Small Cell, Big Risk: A Security Assessment of 4G LTE Femtocells in the Wild

Yaru Yang*, Yiming Zhang*✉, Tao Wan†‡, Haixin Duan*¶✉,
Deliang Chang§, Yishen Li*, Shujun Tang*§
* Tsinghua University, † CableLabs, ‡ Carleton University, § QI-ANXIN Technology Research Institute,
¶ Quancheng Laboratory

***Abstract*—Femtocells are small, operator-deployed base stations designed to extend mobile network coverage, but their integration into operator mobile infrastructure introduces significant new attack surfaces. While 5G femtocell standards were only recently finalized, 4G LTE femtocells have already been standardized and widely implemented. In this work, we conducted the first systematic security evaluation of 4G LTE femtocells based on both real-world commercial devices and large-scale Internet measurements. We systematically analyzed both the software and hardware of 4 commercial femtocell devices and identified 5 critical and common vulnerabilities that can lead to local or remote compromise. Our Internet-wide measurement identified 86,108 suspected femtocell deployments, many of which are exposed to remote attack. Further, we experimentally validated in a real operator network that a single compromised femtocell can serve as a powerful entry point for attacks on both the mobile core network and its subscribers. Our findings highlight that femtocell security in operational 4G LTE networks remains an urgent concern. We reported our results to Global System for Mobile Communications Association (GSMA) and the 3rd Generation Partnership Project (3GPP) Service and System Aspects Working Group 3 (SA3). 3GPP SA3 has subsequently approved both a study item to further enhance the security of 5G femtocells and a work item to define the Security Assurance Specification (SCAS) for 5G femtocells.**

I. Introduction

Femtocells are small, low-power cellular base stations originally designed to improve indoor mobile coverage. Due to their low cost, ease of deployment, and traffic offloading capabilities, they have seen widespread adoption in homes, enterprises, rural areas, and dedicated networks. According to recent estimates [1], the femtocell market reached a size of $6.49 billion in 2024. As such, femtocells have become an essential component of global cellular networks.

Despite their benefits, femtocells introduce new security risks stemming from their physical accessibility. Unlike traditional base stations, which are typically deployed within operators' secured perimeters, femtocells are often installed in homes or offices where attackers can more easily access the hardware. Golde et al. [2] first investigated the security threats of 3G femtocells in 2012. They demonstrated that a femtocell could be compromised via multiple vulnerabilities and converted into a significantly more powerful "rogue" base station, enabling attacks against both subscribers (e.g., voice call interception) and the mobile network.

With the global deployment of 4G LTE (hereafter 4G) and 5G, 3G infrastructure has been gradually phased out. It raises a natural question: has the security of femtocells in the current 4G networks improved? Despite having significantly greater capabilities, femtocells have received far less research effort than rogue base stations [3], [4], [5], [6]. A compromised femtocell is inherently trusted by the core network, thus can be used to monitor and modify nearly all user traffic, and even to attack the core network. To date, previous work on 4G femtocells [7], [8] focus primarily on local attacks and require physical access. There has been no systematic research on the broader security risks of 4G femtocells, particularly their remote attack surface.

In this work, we aim to fill the gap by answering the following research questions: (Q1) Are there security vulnerabilities common among 4G femtocells that allow for local or remote compromising? (Q2) Do security threats against users and network infrastructure from compromised femtocells, as described in [2] for 3G, still persist in the evolved 4G networks? (Q3) How many femtocells deployed by operators are accessible from the public Internet, thus posing real threats to users and operators? Note that, we focus on 4G femtocells since the initial 5G femtocell specifications were only recently finalized in the 3rd Generation Partnership Project (3GPP) Release 19, and commercial deployment has not yet begun. Moreover, the 3GPP security requirements for 5G femtocells [9] largely inherit those of 4G [10]. Therefore, our findings are expected to generalize to upcoming 5G femtocell deployments.

To answer Q1, we conducted a systematic security assessment of commercial 4G femtocells. Through hardware and software analysis, we identified five key vulnerabilities, including accessible debug interfaces and exposed management services, which could allow local or remote attackers to compromise affected devices. We empirically tested those vulnerabilities on 4 commercially available femtocell devices, and found that all devices exhibited at least 2 vulnerabilities, with some affected by all 5 (Table I in Section IV). Furthermore, we discovered a total of 20 Original Equipment Manufacturers (OEM) vendors represented across the tested devices,

Network and Distributed System Security (NDSS) Symposium 2026
23 - 27 February 2026 , San Diego, CA, USA
ISBN 979-8-9919276-8-0
https://dx.doi.org/10.14722/ndss.2026.241968
www.ndss-symposium.org

suggesting broader impact across the femtocell ecosystem. To answer Q2, we revisited the threat compromised femtocells pose to both 4G end users and core networks. In controlled environments, we demonstrated user-side attacks on mobile data, voice, and SMS services by exploiting vulnerabilities in femtocells. Although testing live core networks was not feasible due to ethical constraints, we analyzed how critical interfaces such as GPRS Tunnelling Protocol User Plane (GTP-U) and S1 Application Protocol (S1AP) could be abused to target core components. Our findings show that key issues first identified in 3G femtocells [2] remain unresolved in 4G, and that existing protections such as Closed Subscriber Group (CSG) and even IPsec can be bypassed or hijacked. To answer Q3, we performed the first large-scale identification of Internet-exposed femtocells based on protocol-level characteristics, including Internet Key Exchange version 2 (IKEv2), TR-069 (Customer-Premises Equipment WAN Management Protocol, CWMP), and web-based administrative interfaces. Leveraging these features, we detected exposed femtocells and assigned confidence levels based on response patterns. Our measurement in the global IPv4 space uncovered 86,108 suspected femtocell devices, including 1,598 with open management interfaces and 52,768 (61.28%) classified as *highly confident*. Notably, 185 devices fully matched one of our testbed models even with identical SSH host keys. We confirmed that this model exposes recoverable credentials, allowing unauthenticated remote access via SSH. This highlights that not only are femtocells widely exposed, but known-vulnerable models are actively deployed in real-world networks. Based on our findings, we offer targeted mitigation recommendations and are actively disclosing the vulnerabilities to affected vendors and organizations, including 3GPP and Global System for Mobile Communications Association (GSMA).

Our key contributions include:

- We conducted a systematic analysis of femtocells and identified 5 common vulnerabilities that can lead to compromise. We tested 4 commercial femtocells and found that each device exhibited multiple vulnerabilities.
- We demonstrated threats from compromised femtocells in controlled environments, including a practical IPsec Man-In-The-Middle (MITM) attack that requires no root access and enables potential attacks on core networks.
- We present the first large-scale measurement of femtocells exposed on the Internet, identifying 86,108 femtocell candidates and revealing critical security risks, including exposed management interfaces.
- We shared some results with 3GPP SA3, which subsequently approved a study item to further enhance the security of femtocells in 5G [11], and a work item to define SCAS for 5G femtocells [12].

The rest of the paper is organized as follows. Section II provides background information on femtocell and related work. Section III outlines the threat model and common vulnerabilities of our interest. Section IV presents security testing results of femtocell devices. Security implications against subscribers and core networks are discussed in Section V. Section VI reports Internet-scale measurements of femtocell exposure, and Section VII outlines mitigation strategies. We discuss disclosure, limitations, and other aspects of our study in Section VIII. We conclude in Section IX, and address ethical considerations in Section X.

II. Background

A. Femtocell Architecture in 4G Networks

To improve indoor coverage, mobile operators have adopted femtocells, which are compact and low-cost base stations typically comparable in size to a home router. 3GPP TS 22.220 [13] (published in 2008) defines femtocells under the terms Home Node B (HNB) for 3G and Home evolved Node B (HeNB) for 4G. 5G femtocells are specified in 3GPP Release 19 as part of the Work Item *5G NR Femto* [14]. At the time of writing, these specifications have only recently been finalized, so 5G femtocells have not yet been widely deployed. This paper therefore focuses on 4G femtocells, which are already standardized and widely deployed in the real world.

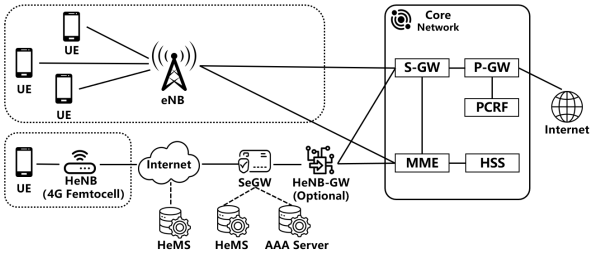

Fig. 1: 4G architecture including eNB and HeNB access.

Figure 1 illustrates a 4G network where User Equipment (UE) connects via either an evolved Node B (eNB) or a Home evolved Node B (HeNB), i.e., a 4G femtocell. The eNB provides wide-area coverage and connects to the core network over dedicated links, while the HeNB targets indoor use and typically connects via the public Internet, requiring additional security measures. As shown in the figure, the HeNB and the Security Gateway (SeGW) perform mutual certificate-based authentication and typically establish IPsec Encapsulating Security Payload (ESP) tunnels to ensure the confidentiality and integrity of the backhaul link. In some deployments, the SeGW also interfaces with an Authentication, Authorization, and Accounting (AAA) server to authenticate the HeNB's hosting party via Extensible Authentication Protocol Method for 3rd Generation Authentication and Key Agreement (EAP-AKA) [15]. This setup requires the HeNB to include an independent Hosting Party Module (HPM) for authentication, provided in the form of a Universal Integrated Circuit Card (UICC). Once authenticated, the HeNB connects to the core network either via an optional Home eNodeB Gateway (HeNB-GW) or, in its absence, directly through the SeGW. Both the eNB and the HeNB interact with the Mobility Management Entity (MME) over the S1-MME interface using the S1 Application Protocol (S1AP) [16], and with the Serving

Gateway (S-GW) over the S1-U interface using the GPRS Tunnelling Protocol User Plane (GTP-U) [17]. In addition, the HeNB is remotely configured, managed, and monitored via the HeNB Management System (HeMS), which may be deployed either behind the SeGW or directly on the public Internet.

B. Related Work

Radio Access Network (RAN) security has been widely studied due to the openness of the air interface, which enables eavesdropping or jamming by external attackers. Base station security is the topic most closely aligned with our work.

Base station security. Early cellular networks (e.g., 2G/GSM) relied on one-way authentication from the base station to the user, exposing users to fake base station attacks. Adversaries can exploit fake base stations to send spoofed SMS messages [18], [19], [4] and harvest user IMSIs for follow-on attacks such as user tracking [20], [21]. Although 4G introduced two-way authentication, attackers can exploit backward compatibility to downgrade users to 2G and continue launching fake base station attacks [22]. Prior work has focused primarily on detecting fake base stations, using features like anomalous signaling data interaction [23] and signal strength [24], [25], [26]. Certificate and digital signature based approaches [27], [5], [28], [29], [30] have been proposed for defense, but they require changes to specifications, limiting their near-term deployability. In contrast, the security of base station devices has received limited attention, as such devices are often assumed to be physically secure. Schmidt et al. [7] challenged this view by purchasing outdated Base Transceiver Station (BTS) modules online and discovering issues such as weak passwords and misconfigurations.

Other notable RAN studies include attacks on identity and location privacy [31], [32], [33], [34], [35], [36], low-layer threats [37], [38], [39], [40], and other protocol-level vulnerabilities [41], [42]. Related work has also examined core network security [43], [44], [45], [46], [47], [48], [49], [50], [51], [52], [53].

Femtocell security. Unlike traditional base stations, the smaller size and user-proximate deployment of femtocells make them more vulnerable to physical access. Despite this, femtocell security has received limited research effort. Early work [54] on H(e)NB security primarily focused on theoretical analysis and lacked experimental validation. Some industrial conference papers [55], [8] evaluated real femtocell devices or tested in real operator networks. But these works typically emphasized findings without offering systematic analysis, methodologies and technical details. Borgaonkar et al. [56] analyzed how to compromise a 3G femtocell device while it remains unclear whether similar issues affect 4G devices and how severe are their real-world impact. Golde et al. [2] discussed potential vulnerabilities in femtocells during updates and the implications of a compromised device on UEs and the core network. Still, the work lacks experimental evidence and remains largely descriptive. Janzen et al. [57] analyzed the security of a 5G indoor O-RAN base station and identified several software-level vulnerabilities. As femtocells differ from O-RAN small cells, whether they share similar vulnerabilities remains unclear. Given their broader deployment in the real-world, femtocells likely pose a greater security risk in practice.

