts, and gateway diversity) are enough to detect and classify likely interference patterns and to choose a mitigation strategy quickly. The central trade-off is always the same: stronger links (e.g., higher spreading factors) reduce interference sensitivity but consume capacity and airtime. A workable approach for municipal networks is selective hardening – protect the critical services first – and time-bounded adaptation during interference episodes rather than permanent "max robustness everywhere."

References

[1] A. N. De São José, V. Deniau, C. Gransart, T. Vantroys, A. Boé, and E. P. Simon, "Susceptibility of LoRa communications to intentional electromagnetic interference with different sweep periods," *Sensors*, vol. 22, no. 13, Art. no. 5015, 2022, https://doi.org/10.3390/s22135015.

[2] N. Hou, X. Xia, and Y. Zheng, "Jamming of LoRa PHY and countermeasure," *in Proc. IEEE INFOCOM*, 2021, pp. 1–10, https://doi.org/10.1109/INFOCOM42981.2021.9488774.

[3] M. A. Haque and A. Saifullah, "Handling jamming attacks in a LoRa network," *in Proc. ACM/IEEE IoTDI*, 2024, pp. 146–157, https://doi.org/10.1109/IoTDI61053.2024.00017.

[4] A. Dossa and E. M. Amhoud, "Impact of reactive jamming attacks on LoRaWAN: A theoretical and experimental study," arXiv:2501.18339, Jan. 2025, https://doi.org/10.48550/arXiv.2501.18339.

[5] J. Huang and S. Lahoud, "Physical-layer analysis of LoRa robustness in the presence of narrowband interference," arXiv:2512.01088, Nov. 2025, https://doi.org/10.48550/arXiv.2512.01088.

[6] E. Harinda, S. Celentano, A. B. Pons, T. Plets, and L. Martens, "Performance of a live multi-gateway LoRaWAN and interference measurement across indoor and outdoor localities," *Computers*, vol. 11, no. 2, Art. no. 25, Feb. 2022, https://doi.org/10.3390/computers11020025.

[7] R. Sanchez-Iborra, J. Sanchez-Gomez, J. Ballesta-Viñas, M.-D. Cano, and A. F. Skarmeta, "Performance evaluation of LoRa considering scenario conditions," *Sensors*, vol. 18, no. 3, Art. no. 772, 2018, https://doi.org/10.3390/s18030772.

[8] R. Mena, J. Crespo, E. Hermoso, and C. Lloret, "Comprehensive evaluation of LoRaWAN technology in real environments," *Eng. Proc.*, vol. 77, no. 1, Art. no. 28, Nov. 2024, https://doi.org/10.3390/engproc2024077028.

[9] L. E. M. B. Pereira, G. Quiriconi, J. A. Abreu, A. L. Printes, M. Gondres, and J. C. Torné, "Coverage and performance analysis of a private LoRaWAN network deployed in urban areas," *IEEE Latin America Trans.*, vol. 23, no. 11, pp. 1090–1098, Nov. 2025, https://doi.org/10.1109/TLA.2025.11195785.

[10] Netline Technologies, "LoRa protocol and the challenge it raises for counter-IEDs," white paper, 2024.

[11] D. Hambling, "Inside the 'magic radio' protecting Russian drones from jamming," *Forbes*, Dec. 20, 2023. [Online]. Available: https://www.forbes.com/sites/davidhambling/2023/12/20/inside-the-magic-radio-protecting-russian-drones-from-jamming/

[12] B. Remler and J. Segre, "The Curran Papers – No. 3: From factory floor to battlefield: How the Internet of Things democratised precision guidance in Ukraine," *Armada International*, Nov. 10, 2025. [Online]. Available: https://www.armadainternational.com/2025/11/the-curran-papers-no-3-milcom/

Received: December 19, 2025
Accepted: January 22, 2026
Published: February 23, 2026

Chernikov Pavlo. ORCID 0009-0006-7390-9698. Specialist in Law. Master of Public Service.
Head of the Specialized Municipal Enterprise "Kyivteleservis", Kyiv, Ukraine.
Education: Taras Shevchenko National University of Kyiv, Kyiv, Ukraine, (2009); Academy of Municipal Management, Kyiv, Ukraine, (2010).
Research area: cybersecurity of critical infrastructure, urban digital resilience, LoRaWAN networks in wartime environment.
E-mail: pavlo.chernikov@gmail.com

П. О. Черніков. Стійкість міського LoRaWAN до використання РЕБ: операційна модель для муніципальних мереж

У статті наведено практично орієнтоване технічне дослідження міської мережі LoRaWAN, розгорнутої для муніципальних сервісів Києва, та розглянуто її стійкість за правдоподібних сценаріїв радіоелектронної боротьби, зокрема за умов широкосмугових перешкод і реактивного (протокольно-обізнаного) глушіння. Описано чинники розгортання та міграції від GSM-сенсорів, запропоновано модель перешкод, проаналізовано вплив на типові застосування «розумного міста» й узагальнено заходи підвищення надійності, які не погіршують енергоефективність.
Ключові слова: LoRaWAN; LPWAN; радіоелектронна боротьба; глушіння; перешкоди; розумне місто; муніципальні послуги; операційна стійкість.

Черніков Павло Олександрович. ORCID 0009-0006-7390-9698. Спеціаліст права. Магістр державного управління.
Керівник спеціалізованого комунального підприємства «Київтелесервіс», Київ, Україна.
Освіта: Київський національний університет імені Тараса Шевченка, Київ, Україна, (2009). Академія муніципального управління, Київ, Україна, (2010).
Напрям наукової діяльності: кібербезпека критичної інфраструктури, цифрова стійкість міст, мережі LoRaWAN в умовах воєнного часу.
E-mail: pavlo.chernikov@gmail.com