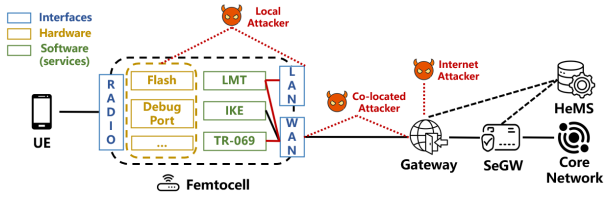

Fig. 2: Threat model of compromising femtocells.

III. Security Analysis

A. Femtocell Internal Structure

To understand the internal structure of femtocells for subsequent security analysis, we examined both the relevant technical specifications [13], [10] and real-world commercial devices. This is because vendors may introduce additional functionality not covered by the specifications to facilitate deployment or management. Therefore, we obtained four commercial femtocell devices from different vendors through legitimate means (e.g., second-hand markets), anonymized as Femto-I (FT-I; similarly hereafter), FT-II, FT-III and FT-IV. We conducted software service scanning (i.e., port scanning and protocol identification) to identify exposed services, and disassembled the devices to examine their hardware components and physical interfaces. Based on our analysis, we constructed the internal structure of a femtocell, as illustrated in the femtocell portion of Figure 2, where components are categorized as follows: interfaces are shown in blue boxes, hardware components in yellow, and software services in green. Note that, compared to Figure 1, it omits the HeNB-GW and simplifies the representation of the core network, while expanding on the internal structure of the femtocell.

Software Services. Femtocells expose various software-level services for communication and management. For local management, they typically enable services such as SSH, Telnet, and web-based management, collectively referred to as the Local Management Terminal (LMT) in Figure 2. To securely connect with the operator's core network, femtocells use IKE to establish IPsec tunnels with the SeGW. Since IKE is a peer-to-peer protocol, femtocells also run a corresponding IKE service to handle incoming handshake messages. For remote management, femtocells initiate HTTP-based TR-069 sessions to the Auto Configuration Server (ACS), which is part of the HeMS, and expose a TR-069 *Connection Request* interface to allow ACS-triggered sessions.

Hardware Components. Femtocells are typically built with hardware architectures similar to those of common Internet of Things (IoT) devices. As such, they include standard components such as a system-on-chip (SoC), flash memory, and RAM. Interestingly, we also identified hardware debug ports in femtocells we obtained.

Interfaces. Femtocells typically expose three types of physical interfaces. The Local Area Network (LAN) interface is used

for local management, while the Wide Area Network (WAN) interface connects the femtocell to the operator's core network. Additionally, the radio interface provides cellular access, allowing UE to connect to the femtocell via the air interface.

B. Threat Model

Compared to prior work, we construct a more comprehensive threat model that, in addition to local attacks, also considers remote compromise of femtocell devices, as shown in Figure 2. The attacker aims to exploit vulnerabilities in hardware design or software services to gain control over the device. Attackers are categorized based on their network position and capabilities. A *local attacker* has physical access to the femtocell, enabling interaction with internal hardware components as well as the LAN interface. A *co-located attacker* resides on the same local network segment (e.g., behind the same gateway) and can access the femtocell's WAN interface without Network Address Translation (NAT) or firewall constraints. In some scenarios, we further assume that the co-located attacker can perform on-path attacks (e.g., by ARP or DHCP spoofing) between the femtocell and the gateway. An *Internet attacker*, in contrast, is located on the public Internet and can only discover and interact with exposed WAN interfaces via Internet scanning (see Section VI for details). In this paper, we use the term *local compromise* to refer to compromise by a local attacker, and *remote compromise* to refer to compromise by either a co-located attacker or an Internet attacker. Once the femtocell is compromised, the attacker can launch further attacks against both the core network and subscribers.

C. Vulnerability Analysis

To identify common security issues that could enable local or remote compromise of femtocells, we conducted a detailed analysis of femtocell devices from both hardware and software levels. At the software level, we focused on analyzing the services exposed by the femtocell over its LAN or WAN interfaces, as these represent the main vectors through which software-level attacks can be carried out.

Hardware level. Our analysis focuses on the debug interfaces and flash memory components, as these are frequently exploited in practical attacks against IoT devices [58], [59]. Common debug interfaces include the Universal Asynchronous Receiver-Transmitter (UART), which allows serial communication with the device, and the Joint Test Action Group (JTAG) interface, which enables low-level hardware debugging and control. Secure practices for these debug interfaces typically require that UART interfaces be disabled or protected by strong authentication credentials, while JTAG interfaces should be permanently disabled (e.g., using eFuse) in commercial, production-grade devices. As a result, we identified two common and easily exploitable vulnerabilities at the hardware level: (V1) accessible debug interfaces and (V2) credential extraction via flash memory dumping. While other hardware-level threats such as fault injection or firmware reflashing may also enable compromise [60], [61], [62], these approaches are either highly device-specific or require significant effort. Firmware reflashing, for instance, typically requires bypassing signature verification, a process that is often highly device-specific and technically demanding. Therefore, we do not focus on these threats in our analysis.

TABLE I: Summary of identified vulnerabilities.

Vulnerability	Device			
	FT-I	FT-II	FT-III	FT-IV
V1. Accessible Debug Interfaces	✔		✔	
V2. Credential Extraction	✔	✔	✔	✔
V3. Credentials Shared across Devices	✔	✔	✔	✔
V4. Exposure of Management Services	✔	✔	✔	
V5. TR-069 Authentication Weakness	✔	✔	—	—

— denotes not applicable (device lacks valid operator provisioning).

Software level. In particular, credentials that are predictable by remote attackers, such as hardcoded values or those that can be computed or guessed from device-specific information, may allow an adversary to compromise multiple femtocells once the credentials of a single device are exposed. Therefore, we analyze (V3) predictable credentials. In addition, since these services are intended for local management, exposing them on the WAN interface can further enable remote compromise. Consequently, we analyze (V4) the exposure of these management services to the WAN interface. In contrast, other threats such as web service implementation flaws are generally device-specific, thus are excluded from our analysis. For the IKE service, since femtocells generally act as IKE initiators and reject inbound IKE connection requests, we believe that this does not pose a security issue that could lead to femtocell compromise. For the TR-069 service, since the ACS may be deployed on the public Internet and is therefore not protected by IPsec, authentication becomes particularly important in this context. Therefore, we analyze (V5) potential authentication weaknesses in TR-069 deployments.

In summary, different types of attackers can exploit various vulnerabilities in femtocell devices. A local attacker with physical access can target vulnerabilities V1 (debug interfaces) and V2 (credential extraction). A co-located or Internet attacker may exploit vulnerabilities V3 and V4 to gain access to exposed management services. Additionally, a co-located attacker can leverage vulnerability V5 to intercept or spoof TR-069 connections.

IV. Vulnerability Testing of Femtocell Devices

In this section, we detail our analysis of five potential vulnerabilities based on testing four commercial femtocell devices. The overall results are presented in Table I.

A. Experimental Settings

We conducted our experiments on 4 commercial femtocell devices to validate the 5 vulnerabilities identified through our analysis. For the hardware analysis, we used soldering tools (e.g. a soldering iron and a heat gun) to attach headers or desolder flash chips, multimeters and oscilloscopes to locate debug pins, and an open source hardware tool *JTAGulator* [63] to identify potential JTAG interfaces. Specifically,

we inspected the Printed Circuit Board (PCB) for silkscreen labels suggesting UART or JTAG signals (e.g., "TX", "RX", "TMS"), as well as unlabeled pin rows or headers (e.g., 4-pin rows). We then used multimeters, oscilloscopes and the *JTAGulator* to confirm the pin functions. After that, we connected these debug interfaces using a USB-to-UART converter or a JTAG debugger (e.g., *J-Link* [64]) to verify their activity. In addition to interface probing, we extracted the contents of the flash memory using a universal programmer *RT-809H* [65] to obtain the firmware image. We then applied tools such as *binwalk* [66] and *cpio* [67] to unpack the firmware file system, and manually identify and extract embedded credentials. We also searched for indicative strings (e.g., `passwd` for modifying Linux user passwords), and reverse-engineered the corresponding binaries to identify hardcoded file paths that store credentials or password-generation logic used to derive them. For hashed credentials, we applied offline cracking techniques, including dictionary-based or brute-force techniques, using tools like *hashcat*[68] and *CMD5*[69]. For the software analysis, we first analyzed the extracted credentials to determine whether they were shared across devices. To assess exposed services and evaluate potential weaknesses in TR-069 authentication, we then connected the femtocell's WAN interface and a laptop simultaneously to the LAN interface of a router running *OpenWrt* [70]. From the laptop, we conducted port scanning on the femtocell, while the *OpenWrt*-based router was used to capture traffic and host a fake TR-069 server for testing.

B. Accessible Debug Interfaces (V1)

Physical debug interfaces like UART and JTAG are widely retained in embedded IoT devices for development and manufacturing purposes. If left active in production systems without adequate protections, these interfaces can provide attackers with low-level access to bootloaders, operating systems, or firmware content. Given that femtocells share similar hardware design features with embedded IoT systems, we systematically examined all acquired devices for residual debug interfaces.

Results. On the FT-I, we identified an accessible JTAG interface (12-pin) and established a connection using a J-Link debugger (Cortex-A7 configuration), enabling low-level CPU access (e.g., registers, memory). On the FT-II, we identified an accessible UART interface (6-pin, 3.3V, 115200 baud) that, when connected via USB-to-UART adapter, output bootloader and kernel logs. Although it lacked an interactive shell, the logs revealed sensitive information such as init scripts and file system structure. The FT-III and FT-IV exhibited the same debug interface behavior as FT-I and FT-II, respectively.

Security impact. Our results confirm that multiple commercial femtocells expose accessible debug interfaces. Each device examined includes at least one functional UART or JTAG interface. UART interfaces enable serial communication and, in our tests, revealed console output such as initialization logs and file system paths. These outputs expose details about the femtocell's software components and directory structure, which may facilitate further attacks such as path-traversal access to sensitive files. While not observed in our setup, unsecured UARTs may even permit unauthenticated operating system access [8]. JTAG interfaces pose a greater risk by allowing low-level access to memory and processor state, potentially enabling full device compromise including root access. Crucially, these vulnerabilities originate from production-stage hardware configurations and cannot be fully addressed via firmware updates alone.

TABLE II: Summary of V2. Identical upper-right label indicate that the corresponding credentials share the same password. A red checkmark in the AKA column indicates that all IMSI, OPc, and Ki values are stored in plaintext or are recoverable.

Device	LMT			AKA
	SSH	Web	Telnet	
FT-I	root, admin[¶] anonymous[†]	admin[¶] user	-	✔
FT-II	root[§]	user[*], admin	Omu...[‡]	✔
FT-III	root, admin, anonymous[†]	admin	-	✔
FT-IV	root[§]	femtodebug, admin	Omu...[‡]	✔

* indicates an anonymized vendor-specific username.
■ Static and recoverable credentials; ■ Dynamic but derivable credentials (e.g., via reverse-engineering password generation logic)

C. Credential Extraction via Flash Memory Dumping (V2)

In this paper, we focus on two categories of credentials that can be exposed in onboard flash memory: (1) Local management credentials, such as those used for SSH access, which may allow attackers to gain administrative control over the femtocell. It is important to note that operators typically do not provide local management credentials to device owners (e.g., residential users), or offer only accounts with minimal privileges (e.g., configuring the WAN interface IP address). However, a malicious device owner with physical access can extract the full file system and either recover stored credentials or reverse-engineer programs that generate them. (2) AKA credentials, including International Mobile Subscriber Identity (IMSI), Derived Operator Code (OPc), and Authentication Key (Ki). Since these AKA credentials represent the identity of the femtocell's hosting party (i.e., the entity responsible for its deployment and management) and are used to authenticate it to the mobile core network, their leakage could enable unauthorized access to the core network.

Results. Table II summarizes extracted credentials, listing usernames (e.g., `root`) and color-coded password extractability. "Recoverable" means the plaintext password can be obtained from hashes or encrypted values; "Derivable" means computability from device-specific inputs (e.g., serial numbers). For ethical reasons, we omit plaintext passwords and exact storage paths and anonymize usernames that would reveal the vendor identity. Across all tested femtocells, we extracted 17 LMT credentials. Although the FT-I's SSH `root` and `admin` passwords were not directly recoverable, we derived them by reversing a password-generation binary. Its web password was identical to the SSH `admin` password and thus also derivable. The FT-III's SSH `root` password could be derived in a similar manner. Furthermore, credential reuse was common: FT-I and FT-III shared the SSH `anonymous`

account, and FT-III's SSH `root` password matched legacy values in FT-I's binary, suggesting firmware inheritance. FT-IV also reused FT-II's SSH and Telnet credentials. Besides, all devices stored recoverable AKA credentials.

Additional Findings. We found the Web login API of FT-I accepts hashed passwords directly, allowing an adversary to authenticate even when the credential is unrecoverable and underivable. We also found other evidence of firmware reuse. For example, FT-I's firmware bundles UI assets from 12 other OEM vendors, including the vendor of FT-III. FT-II's firmware includes configurations for 6 OEM variants, containing 10 additional recoverable Web account credentials.

Security impact. These findings confirm that credential storage in femtocell firmware remains insufficiently protected in practice, with most credentials stored in plaintext, recoverable, or derivable. With dumping tools and moderate effort, attackers can extract sensitive secrets. More critically, most of these credentials are reused across devices or can be easily derivable by remote attackers (Section IV-D) and can be exploited remotely due to the unintended exposure of local management services (Section IV-E). Additionally, we identified a total of 20 OEM vendors (including the devices we tested), which could amplify the scale of security issues [71].

D. Predictable Credentials (V3)

Credential security for local management services relies heavily on unpredictability. *Predictable* means the credentials are not randomly generated, e.g., with certain patterns that can be learned. To assess this risk, we evaluated whether the credentials extracted in Section IV-C are predictable.

Results. All the LMT credentials discussed in Section IV-C were predictable. Specifically, all recovered credentials, including the web passwords used by FT-II's OEM vendors, were static and consisted of a fixed structure such as a vendor identifier followed by a numeric value. Moreover, the remaining unrecoverable credentials were derivable. For example, the FT-I's SSH `admin` and web `admin` passwords can be derived from the last four digits of the serial number, whereas the SSH `root` password can be derived from the full serial number and software version. Therefore, an attacker could construct the `admin` password by obtaining the serial number from alternative sources, such as TR-069 messages or physical device labels. Brute-force guessing was also feasible, given that the password space consisted of only 10,000 possible combinations. With `admin` privilege, the attacker could then retrieve the full serial number and software version and derive the `root` password.

Security impact. These findings indicate that credential randomization is rarely enforced across femtocell devices, or is implemented in a predictable manner such that credentials can be easily guessed or derived. As a result, the use of predictable credentials substantially increases the risk of both local and remote compromise.

E. Exposure of Management Services to WAN Interface (V4)

Femtocell devices typically run local management services such as web, Telnet, or SSH for administrative purposes. While these services are intended for operator-side access via LAN or private infrastructure, we found multiple cases where management interfaces were directly accessible from the WAN interface, potentially enabling remote access without prior authorization.

Results. The FT-I exposed the SSH interface on TCP port 27149. The FT-II exposed the proprietary Telnet-based LMT service on TCP port 50000. The FT-III exposed the web interface on TCP port 443 and an SSH service on TCP port 27149. The FT-IV did not expose any management port we identified to the WAN interface.

Security impact. These WAN-exposed management services introduce significant remote attack surfaces. Notably, all the 4 LMT services we identified on the WAN interfaces of the 3 femtocell devices have predictable credentials (see Table II). So femtocells directly connected to the public Internet are susceptible to remote compromise. Even when deployed behind NAT or firewalls, they remain vulnerable to co-located attackers (from the same subnet).

F. TR-069 Authentication Weakness (V5)

TR-069 (CWMP) is widely used for remote provisioning of femtocells by connecting to an Auto Configuration Server (ACS). According to 3GPP TS 33.320 [10], when the ACS resides on the public Internet and is not accessed via a SeGW, the femtocell must authenticate the ACS via Transport Layer Security (TLS). This requirement is critical to prevent impersonation and unauthorized remote control. If ACS authentication is absent, an on-path attacker can impersonate the ACS and deliver crafted TR-069 commands. This enables a range of attacks, including modifying configuration parameters such as the TR-069 server address, SeGW address, management interface credentials, subscriber access policies, or radio frequency settings; forcing device reboots or resets; extracting sensitive parameters (e.g., IMSI, Ki); or issuing firmware download commands. If the device fails to validate the download source or integrity, the attacker may install a malicious firmware image, resulting in persistent compromise.

We first attempted to connect all four femtocell devices in our testbed to the operator core network, by connecting their WAN interfaces to a router with Internet access. Among them, only the FT-I and FT-II successfully established connections with the core network of Operator-I (referred to as OP-I), a major operator with hundreds of millions of subscribers. In contrast, the FT-III and FT-IV failed to connect, as the SeGW did not respond. This behavior is expected for factory-default devices without operator provisioning or with outdated SeGW configuration. Since TR-069 authentication relies on device-side configuration, we restricted our analysis to the FT-I and FT-II. For these devices, we captured TR-069 traffic and examined whether ACS authentication was performed. Due to ethical considerations, we configured the femtocells to be accessible only to UEs under our control and refrained from sending or injecting any packets to the ACS. We will provide a detailed discussion of our ethical considerations in Section X.

Results. The FT-I connects to a vendor-operated ACS on the public Internet and transmits TR-069 requests in plaintext without any authentication. These messages include sensitive metadata such as the device's serial number, model, firmware version, IMSI, and MME address. In contrast, the FT-II uses an ACS endpoint on a private operator network accessed via an IPsec tunnel to the SeGW. Although the TR-069 session is protected by IPsec, weaknesses in the tunnel establishment process (see Section V-B) allow an attacker to hijack the connection. After hijacking, we confirmed that the device did not authenticate the ACS during TR-069 sessions. While this behavior does not violate the requirements of TS 33.320, it enables an attacker to issue TR-069 commands, for example to modify the femtocell configuration.

Security impact. On both devices we analyzed, TR-069 sessions were established without authenticating the ACS. On the FT-I, this occurred over the public Internet without encryption, while on the FT-II, the ACS was reached via IPsec but still unauthenticated. These behaviors allow an on-path co-located attacker to issue arbitrary provisioning commands. In particular, the attacker could read or modify critical femtocell configuration, such as retrieving AKA credentials (IMSI, OPc, and Ki), starting the `sshd` or `telnetd` services, changing web credentials, or altering firewall rules to grant external access to the LMT interface. Although the specification mandates proper authentication, whether it is enforced depends on operator-controlled deployment choices.

To demonstrate exploitability, we deployed a fake ACS server on the local network and redirected the femtocell's TR-069 traffic to it. The device initiated a session without authentication. Our server responded with crafted messages to retrieve parameter lists via `GetParameterNames`. We then used `GetParameterValues` to extract AKA credentials, and `SetParameterValues` to modify configuration parameters, followed by a `Reboot` command, which was immediately accepted and executed. This successful attack confirms that the TR-069 client did not validate the ACS.

G. Summary of Findings

Our findings reveal that 4G femtocells continue to suffer from widespread and critical security flaws, with little progress since the 3G era, and even worse. Among the 5 vulnerabilities we examined, V3 (predictable credentials) and V4 (exposed management interfaces) represent novel attack surfaces unexamined in previous works. While prior research on 3G femtocells examined debug interfaces (V1), they found the interfaces to be either inactive or protected by root credentials [56], [2]. We observed fully active and exploitable interfaces (e.g., JTAG) in 4G femtocells. Moreover, while flash-based credential extraction [8] was previously demonstrated in local settings, our V2 analysis exposes this issue across multiple vendors and highlights its potential for remote exploitation. In addition, while previous work [2] exploited insufficient TR-069 authentication (V5) to compromise a single local device in their test setup, we show that a co-located attacker can leverage the same issue to perform remote attacks. In summary, our study uncovers new vulnerabilities, confirms that legacy flaws persist, and reveals an even broader attack surface in modern 4G femtocell deployments.

Root causes. The root cause of V1 is a hardware deployment issue attributable to the vendor, who did not disable the JTAG interface (e.g., via eFuse). The root causes of V2, V3, and V4 are due to deficiencies in the vendor's software implementation. Specifically, the vendor did not adopt sufficiently secure methods for per-device credential generation and secure hash algorithms, and did not adequately restrict management services to the LAN interface through software configurations. The root cause of V5 (TR-069 Authentication Weakness) is attributable to a deployment oversight, as authentication of TR-069 server was not enabled even when the HeMS was deployed in public network environments.

In addition, we observed notable OEM relationships that reveal a broader supply chain impact. Specifically, we identified 13 OEM vendors associated with FT-I and 7 OEM vendors associated with FT-II (including FT-I and FT-II themselves). These findings suggest that the vulnerabilities we discovered may not be limited to the specific devices analyzed but could affect a wider range of femtocell devices, which highlights the need for improved femtocell security practices in both hardware and software development and deployment.

V. Security Implications

Based on the security issues identified in Section IV, an attacker can compromise a femtocell and cause severe real-world impact on mobile subscribers and potentially even the core network. This section presents a series of controlled attack experiments to demonstrate the feasibility and consequences of such compromise. Specifically, Section V-A illustrates how a compromised femtocell can be used to attack subscribers and potentially the core network. Section V-B shows that an attacker can hijack the IPsec tunnel without requiring root privileges on the femtocell. This approach enables more flexible and efficient attacks compared to executing them directly on the constrained femtocell device. Section V-C demonstrates that, despite the presence of standard security mechanisms, these protections can often be easily bypassed in practice. Section V-D presents an end-to-end example demonstrating how our findings can be leveraged by an attacker.

A. Threats to Subscribers and the Core Network

Golde et al. [2] discussed how compromised 3G femtocells can be used to threaten both mobile subscribers and the core network. Given the substantial differences in core network architecture and fundamental services between 3G and 4G, we revisit the threats posed by compromised 4G femtocells to users and the core network. We implemented a Man-In-The-Middle (MITM) attack on the FT-II connected to the OP-I network, as discussed in Section V-B, to enable convenient testing of eavesdropping, injection, and hijacking attacks. Note that we only conducted attack tests targeting a test UE under our control and did not perform any attack testing against the core network. Therefore, these tests do not pose any threat to

the operational core network or real users. We describe our ethical considerations in detail in Section X.

Data service. We connected a UE (i.e., a smartphone) to the FT-II and observed that all data service traffic of this UE was transmitted over GTP-U between the femtocell and the S-GW, without integrity or confidentiality protection. This allows attackers to eavesdrop on, inject, and hijack user traffic. To demonstrate the feasibility of traffic hijacking, we successfully intercepted DNS requests initiated by the UE and responded with a forged IP. Figure 3 illustrates this by tampering with the DNS response for a visited website to point to an attacker-controlled server. We also successfully spoofed the user's identity to establish direct communication with a server under our control. This enables Economic Denial-of-Sustainability (EDoS) attacks, where large volumes of traffic are consumed under the user's account, causing data exhaustion or billing overages without the user's awareness.

Fig. 3: Hijacking subscriber data service via a compromised femtocell. The UE was tricked into visiting an attacker-controlled server by tampering with the DNS response for `jd.com`, which then returned spoofed content (i.e., `Hacked`).

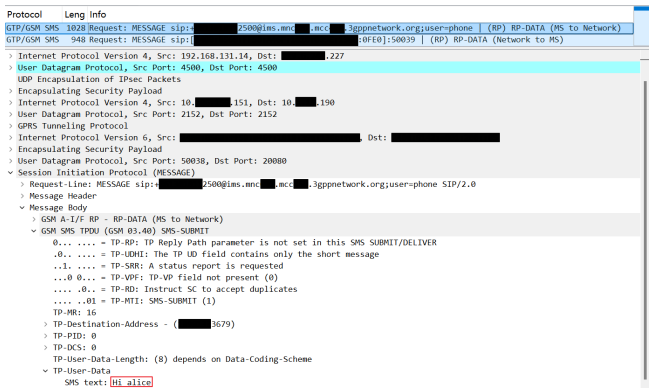

Fig. 4: Eavesdropping on subscriber SMS through a compromised femtocell.

Voice calls. 4G networks introduced Voice over LTE (VoLTE), which is based on the IP Multimedia Subsystem (IMS) for managing voice call signaling and media transport. We connected a VoLTE-enabled smartphone to the FT-II and conducted experiments involving making and receiving calls. We successfully captured plaintext Real-time Transport Protocol (RTP) traffic, which can be used to reconstruct the content of the voice calls.

Short messages. In 4G networks, Short Message Service (SMS) is increasingly delivered via the IMS rather than the traditional NAS-based SMS. We connected a UE with IMS-based SMS enabled to the FT-II and performed sending and receiving of SMS messages. We observed that the full SMS contents are contained in the plaintext payloads of Session Initiation Protocol (SIP) MESSAGE requests. This exposure allows an attacker to obtain verification codes, password reset links, and private conversations, posing threats to both personal privacy and account security. Figure 4 shows an example where an attacker eavesdrops on such messages through a compromised femtocell.

Core network threats. Due to ethical considerations, we did not conduct active exploitation of the core network. Instead, we analyze the security risks based on the exposure of core network protocols to a compromised femtocell. As shown in Figure 1, an attacker can at minimum leverage the S-GW's user-plane GTP-U protocol and the MME's control-plane S1AP protocol, both of which are exposed to the femtocell. Depending on the operator's firewall and routing configurations, the attacker may also be able to access additional control-plane protocols such as GTP-C and Diameter. All of these protocols, originally designed for deployment in trusted environments, may be vulnerable when exposed to untrusted entities and collectively create a broad attack surface in the core network. Previous research has highlighted the severity of these risks. Zhang et al. [47] analyzed the security risks associated with publicly exposed GTP-U and GTP-C ports. They proposed and experimentally demonstrated several practical attacks, including DoS, session hijacking, and user-plane data injection. Bennett et al. [48] conducted fuzzing of core network protocols on six open-source and one commercial core network implementation, demonstrating that crafted GTP and S1AP messages (e.g., malformed messages) can lead to memory corruption and DoS. This highlights the potential for compromised femtocells to significantly undermine the security and stability of the mobile core network.

B. Hijacking IPsec Tunnel

Femtocells typically establish IPsec tunnels to securely connect to the operator's Security Gateway (SeGW), using IKE for key negotiation followed by ESP for encrypted communication. One approach for attackers to control traffic between the femtocell and the core network is to hijack the IPsec tunnel. Once the attacker obtains the root privilege, it becomes possible to hijack the IPsec tunnel [2]. Borgaonkar et al. [56] presented another practical attack against the implementation of the EAP-SIM protocol in a 3G femtocell device they tested to decrypt IPsec traffic. EAP-SIM is a mobile network authentication protocol used in 2G and 3G networks.

We further demonstrated a novel IPsec hijacking attack that requires no root privileges, and validated it on the FT-II. This attack exploits weaknesses in the authentication configuration: the certificate-based device authentication is not used; instead, the femtocell and SeGW perform mutual authentication using a combination of non-EAP pre-shared keys and EAP-AKA, the latter of which is based on recoverable AKA credentials. To hijack the IPsec tunnel, we implemented a MITM attack on the IKE-based authentication exchange between the FT-II and the SeGW. By extracting the necessary credentials (the pre-

shared key, IMSI, Ki, and OPc) from firmware, and reverse engineering the proprietary Milenage implementation, we were able to establish ourselves as the legitimate SeGW to the femtocell and as the legitimate femtocell to the SeGW, achieving a full MITM. This enabled decryption and real-time interception of both user and control plane traffic, compromising the confidentiality and integrity of all communications traversing the IPsec tunnel. It also allowed us to inject packets (e.g., ping, GTP packets) into the IPsec tunnel, enabling potential attacks against the core networks that would otherwise be impossible without root privileges on the femtocell.

C. Bypassing Security Mechanisms

3GPP TS 33.320 [10] defines several additional security mechanisms, including location verification and Closed Subscriber Group (CSG)-based access control, which can limit the threats posed by compromised femtocells to users. However, we experimentally verified that these security restrictions can be bypassed, enabling more impactful attacks.

Location verification. According to TS 33.320 Section 8.1 and Section 5.4 [10], the HeMS is required to perform location verification for HeNB deployments to meet security, regulatory, and operational requirements. This process relies on at least one of several types of information, including the femtocell's public IP address, broadband access identifiers, neighboring macro-cell information, or GNSS-based coordinates. Upon analysis, we found that both femtocell devices we obtained (i.e., FT-I and FT-II) only employed public IP location-based restrictions, requiring the femtocell to use a public IP address from a specific administrative region to connect to the SeGW. However, this security measure can be easily bypassed using a proxy. To demonstrate this, we conducted an experiment with the FT-I originally deployed in Region-A. We physically relocated the femtocell to Region-B, approximately 1000 km away, and connected it to a gateway deployed in Region-B, thereby obtaining Internet access with a public IP address geolocated to Region-B. Packet captures confirmed that the SeGW rejected IKE requests from the new public IP address. We then placed an *OpenWrt*-based router between the femtocell and the gateway, forwarding all traffic through a server with a public IP address from Region-A. *OpenWrt* allowed us to implement the required proxying and traffic redirection in our experiment. Figure 5 demonstrates our experimental settings. As a result, the femtocell was successfully connected to the SeGW. We then connected our own UE to this femtocell and confirmed that it operated as expected. Given the compact size and portability of femtocell devices, an attacker can easily deploy movable rogue base stations at arbitrary locations by leveraging proxy-based redirection and a portable power supply. This significantly expands the potential attack surface and facilitates flexible, location-independent attacks on mobile subscribers.

Closed Subscriber Group (CSG). The CSG mechanism enables access control for subscribers connecting to femtocells. In the HeNB access scenario, access control is enforced by the MME, as specified in TS 36.300 [72]. Femtocells can be configured to operate in different access modes, and when configured in closed access mode, only a set of specifically provisioned subscribers are permitted to connect. This limits the attacker's ability to target arbitrary mobile subscribers through a compromised femtocell. However, in our experimental environment, all the femtocells were not configured in closed access mode, which is consistent with practical deployments where femtocells may be intended for use in public spaces (e.g., underground parking garages). Moreover, we experimentally verified that all four femtocell models allow modification of the access mode after compromise. For example, the FT-I exposes this configuration via its web interface under *LTE Setting* → *Advanced* → *eNodeB Settings* → *accessMode*. When the femtocell is not operating in closed access mode, the attacker is able to target any nearby subscribers who connect to the compromised femtocell, enabling the attacks described in Section V-A.

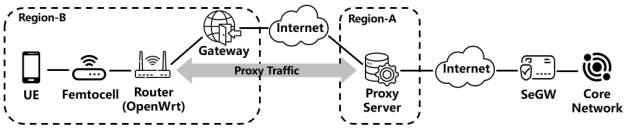

Fig. 5: Demonstration setup for location verification bypass.

D. An End-to-End Example

An attacker aims to eavesdrop on SMS messages and voice calls of a specific target. The attacker first applies for a femtocell from the operator and exploits V2 (credential extraction) to recover the device's SSH credentials, thereby gaining root access. The femtocell is then connected to an *OpenWrt*-based router that proxies traffic to an operator-approved IP address, bypassing the location verification mechanism. The attacker powers both devices using a portable battery pack and provides Internet connectivity through a mobile data modem. These components can be concealed within a backpack. By physically approaching the target so that the target's handset camps on the attacker-controlled femtocell, the attacker can eavesdrop on the target's SMS messages and voice calls.

VI. Identify Internet-Exposed Femtocells

To assess real-world femtocell exposure and associated potential security risks, we conducted an Internet-wide measurement. While prior work has examined protocol-level vulnerabilities for individual devices, to the best of our knowledge, no study has systematically measured femtocell exposure at scale. Existing cyberspace search engines (e.g., Shodan [73], Censys [74]) do not support femtocell-specific identification, further motivating the need for dedicated measurement.

A. Feature Analysis

To identify unknown femtocell deployments, we analyzed network-facing protocol features that may serve as potential indicators. This analysis was guided by standard specifications as well as empirical observations from the femtocell devices we acquired. Specifically, we considered four categories: a)

IKE: Femtocells typically establish IPsec tunnels to the operator's SeGW, with tunnel negotiation handled by the IKE protocol (specifically, IKEv2 as specified in TS 33.320 Section 4.4.5[10]). As IKE is a peer-to-peer protocol, femtocells commonly listen on UDP port 500 to accept inbound IKE requests, making this port a potential indicator of IPsec-related femtocell functionality. b) **TR-069:** Femtocells are remotely managed by the HeMS via the TR-069 protocol. The TR-069 specification [75] mandates support for the Connection Request mechanism, which requires the device to expose an HTTP-based endpoint, typically on TCP port 7547. The presence of this service provides a reliable signal of TR-069 compliance. c) **Web-based LMT:** Web-based LMTs on femtocells may disclose identifying information through HTTP responses or TLS certificates, revealing textual indicators (e.g., `HeNB`) that the device is a femtocell. d) **SSH, Telnet, and other services:** Services such as Telnet and SSH are not standardized in femtocell specifications and vary widely across vendors. Moreover, unlike web-based LMTs, these interfaces typically do not expose any information that clearly indicates the device is a femtocell. As a result, SSH and Telnet are not suitable for identifying unknown femtocell deployments. Nonetheless, we used them to identify known devices in our testbed to support validation. In contrast, other common services such as Network Time Protocol (NTP) are too prevalent in unrelated device types to provide meaningful discrimination.

B. Protocol-Based Discovery of Femtocells

Empirical analysis of scan ports. As discussed in Section VI-A, our scan targets included IKEv2, TR-069, web-based LMT, SSH, and Telnet. We now describe how we selected target port numbers for these services. For IKEv2, which runs over UDP and requires sending data-bearing packets without prior port discovery, we aimed to minimize scanning load by probing only port 500, which is typically used as the initial negotiation port according to RFC 7296 [76]. For TCP-based services, we extended our scan beyond standard and known femtocell ports to account for potential non-standard deployments. To identify additional candidate ports for each protocol, we evaluated three major cyberspace search engines: Censys [74], Shodan [73], and FOFA [77]. Due to Shodan's lack of IKE filtering and Censys's limited port statistics, we relied on FOFA to extract the top 20 most common ports for HTTP, HTTPS, SSH, and Telnet. In addition, since TR-069 is not a typical web protocol and is not directly identifiable by the major cyberspace search engines, we aimed to include additional scan ports in order to improve the coverage of our method. We first manually assembled a list of HTTP header keywords based on an analysis of the TR-069 specification. Representative keywords include `CPE`, `CWMP`, and `SOAP`, where CWMP is an alternative name for TR-069, and SOAP defines the XML-based envelope used to encode TR-069 messages. We then used FOFA to filter for devices whose HTTP headers contained these keywords, aggregated the top ports observed in the results, and removed duplicates, yielding a final set of 19 additional candidate port numbers. Combining these results, we finalized the target port sets as follows: 45 ports for HTTP/HTTPS, 21 for Telnet, and 21 for SSH.

Large-scale measurement. We conducted an Internet-wide IPv4 scan to identify femtocell exposure. Each IP was probed for IKEv2 on UDP port 500. If IKEv2 was detected, we further scanned HTTP/HTTPS (45 ports), Telnet (21 ports), and SSH (21 ports). Scanning was performed from four servers (two in the US and two in France). We used *XMap* [78] to scan for IKEv2 presence and to probe port availability of other target protocols, because it supports custom UDP payloads and efficient multi-port scanning. After confirming port availability, we used *ZGrab2* [79] to perform application-layer handshakes and collect protocol-specific responses, as it is optimized for layer-7 interaction. We limited the aggregate scanning rate across all four servers to 10,000 packets per second for *XMap* and 200 target ports per second for *ZGrab2*, to reduce the potential impact on scanned systems. For IKEv2, we sent only an `IKE_SA_INIT` request. For TR-069 and web-based LMT, we issued HTTP `GET` requests; for Telnet, we completed the TCP handshake and recorded banners; for SSH, we performed key exchange and extracted metadata (e.g., server key and algorithm suite). Our measurement adhered to the Menlo Report's ethical guidelines [80]. Detailed discussions of ethical considerations are provided in Section X.

Classification. Our goal was to discover femtocells that are accessible from the public Internet, as such exposure may enable remote attacks. However, a key challenge is that some devices expose only a subset of services, and the protocol-level features of femtocells are not always distinctive, making it difficult to accurately identify femtocell deployments. To avoid overclaiming femtocell exposure, we adopted a conservative design that prioritized minimizing false positives (i.e., incorrect inclusion) over minimizing false negatives (i.e., missed cases). Accordingly, we applied a classification approach based on both protocol exposure and textual indicators.

We define *negative indicators* as textual cues suggesting a device is not a femtocell (e.g., `VPN`, `Firewall`), and *positive indicators* as those suggesting it is (e.g., `femto`, `FAP`, where `FAP` stands for `Femto Access Point`). These indicators were curated through empirical analysis and manual inspection of specifications and tested femtocells. We searched for them in HTTP headers (e.g., `Server`, `WWW-Authenticate`), HTTP bodies (e.g., `<title>` tag), and TLS certificate fields (e.g., `Organizational Unit`). Overall, this approach assigns each device a confidence level: *potential*, *confident*, or *highly confident*. Only IKE-exposing devices were considered. Devices with negative indicators were excluded. The remaining were labeled as *potential* if TR-069 was present, *confident* if positive indicators were found, and *highly confident* if both TR-069 and positive indicators were observed. Devices lacking both were discarded to reduce false positives. We use the term *femtocell candidate* to refer to any device not excluded by this process—that is, any device labeled as *potential*, *confident*, or *highly confident*. As a result, we identified a total of 86,108 femtocell candidates, of which 52,768 were labeled as *highly confident*, 720 as *confident*, and 32,620 as *potential*.

C. Clustering Analysis

To better understand the characteristics of identified devices, we performed clustering analysis. Devices of the same model or vendor often exhibit similar protocol behaviors due to shared or closely related firmware and software stacks. Clustering thus helps group femtocells likely originating from the same or related product lines. Our pipeline consisted of five steps: data preparation, feature extraction, dimensionality reduction, feature concatenation, and clustering. Specifically, we applied unsupervised clustering based on protocol responses (e.g., HTTP headers, TLS certificates, and IKEv2 messages). These responses were first transformed into structured features using TF-IDF encoding, reduced in dimensionality, then augmented with service availability indicators (i.e., open ports) and normalized to construct the final feature vectors. We evaluated four clustering algorithms: K-Means, Gaussian Mixture Models (GMM), Density-Based Spatial Clustering of Applications with Noise (DBSCAN), and Hierarchical DBSCAN (HDBSCAN), each tested over a wide range of parameter configurations. To determine the best-performing model, we used three selection criteria: noise ratio, i.e., the percentage of unclustered points (ideally less than 5% to avoid inflated scores), the external metric *purity* and the internal metric *silhouette score*. DBSCAN provided the best trade-off between purity and clustering quality while keeping the noise ratio low. We selected DBSCAN (ε = 0.02, min_samples = 2) for our final clustering. Due to space limits, details of the clustering can be found in Appendix A.

Clustering results. Applying the above method, we obtained a total of 1162 clusters. The five largest clusters contain 52,538 (61.01%), 4,412 (5.12%), 3,118 (3.62%), 1,943 (2.26%), and 1,495 (1.74%) devices, respectively. The overall distribution reveals a pronounced long-tail shape. In this paper, we use *Cluster X* to refer to the X-th largest cluster. Table III lists the ten largest clusters. Although geolocation was not used as a clustering feature, devices in each of the top 10 clusters are predominantly located within a single country, indicating a strong degree of homogeneity within each cluster. Furthermore, all clusters except Cluster 7 are associated with entities whose names clearly suggest telecommunication providers, implying that these clusters likely correspond to deployments managed by telecom operators.

Case study-I. We begin by analyzing Cluster 1, the largest group in our dataset with 52,538 devices. As shown in Table III, all devices are located in South Korea and expose a TR-069 service on port 8082. The responses of TR-069 consistently return `401 Unauthorized`, include the string `realm="femtoAP"` (indicates a femtocell access point) in the `WWW-Authenticate` header and `gSOAP/2.8` in the `Server` header. Among them, 77 devices also expose an SSH port. The different characteristics of these SSH services are summarized in Table V (Appendix B), which may reflect multiple firmware versions or device variants from the same vendor. While no web services are exposed in Cluster 1, we identified two smaller clusters with a total of 7 devices, as shown in Table VI (Appendix B), that share the same TR-069 response pattern as Cluster 1. These devices expose web interfaces that display the name of a Korean femtocell vendor directly in the webpage content. Figure 6a shows one example. Notably, 3 of the 7 devices are also located in South Korea. Given the consistent TR-069 responses and geographic overlap, we infer that Cluster 1 likely consists of femtocells from the Korea-based vendor (or OEM variants) with disabled web interfaces.

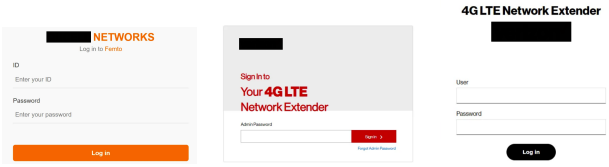

(a) Cluster 1. (b) Cluster 15. (c) Cluster 20.

Fig. 6: Examples of exposed web-based LMT pages.

Case study-II. Cluster 15 and Cluster 20, which contain 391 and 254 devices respectively, are the two largest clusters featuring visual web-based LMT interfaces, where "visual" refers to fully rendered web pages as opposed to mere `WWW-Authenticate` headers. Based on the content of the web pages, these clusters appear to correspond to two different models or firmware versions from the same femtocell vendor, as their web pages display the same vendor name. The devices in these two clusters are all classified as *confident* in our method, as their TLS certificates have Issuer Common Names `enodeb.askey-cons4g.vzwfemto.com` and `Casa Systems Small Cells SubCA`, respectively. Figure 6b and Figure 6c show the web pages of devices from Cluster 15 and Cluster 20, respectively.

D. Evaluation

In this section, we evaluate both our measurement results and clustering quality. As noted earlier, accurately identifying femtocells is challenging due to incomplete service exposure, limited protocol-level distinctiveness, and the absence of ground truth. To address this, we adopted a two-part strategy. First, we assessed whether our method could identify Internet-facing devices matching the tested femtocell models and examined the assigned confidence labels. Second, we manually sampled and validated candidates using cyberspace search engines to evaluate overall effectiveness.

Identification of known models. We used the SSH and Telnet features of our tested femtocells to identify known models found in the measurement. Among the four devices analyzed, FT-III and FT-IV were determined to be OEM variants of FT-I and FT-II, respectively. We therefore center our analysis on FT-I and FT-II devices. For FT-I, as described in Section IV-E, our tested model exposes an SSH interface on TCP port 27149. Our measurements identified 185 devices that ran SSH on this port and matched the host key of the FT-I. All the matched devices were grouped into three clusters, distinguished by the presence or absence of TR-069 and web-based LMT. Among

TABLE III: Characteristics of the ten largest clusters. ASN Org shows the dominant AS organization. Mgmt. Port is the number of devices with exposed SSH or Telnet management ports.

Rank	Cluster Size (#; %)	Cumulative (#; %)	Dominant IP Owner*(%)	Dominant Country (%)	TR-069 Port	Mgmt. Port (#; %)
1	52,538 (61.01)	52,538 (61.01)	Korea Telecom (99.90)	South Korea (100.00)	8082	77 (0.15)
2	4,412 (5.12)	56,950 (66.14)	Telefonica (100.00)	Germany (100.00)	7170	1 (0.02)
3	3,118 (3.62)	60,068 (69.76)	Telefonica (99.94)	Germany (100.00)	7170	0 (0.00)
4	1,943 (2.26)	62,011 (72.02)	Telefonica (100.00)	Germany (100.00)	7170	0 (0.00)
5	1,495 (1.74)	63,506 (73.75)	Telecom Italia S.p.A. (54.78)	Italy (99.93)	7170	17 (1.14)
6	1,460 (1.70)	64,966 (75.45)	Telstra Limited (97.60)	Australia (100.00)	7547	1 (0.07)
7	981 (1.14)	65,947 (76.59)	UBS ADSL range (16.51)	New Zealand (90.72)	30005	1 (0.10)
8	923 (1.07)	66,870 (77.66)	Telecom Italia S.p.A. (55.36)	Italy (100.00)	7170	15 (1.63)
9	859 (1.00)	67,729 (78.66)	Eircom (98.95)	Ireland (98.95)	7547	0 (0.00)
10	857 (1.00)	68,586 (79.65)	Telstra Limited (97.78)	Australia (100.00)	7547	1 (0.12)

* We manually consolidated IP owners belonging to the same organization. For example, "Eircom" and "Eircom Limited" were merged to "Eircom".

them, 8 devices without TR-069 were identified as *confident* femtocells, while the remaining 177 were identified as *highly confident*. For FT-II, we did not identify any public devices with Telnet behavior exactly matching that of our tested model. Nonetheless, the ability to successfully rediscover FT-I devices demonstrates that our method is effective in locating real femtocell deployments on the public Internet.

Manual validation. To further validate the devices identified by our method, we randomly sampled 20 IPs from each of the top 20 largest clusters for manual inspection. For each sample, we used FOFA and Shodan to assess: (1) whether the IP address had been labeled with any device-type labels; (2) whether protocol-level responses within the cluster were consistent; and (3) whether other open ports revealed identifiable device traits. As a result, none of the sampled devices had associated device-type labels, and protocol responses were highly consistent within each cluster. In addition to the previously analyzed *confident* and *highly confident* clusters (i.e., Cluster 1, Cluster 15, and Cluster 20), which exposed clear femtocell identifiers, we found that 9 *potential* clusters contained vendor-related identifiers such as `FRITZ!OS`, and 8 *potential* clusters lacked any identifiers that could reveal device type or vendor affiliation. Although vendor-related identifiers do not conclusively determine device type, they allow partial inference. For example, the presence of `FRITZ!OS` suggests that the corresponding devices may be `FRITZ!Box` routers with TR-069 capabilities. Overall, this validation result aligns with our classification design: *highly confident* and *confident* clusters correspond to devices exposing strong femtocell indicators, whereas *potential* clusters consist of devices with limited or ambiguous identifiers.

Sources of false negatives and false positives. While our validation supports the overall plausibility of our method, both false positives and false negatives may still occur. For false negatives, devices that disable IKE or both TR-069 and web-based LMT on their public interfaces, that suppress all femtocell-identifying information in their responses, or that restrict responses to known IP ranges, may be missed by our method. For false positives, some non-femtocell devices may present network features similar to those of femtocells. For instance, broadband routers or home gateways that support TR-069 for Internet Service Provider (ISP) remote management, TR-069-enabled network cameras or firewalls without explicit filtering indicators (e.g., `CAMERA`), and devices misconfigured to expose both IKE and TR-069 may all be mistakenly identified. However, such devices are typically assigned to the *potential* category due to the lack of strong femtocell-specific indicators, minimizing their impact on overall accuracy.

TABLE IV: Top PTR FQDN naming patterns.

eTLD+1	Count (#; %)	Representative FQDN patterns
telefonica.de	9,821 (11.41)	dynamic-[*].pool.telefonica.de
telecomitalia.it	3,997 (4.64)	host-[*].[retail|business].telecomitalia.it
telstra.net	2,327 (2.70)	cpe-[*].asp.telstra.net
bell.ca	1,198 (1.39)	[bras-base|ipagstaticip]-[*].dsl.bell.ca
inspire.net.nz	888(1.03)	[*].[dsl.]sta.inspire.net.nz

E. Characterizing Femtocell Candidates

We further analyze the discovered devices' network characteristics, including IP ownership distribution, geographic distribution, and reverse DNS (PTR) records.

IP ownership and geographic distribution. We identified a total of 2,450 distinct IP owners among the discovered devices. The top five IP owners are Korea Telecom (52,513, 60.99%), Telefonica (10,077, 11.70%), Telstra (3,392, 3.94%), Telecom Italia S.p.A. (2,619, 3.04%), and Eircom (867, 1.01%), together accounting for 80.68% of all observed devices. We examined the top 50 IP owners and found that the vast majority are telecommunications service providers. Only a few entries which use descriptive or infrastructure-level labels are not telecom-related, such as *IP pools* (202, 0.23%). This observation suggests that the majority of the discovered devices are likely femtocell deployments, although we cannot rule out the presence of other telecom-related services. Figure 7 shows the global distribution of femtocell candidates, as inferred from their IP address allocation. These devices span 103 countries, with the top three being South Korea (52,751, 61.05%), Germany (10,904, 12.66%), and Italy (4,934, 3.38%). We suspect that the geographical bias may be influenced primarily by two factors. First, regions with more developed cellular infrastructure tend to have higher coverage demands and broader femtocell deployment. Second, some femtocells are not exposed on the Internet due to their configuration or implementation, yet they may still be susceptible to attacks from local or co-located adversaries.

PTR analysis. We performed reverse DNS lookups to obtain Fully Qualified Domain Names (FQDNs) of femtocell

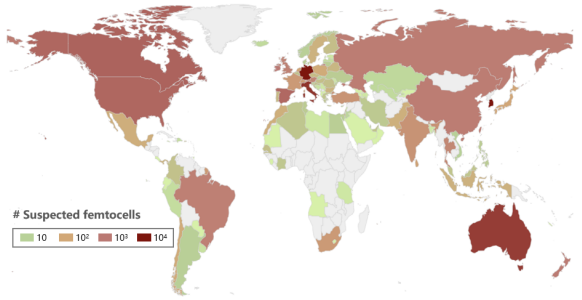

Fig. 7: Global distribution of identified femtocell candidates.

candidates, as FQDNs may reveal organizational or device-type information. Among 86,108 femtocell candidates, 27,763 (32.24%) had valid PTR records, spanning 1,107 distinct eTLD+1 domains. Table IV lists the top 5 eTLD+1 domains and representative FQDN patterns, where `[*]` denotes placeholders (e.g., IP or city names), `[<content>]` indicates optional components, and `[<a>|<b>]` specifies mutually exclusive alternatives. Specifically, labels such as `dynamic`, `host`, `cpe`, `ipagstaticip`, and `dsl` are commonly associated with customer-premises equipment, suggesting that the devices are positioned at the network edge and align with typical femtocell deployment scenarios. In contrast, `bras-base` suggests affiliation with ISP infrastructure (i.e., BRAS, or Broadband Remote Access Server) rather than femtocells.

F. Security Analysis

The exposure of femtocells on the public Internet, particularly the exposure of management services, can lead to severe security threats. Among the 86,108 femtocell candidates identified in our protocol-based measurement, although we did not have direct access to these devices for testing, we found that 1,459 exposed SSH services and 139 exposed Telnet services, suggesting that they may be vulnerable to remote attacks. In addition, using the SSH protocol feature, we discovered 185 femtocells matching the model of FT-I, with publicly accessible management ports. As described in Section IV-C and Section IV-D, we successfully extracted management credentials from these devices and confirmed that many used predictable passwords, directly enabling remote attacks. Moreover, even devices that do not expose their management interfaces to the Internet (for example, due to NAT) can still be targeted by co-located attackers on the same local network. As discussed in Section V, the compromise of a single femtocell can introduce significant risks to both the core network and subscribers. Given their potential impact, such risks warrant serious attention.

VII. Mitigation

A. Enhancing Femtocell Security

3GPP femtocell security specification [10] lacks requirements on platform security, such as hardware and software hardening. For example, it does not mandate disabling or strict protection of production-stage debug interfaces (V1). It also provides no guidance on securing vendor-customized services like LMT services and their credentials (V2, V3), or on addressing risks from unintended service exposure (V4).

Further, 3GPP does not appear to have any Security Assurance Specification (SCAS) for femtocells. Note that 3GPP does have SCAS specifications for regular base stations (i.e., TS 33.216 [81] for 4G eNodeB, TS 33.511 [82] for 5G gNodeB). The lack of SCAS specifications may explain why some femtocell products do not comply with the femtocell security specification TS 33.320 [10], since there are no security test cases defined. For example, although TS 33.320 mandates TLS authentication when the ACS resides on the public Internet, the FT-I established TR-069 connections without TLS (V5). Besides, TS 33.320 requires AKA credentials for hosting party authentication to be stored in a UICC, but none of the tested femtocells comply, storing them instead in flash memory.

We believe it is necessary for 3GPP to further improve femtocell security standards and to define SCAS specifications to motivate and facilitate vendors for security enhancement.

B. Mitigating Threats to Subscribers

The root cause of threats from misbehaving femtocells to subscribers is that UEs cannot authenticate the femtocell, risking unintended connections. The current solution in femtocell security standard is to use CSG to inform UEs that are not part of the groups to not connect. However, CSG advertised by a femtocell is not authenticated either, and thus can be forged.

As a short-term fix, the UE and core network can verify femtocell identity and CSG via integrity-protected NAS messages during connection setup, enabling detection of forged identity/CSG and redirection to legitimate cells. A long-term mitigation is to ensure the authenticity of network broadcasting messages so that UEs can verify them and only connect to legitimate and authorized cells. The lack of authenticity of broadcasting messages exists in all generations of mobile networks from 2G to 5G. Hopefully, this study, along with prior efforts [27], [3], [5], [83], could motivate 3GPP to solve this problem in 6G.

C. Mitigating Threats to Core Networks

As a short-term solution to mitigate abuse of core network interfaces (e.g., GTP-C and Diameter) that are unintentionally exposed to femtocells, SeGWs should strictly filter femtocell-originated traffic, allowing only authorized destinations and message types. A long-term solution is to enhance femtocell architecture by mandating the deployment of a femtocell gateway (HeNB-GW) for both control and user plane to further shield the core network from femtocells. In the current 3GPP femtocell architecture, HeNB-GW is optional for the control plane, and not applicable to the user plane, leaving critical user plane entities (e.g., the SGW in 4G and the UPF in 5G) directly exposed. Control plane entities (e.g., the MME in 4G and the AMF in 5G) may also be directly exposed to femtocells if the HeNB-GW is not deployed. Note that the SeGW does not

terminate any control plane or user plane interface, thus cannot fully protect core networks.

VIII. Discussions

A. Vulnerability Disclosure

We have disclosed our findings to affected vendors, GSMA, and 3GPP. Among the vendors, FT-I's vendor has acknowledged the reported vulnerabilities and confirmed that they have been fixed in the latest firmware version. GSMA has acknowledged our report and assigned us a CVD identifier (CVD-2025-0106). GSMA also notified its member operators about our research and asked them to read our paper to ensure the security of their femtocell implementations. We discussed some of our results at several 3GPP SA3 meetings, starting at SA3#122 (May 2025), and highlighted that the current security standards for 4G femtocells are insufficient. As a result, 3GPP SA3 approved both a study item to further enhance the security of femtocells in 5G [11], and a work item to define SCAS for 5G femtocells [12]. In parallel, we also shared our findings with several operators whose production femtocells were confirmed to be affected.

B. Limitations

Femtocell devices coverage. Due to constraints in obtaining femtocell devices (e.g., sourced from second-hand markets), the range of vendors and models available to us was limited. Furthermore, only some of the acquired devices (specifically, FT-I and FT-II) were able to connect to the operator's core network. In addition, we could not determine whether these devices were running the latest firmware versions. Nevertheless, since they remain in use in real-world deployments (i.e., capable of connecting to the core network) and lack timely remote updates, our findings still carry important real-world security implications. We will further discuss the generalizability of our findings in Section VIII-C.

Security analysis in compromising femtocells. Although our analysis was relatively systematic, it focused on common security issues that are more likely to be found in devices from other vendors. As a result, device-specific issues such as web service vulnerabilities were not covered in our testing.

Threats to core networks. For ethical reasons, we did not perform security tests against production core networks. However, as discussed in Section V-A, an attacker can at least inject traffic directly into core-network elements by abusing the S-GW's user-plane GTP-U protocol and the MME's control-plane S1AP protocol. An open-source testbed for validating such threats is currently infeasible: essential components such as the SeGW and the HeMS are not available from existing open-source implementations (e.g., Open5GS [84], OAI [85]). Implementing these components may require significant engineering effort, and we therefore leave this to future work.

Protocol-based measurement. While our study presents a detailed analysis of Internet-exposed femtocell devices, several limitations remain. Specifically, our protocol-based discovery relies on IKEv2, TR-069, and web-based LMT features, which may result in false positives (such as non-femtocell devices exposing similar services) and false negatives (such as femtocells with these services disabled or blocked by firewalls) as discussed in Section VI-D. Moreover, as we applied a keyword-based method to identify TR-069 responses, we may miss those with vendor-specific or obfuscated behavior. In addition, our Internet measurements are subject to timing and network coverage limitations. Some devices or networks may filter scanning traffic based on source IP address or other factors, leading to incomplete visibility. Despite these limitations, our work provides valuable insights into the Internet-exposed femtocell landscape and highlights the urgent need for enhanced security mechanisms.

C. Generalizability of Our Findings

Although the number of femtocells we could obtain is limited, the identified vulnerabilities stem from common design and deployment patterns across femtocells, thus indicating some level of generalizability. The vendors of FT-I and FT-II, whose femtocells we evaluated in our security analysis, are also recognized as "Top Players in the Femtocell Market" in an independent industry report by SkyQuest [86], suggesting that our findings may extend to other influential vendors. Moreover, we observed extensive firmware reuse across 20 femtocell OEM vendors (Section IV-C). To further examine the generalizability, we analyzed two additional femtocell models, FT-V and FT-VI. These devices belong to the same OEM family, which differs from the four previously tested models. Our experiments show that FT-V exhibits vulnerabilities V1, V2, V3 and V4, whereas FT-VI exhibits V1, V2, V3 and V5, through which root access can be obtained. In addition, FT-VI can connect to the core network of OP-II, another major operator with hundreds of millions of subscribers. Taken together, these observations suggest that our findings are likely to generalize across a broader range of femtocell deployments.

IX. Conclusion

This paper presents a systematic security analysis of femtocell devices, identifying 5 critical hardware and software vulnerabilities. These vulnerabilities can be exploited by attackers with different capabilities, including local, co-located, and remote attackers. We validated all 5 issues on 4 commercial femtocells, each of which was susceptible to multiple issues and compromisable. Our controlled experiments and analysis further demonstrate that compromised femtocells can severely threaten mobile users and the core network. We also show that location verification can be bypassed, allowing attacker-controlled femtocells to be deployed arbitrarily, amplifying the threat to users. We further present the first measurement of Internet-exposed femtocells, identifying 86,108 candidates, including 1,598 with publicly accessible management interfaces potentially affected by the discovered vulnerabilities. These findings underscore the urgent need for enhanced security in femtocell design, deployment, and management.

X. Ethics Considerations

To ensure that our research conforms to ethical principles, we follow the ethical principles outlined in the Menlo Re-

port [80], which emphasize respect for persons, beneficence, justice, and respect for law and public interest. Our study was designed and conducted in accordance with these principles to avoid potential harms and ensure responsible research practices. We also reflected on broader ethical issues discussed in prior work on cybersecurity research ethics [87], which offered additional perspectives on the responsibilities and real-world implications of conducting security research.

Internet-wide scanning. For the Internet-wide scanning portion of our study, we employed well-established tools (i.e., *XMap* and *ZGrab2*) that incorporate scanning best practices, including randomized target selection to reduce the likelihood of overloading any particular network. We used a scanning rate of no more than 10,000 packets per second for *XMap* and 200 targets per second for *ZGrab2*, and performed all scans with the knowledge and permission of our hosting provider. We did not send any attack-oriented or intrusive payloads; instead, all payloads were harmless and designed solely to elicit basic protocol responses. Furthermore, we hosted an HTTP page on our scanning servers describing the purpose of the research and our ethical commitments. This page also included a dedicated contact email address, allowing owners of scanned systems to request exclusion from future scans or removal from aggregated results.

Sensitive information disclosure. Given that vulnerable femtocell devices may still be deployed in operational networks and considering the potential severity of the identified security issues, we deliberately anonymize all femtocell vendor identities throughout this paper to avoid creating unnecessary risk for deployed systems. In addition, throughout the preparation and writing of this paper, we ensured that no password details or other sensitive information were disclosed. In accordance with established practices for mobile network security research, we reported the vulnerabilities to GSMA and 3GPP. In parallel, we independently disclosed the issues to affected operators and vendors. These disclosures are also mentioned in Section VIII-A.

Avoiding impact on the core network. All experiments were conducted using legally provisioned femtocells with valid credentials. Interactions involving the core network were strictly limited to passive traffic analysis, and no packets were injected or modified. To further reduce potential resource consumption, we conducted these tests during periods of low user activity, such as late at night. In addition, we constrained both the duration and behavior of the tests: each femtocell remained connected to the core network for no longer than 20 minutes each time, and on the UE side, we avoided high-bandwidth operations and limited activity to basic procedures such as making and receiving phone calls, as well as sending and receiving SMS messages. As such, the experimental setup posed no risk of disruption or undue resource consumption in the core network.

Avoiding impact on normal subscribers. To avoid any impact on normal mobile subscribers, we exclusively used our own test UE in all testing. To further prevent nearby subscriber UEs from accidentally connecting to our femtocell, we enabled whitelist-based access control and configured the femtocell to accept only the IMSI provisioned on our test Universal Subscriber Identity Module (USIM) card. This ensured that no potential harm was caused to other subscribers.

While our institution does not have an Institutional Review Board (IRB), the experimental controls described above ensure that our work adheres to ethical principles and avoids any potential impact on users or the live network.

ACKNOWLEDGMENT

We are grateful to the anonymous reviewers for their constructive feedback and valuable suggestions. We also appreciate GSMA and 3GPP for discussions on disclosure, mitigation, and specification aspects of our findings. We thank the operators and femtocell vendors involved for their engagement. This work is supported by Taishan Scholar Program. Yiming Zhang is in part supported by NSFC #62302258.

REFERENCES

[1] "Femtocell market report 2025," https://www.thebusinessresearchcompany.com/report/femtocell-global-market-report, 2025.

[2] N. Golde, K. Redon, and R. Borgaonkar, "Weaponizing femtocells: The effect of rogue devices on mobile telecommunications," in *19th Annual Network and Distributed System Security Symposium, NDSS 2012, San Diego, California, USA, February 5-8, 2012*. The Internet Society, 2012. [Online]. Available: https://www.ndss-symposium.org/ndss2012/weaponizing-femtocells-effect-rogue-devices-mobile-telecommunications

[3] H. Yang, S. Bae, M. Son, H. Kim, S. M. Kim, and Y. Kim, "Hiding in plain signal: Physical signal overshadowing attack on LTE," in *28th USENIX Security Symposium, USENIX Security 2019, Santa Clara, CA, USA, August 14-16, 2019*, N. Heninger and P. Traynor, Eds. USENIX Association, 2019, pp. 55–72. [Online]. Available: https://www.usenix.org/conference/usenixsecurity19/presentation/yang-hojoon

[4] Y. Zhang, B. Liu, C. Lu, Z. Li, H. Duan, S. Hao, M. Liu, Y. Liu, D. Wang, and Q. Li, "Lies in the air: Characterizing Fake-Base-Station spam ecosystem in China," in *Proceedings of the 2020 ACM SIGSAC Conference on Computer and Communications Security*. New York, NY, USA: Association for Computing Machinery, 2020, p. 521–534. [Online]. Available: https://doi.org/10.1145/3372297.3417257

[5] A. Singla, R. Behnia, S. R. Hussain, A. A. Yavuz, and E. Bertino, "Look before you leap: Secure connection bootstrapping for 5g networks to defend against fake base-stations," in *ASIA CCS '21: ACM Asia Conference on Computer and Communications Security, Virtual Event, Hong Kong, June 7-11, 2021*, J. Cao, M. H. Au, Z. Lin, and M. Yung, Eds. ACM, 2021, pp. 501–515. [Online]. Available: https://doi.org/10.1145/3433210.3453082

[6] K. S. Mubasshir, I. Karim, and E. Bertino, "Fbsdetector: Fake base station and multi step attack detection in cellular networks using machine learning," *arXiv preprint arXiv:2401.04958*, 2024.

[7] H. Schmidt and B. Butterly, "Attacking basestations," https://media.defcon.org/DEF%20CON%2024/DEF%20CON%2024%20presentations/DEF%20CON%2024%20-%20Hendrik-Schmidt-Brian-Butter-Attacking-BaseStations.pdf, 2016, dEF CON 24.

[8] X. Zou, "4g lte man in the middle attacks with a hacked small cells," https://www.telecomsinfrastructure.com/2019/10/4g-lte-man-in-middle-attacks-with.html, 2019, hITB GSEC 2019.

[9] 3GPP, "Security aspects of NR Femto," 3rd Generation Partnership Project (3GPP), Technical Specification (TS) 33.545, 3 2025, version 19.0.0.

[10] ——, "Security of Home Node B (HNB) / Home evolved Node B (HeNB)," 3rd Generation Partnership Project (3GPP), Technical Specification (TS) 33.320, 4 2024, version 18.0.0.

[11] 3GPP, "Study on security aspects for NR Femto phase2," 3rd Generation Partnership Project (3GPP), Technical Report TR 33.746, December 2025, release 20 (Draft).

[12] ——, "Security Assurance Specification (SCAS) for NR Femto," 3rd Generation Partnership Project (3GPP), Tech. Rep. TS 33.546, December 2025, release 20, V0.2.0.

[13] 3GPP, "Service requirements for Home Node B (HNB) and Home eNode B (HeNB)," 3rd Generation Partnership Project (3GPP), Technical Specification (TS) 22.220, 4 2024, version 18.0.1.

[14] ——, "5G NR Femto (5G_Femto) Work Item," https://portal.3gpp.org/desktopmodules/WorkItem/WorkItemDetails.aspx?workitemId=1060048, 2024, 3GPP Work Item 1060048, Release 19, PCG Approved.

[15] J. Arkko and H. Haverinen, "Extensible Authentication Protocol Method for 3rd Generation Authentication and Key Agreement (EAP-AKA)," RFC 4187, Jan. 2006. [Online]. Available: https://www.rfc-editor.org/info/rfc4187

[16] 3GPP, "Evolved Universal Terrestrial Radio Access Network (E-UTRAN); S1 Application Protocol (S1AP)," 3rd Generation Partnership Project (3GPP), Technical Specification (TS) 36.413, 12 2024, version 18.3.0.

[17] ——, "General Packet Radio System (GPRS) Tunnelling Protocol User Plane (GTPv1-U))," 3rd Generation Partnership Project (3GPP), Technical Specification (TS) 29.281, 1 2025, version 19.1.0.

[18] Z. Li, W. Wang, C. Wilson, J. Chen, C. Qian, T. Jung, L. Zhang, K. Liu, X. Li, and Y. Liu, "FBS-Radar: Uncovering Fake Base Stations at scale in the wild," in *24th Annual Network and Distributed System Security Symposium, NDSS 2017, San Diego, California, USA, February 26 - March 1, 2017*. The Internet Society, 2017. [Online]. Available: https://www.ndss-symposium.org/ndss2017/ndss-2017-programme/fbs-radar-uncovering-fake-base-stations-scale-wild/

[19] Z. Zhuang, X. Ji, T. Zhang, J. Zhang, W. Xu, Z. Li, and Y. Liu, "FBSleuth: Fake Base Station forensics via radio frequency fingerprinting," in *Proceedings of the 2018 on Asia Conference on Computer and Communications Security*, ser. ASIACCS '18. New York, NY, USA: Association for Computing Machinery, 2018, p. 261–272. [Online]. Available: https://doi.org/10.1145/3196494.3196521

[20] E. C. Jimenez, P. K. Nakarmi, M. Näslund, and K. Norrman, "Subscription identifier privacy in 5g systems," in *International Conference on Selected Topics in Mobile and Wireless Networking, MoWNeT 2017, Avignon, France, May 17-19, 2017*. IEEE, 2017, pp. 1–8. [Online]. Available: https://doi.ieeecomputersociety.org/10.1109/MoWNet.2017.8045947

[21] B. Hong, S. Bae, and Y. Kim, "GUTI reallocation demystified: Cellular location tracking with changing temporary identifier," in *25th Annual Network and Distributed System Security Symposium, NDSS 2018, San Diego, California, USA, February 18-21, 2018*. The Internet Society, 2018. [Online]. Available: https://www.ndss-symposium.org/wp-content/uploads/2018/02/ndss2018_02A-4_Hong_paper.pdf

[22] H. Lin, "Lte redirection attack—forcing targeted lte cellphone into unsafe network," *Unicorn Team—Radio and Hardware Security Research*, 2016.

[23] "Catchercatcher - mobile network assessment tools - SRLabs open source projects," https://opensource.srlabs.de/projects/mobile-network-assessment-tools/wiki/CatcherCatcher, 2014.

[24] D. Abodunrin *et al.*, "Detection and mitigation methodology for Fake Base Stations detection on 3G/2G cellular networks." Master's thesis, 2015.

[25] SRLabs, "Snoopsnitch - mobile network security tool," https://github.com/srlabs/snoopsnitch.

[26] C. Zhang, "Malicious base station and detecting malicious base station signal," *China Communications*, vol. 11, no. 8, pp. 59–64, 2014.

[27] S. R. Hussain, M. Echeverria, A. Singla, O. Chowdhury, and E. Bertino, "Insecure connection bootstrapping in cellular networks: the root of all evil," in *Proceedings of the 12th Conference on Security and Privacy in Wireless and Mobile Networks, WiSec 2019, Miami, Florida, USA, May 15-17, 2019*. ACM, 2019, pp. 1–11. [Online]. Available: https://doi.org/10.1145/3317549.3323402

[28] A. Lotto, V. Singh, B. Ramasubramanian, A. Brighente, M. Conti, and R. Poovendran, "BARON: base-station authentication through core network for mobility management in 5g networks," in *Proceedings of the 16th ACM Conference on Security and Privacy in Wireless and Mobile Networks, WiSec 2023, Guildford, United Kingdom, 29 May 2023 - 1 June 2023*, I. Boureanu, S. Schneider, B. Reaves, and N. O. Tippenhauer, Eds. ACM, 2023, pp. 133–144. [Online]. Available: https://doi.org/10.1145/3558482.3590187

[29] H. Gao, Y. Zhang, T. Wan, J. Zhang, and H. Duan, "On evaluating delegated digital signing of broadcasting messages in 5g," in *2021 IEEE Global Communications Conference (GLOBECOM)*, 2021, pp. 1–7.

[30] Y. Dong, T. Wan, T. Wu, and S. R. Hussain, "Evaluating time-bounded defense against rrc relay in 5g broadcast messages," in *18th ACM Conference on Security and Privacy in Wireless and Mobile Networks*, 2025, pp. 236–241.

[31] S. R. Hussain, M. Echeverria, O. Chowdhury, N. Li, and E. Bertino, "Privacy attacks to the 4g and 5g cellular paging protocols using side channel information," in *26th Annual Network and Distributed System Security Symposium, NDSS 2019, San Diego, California, USA, February 24-27, 2019*. The Internet Society, 2019. [Online]. Available: https://www.ndss-symposium.org/ndss-paper/privacy-attacks-to-the-4g-and-5g-cellular-paging-protocols-using-side-channel-information/

[32] M. Kotuliak, S. Erni, P. Leu, M. Röschlin, and S. Capkun, "Ltrack: Stealthy tracking of mobile phones in LTE," in *31st USENIX Security Symposium, USENIX Security 2022, Boston, MA, USA, August 10-12, 2022*, K. R. B. Butler and K. Thomas, Eds. USENIX Association, 2022, pp. 1291–1306. [Online]. Available: https://www.usenix.org/conference/usenixsecurity22/presentation/kotuliak

[33] N. Lakshmanan, N. Budhdev, M. S. Kang, M. C. Chan, and J. Han, "A stealthy location identification attack exploiting carrier aggregation in cellular networks," in *30th USENIX Security Symposium, USENIX Security 2021, August 11-13, 2021*, M. Bailey and R. Greenstadt, Eds. USENIX Association, 2021, pp. 3899–3916. [Online]. Available: https://www.usenix.org/conference/usenixsecurity21/presentation/lakshmanan

[34] M. Chlosta, D. Rupprecht, C. Pöpper, and T. Holz, "5g suci-catchers: still catching them all?" in *WiSec '21: 14th ACM Conference on Security and Privacy in Wireless and Mobile Networks, Abu Dhabi, United Arab Emirates, 28 June - 2 July, 2021*, C. Pöpper, M. Vanhoef, L. Batina, and R. Mayrhofer, Eds. ACM, 2021, pp. 359–364. [Online]. Available: https://doi.org/10.1145/3448300.3467826

[35] T. Oh, S. Bae, J. Ahn, Y. Lee, T. D. Hoang, M. S. Kang, N. O. Tippenhauer, and Y. Kim, "Enabling physical localization of uncooperative cellular devices," in *Proceedings of the 30th Annual International Conference on Mobile Computing and Networking, ACM MobiCom 2024, Washington D.C., DC, USA, November 18-22, 2024*, W. Shi, D. Ganesan, and N. D. Lane, Eds. ACM, 2024, pp. 1530–1544. [Online]. Available: https://doi.org/10.1145/3636534.3690709

[36] A. Paci, G. Bologna, I. Palamà, and G. Bianchi, "Flashcatch: Minimizing disruption in IMSI catcher operations," in *18th ACM Conference on Security and Privacy in Wireless and Mobile Networks, WiSec 2025, Arlington, VA, USA, 30 June 2025- 3 July 2025*, M. Albanese, L. da Silva, A. Ranganathan, and J. Seifert, Eds. ACM, 2025, pp. 124–135. [Online]. Available: https://doi.org/10.1145/3734477.3734705

[37] D. Rupprecht, K. Kohls, T. Holz, and C. Pöpper, "Breaking LTE on layer two," in *2019 IEEE Symposium on Security and Privacy, SP 2019, San Francisco, CA, USA, May 19-23, 2019*. IEEE, 2019, pp. 1121–1136. [Online]. Available: https://doi.org/10.1109/SP.2019.00006

[38] ——, "IMP4GT: impersonation attacks in 4g networks," in *27th Annual Network and Distributed System Security Symposium, NDSS 2020, San Diego, California, USA, February 23-26, 2020*. The Internet Society, 2020. [Online]. Available: https://www.ndss-symposium.org/ndss-paper/imp4gt-impersonation-attacks-in-4g-networks/

[39] J. Xing, S. Yoo, X. Foukas, D. Kim, and M. K. Reiter, "On the criticality of integrity protection in 5g fronthaul networks," in *33rd USENIX Security Symposium, USENIX Security 2024, Philadelphia, PA, USA, August 14-16, 2024*, D. Balzarotti and W. Xu, Eds. USENIX Association, 2024. [Online]. Available: https://www.usenix.org/conference/usenixsecurity24/presentation/xing-jiarong

[40] N. Ludant, M. Vomvas, S. Dimou, and G. Noubir, "Low-layer attacks against 4g/5g networks," in *18th ACM Conference on Security and Privacy in Wireless and Mobile Networks, WiSec 2025, Arlington, VA, USA, 30 June 2025- 3 July 2025*, M. Albanese, L. da Silva, A. Ranganathan, and J. Seifert, Eds. ACM, 2025, pp. 248–255. [Online]. Available: https://doi.org/10.1145/3734477.3734725

[41] S. R. Hussain, M. Echeverria, I. Karim, O. Chowdhury, and E. Bertino, "5greasoner: A property-directed security and privacy analysis framework for 5g cellular network protocol," in *Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security, CCS 2019, London, UK, November 11-15, 2019*, L. Cavallaro, J. Kinder, X. Wang, and J. Katz, Eds. ACM, 2019, pp. 669–684. [Online]. Available: https://doi.org/10.1145/3319535.3354263

[42] E. Bitsikas and C. Pöpper, "Don't hand it over: Vulnerabilities in the handover procedure of cellular telecommunications," in *ACSAC '21: Annual Computer Security Applications Conference, Virtual Event, USA, December 6 - 10, 2021*. ACM, 2021, pp. 900–915. [Online]. Available: https://doi.org/10.1145/3485832.3485914

[43] C. Peng, C.-y. Li, G.-H. Tu, S. Lu, and L. Zhang, "Mobile data charging: New attacks and countermeasures," in *Proceedings of the 2012 ACM Conference on Computer and Communications Security*, ser. CCS '12. New York, NY, USA: Association for Computing Machinery, 2012, p. 195–204. [Online]. Available: https://doi.org/10.1145/2382196.2382220

[44] C. Peng, C.-Y. Li, H. Wang, G.-H. Tu, and S. Lu, "Real threats to your data bills: Security loopholes and defenses in mobile data charging," in *Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security*, ser. CCS '14. New York, NY, USA: Association for Computing Machinery, 2014, p. 727–738. [Online]. Available: https://doi.org/10.1145/2660267.2660346

[45] Y. Go, E. Jeong, J. Won, Y. Kim, D. F. Kune, and K. Park, "Gaining control of cellular traffic accounting by spurious TCP retransmission," in *21st Annual Network and Distributed System Security Symposium, NDSS 2014, San Diego, California, USA, February 23-26, 2014*. The Internet Society, 2014. [Online]. Available: https://www.ndss-symposium.org/ndss2014/gaining-control-cellular-traffic-accounting-spurious-tcp-retransmission

[46] H. Hong, H. Kim, B. Hong, D. Kim, H. Choi, E. Lee, and Y. Kim, "Pay as you want: Bypassing charging system in operational cellular networks," in *International Workshop on Information Security Applications*. Springer, 2016, pp. 148–160.

[47] Y. Zhang, T. Wan, Y. Yang, H. Duan, Y. Wang, J. Chen, Z. Wei, and X. Li, "Invade the walled garden: Evaluating gtp security in cellular networks," in *2025 IEEE Symposium on Security and Privacy (SP)*. IEEE Computer Society, 2024, pp. 28–28.

[48] N. Bennett, W. Zhu, B. Simon, R. Kennedy, W. Enck, P. Traynor, and K. R. B. Butler, "Ransacked: A domain-informed approach for fuzzing LTE and 5g ran-core interfaces," in *Proceedings of the 2024 on ACM SIGSAC Conference on Computer and Communications Security, CCS 2024, Salt Lake City, UT, USA, October 14-18, 2024*, B. Luo, X. Liao, J. Xu, E. Kirda, and D. Lie, Eds. ACM, 2024, pp. 2027–2041. [Online]. Available: https://doi.org/10.1145/3658644.3670320

[49] S. Thorn, K. V. English, K. R. B. Butler, and W. Enck, "5gac-analyzer: Identifying over-privilege between 5g core network functions," in *Proceedings of the 17th ACM Conference on Security and Privacy in Wireless and Mobile Networks*, ser. WiSec '24. New York, NY, USA: Association for Computing Machinery, 2024, p. 66–77. [Online]. Available: https://doi.org/10.1145/3643833.3656134

[50] M. Akon, T. Yang, Y. Dong, and S. R. Hussain, "Formal analysis of access control mechanism of 5g core network," in *Proceedings of the 2023 ACM SIGSAC Conference on Computer and Communications Security, CCS 2023, Copenhagen, Denmark, November 26-30, 2023*, W. Meng, C. D. Jensen, C. Cremers, and E. Kirda, Eds. ACM, 2023, pp. 666–680. [Online]. Available: https://doi.org/10.1145/3576915.3623113

[51] T. Yang, S. K_S, A. Arumugam, and S. Hussain, "Feedback-guided api fuzzing of 5g network," 2025.

[52] M. Akon, M. Toufikuzzaman, and S. R. Hussain, "From control to chaos: A comprehensive formal analysis of 5g's access control," in *IEEE Symposium on Security and Privacy, SP 2025, San Francisco, CA, USA, May 12-15, 2025*, M. Blanton, W. Enck, and C. Nita-Rotaru, Eds. IEEE, 2025, pp. 1081–1100. [Online]. Available: https://doi.org/10.1109/SP61157.2025.00141

[53] Y. Dong, T. Yang, A. Al Ishtiaq, S. M. M. Rashid, A. Ranjbar, K. Tu, T. Wu, M. S. Mahmud, and S. R. Hussain, "Corecrisis: Threat-guided and context-aware iterative learning and fuzzing of 5g core networks," 2025.

[54] R. Rajavelsamy, J. Lee, and S. Choi, "Towards security architecture for home (evolved) nodeb: challenges, requirements and solutions," *Security and Communication Networks*, vol. 4, no. 4, pp. 471–481, 2011.

[55] R. Borgaonkar, N. Golde, and K. Redon, "Femtocells: a poisonous needle in the operator's hay stack," https://media.blackhat.com/bh-us-11/Borgaonkar/BH_US_11_RaviNicoKredon_Femtocells-WP.pdf, 2011, Black Hat 2011.

[56] R. Borgaonkar, K. Redon, and J. Seifert, "Security analysis of a femtocell device," in *Proceedings of the 4th International Conference on Security of Information and Networks, SIN 2011, Sydney, NSW, Australia, November 14-19, 2011*, M. A. Orgun, A. Elçi, O. B. Makarevich, S. A. Huss, J. Pieprzyk, L. K. Babenko, A. G. Chefranov, and R. Shankaran, Eds. ACM, 2011, pp. 95–102. [Online]. Available: https://doi.org/10.1145/2070425.2070442

[57] L. Janzen, L. Becker, C. Wiesenäcker, and M. Hollick, "Oh no, my RAN! breaking into an O-RAN 5g indoor base station," in *18th USENIX WOOT Conference on Offensive Technologies (WOOT 24)*. Philadelphia, PA: USENIX Association, Aug. 2024, pp. 101–115. [Online]. Available: https://www.usenix.org/conference/woot24/presentation/janzen

[58] S. Vasile, D. Oswald, and T. Chothia, "Breaking all the things—a systematic survey of firmware extraction techniques for iot devices," in *Smart Card Research and Advanced Applications*, B. Bilgin and J.-B. Fischer, Eds. Cham: Springer International Publishing, 2019, pp. 171–185.

[59] G. Vishwakarma and W. Lee, "Exploiting jtag and its mitigation in iot: A survey," *Future Internet*, vol. 10, no. 12, 2018. [Online]. Available: https://www.mdpi.com/1999-5903/10/12/121

[60] A. Gangolli, Q. H. Mahmoud, and A. Azim, "A systematic review of fault injection attacks on iot systems," *Electronics*, vol. 11, no. 13, 2022. [Online]. Available: https://www.mdpi.com/2079-9292/11/13/2023

[61] MITRE Corporation, "CVE-2019-15894," https://www.cve.org/CVERecord?id=CVE-2019-15894, 2019, accessed: 2025-06-20.

[62] ——, "CVE-2024-10237," https://www.cve.org/CVERecord?id=CVE-2024-10237, 2024, accessed: 2025-06-20.

[63] "Jtagulator," https://grandideastudio.com/portfolio/security/jtagulator/, accessed: 2025-06-05.

[64] SEGGER Microcontroller GmbH, "SEGGER J-Link debug probe series," accessed: 2025-08-02. [Online]. Available: https://www.segger.com/products/debug-probes/j-link/

[65] "iFix RT809 Series," accessed: 2025-08-02. [Online]. Available: http://doc.ifix.net.cn/@rt809

[66] "Binwalk v3," accessed: 2025-08-02. [Online]. Available: https://github.com/ReFirmLabs/binwalk/tree/master

[67] R. Carleton and S. Poznyakoff, "GNU cpio version 2.14." [Online]. Available: https://www.gnu.org/software/cpio/

[68] "hashcat: advanced password recovery," https://hashcat.net/hashcat/, accessed: 2025-05-27.

[69] "CMD5," https://www.cmd5.org/, accessed: 2025-05-27.

[70] "OpenWrt Project," https://openwrt.org/, accessed: 2025-05-27.

[71] A. Nakajima, "Oem finder: Hunting vulnerable oem iot devices at scale," https://i.blackhat.com/eu-19/Thursday/eu-19-Nakajima-OEM-Finder-Hunting-Vulnerable-OEM-IoT-Devices-At-Scale-2.pdf, 2019, Black Hat Europe 2019.

[72] 3GPP, "Evolved Universal Terrestrial Radio Access (E-UTRA) and Evolved Universal Terrestrial Radio Access Network (E-UTRAN); Overall description; Stage 2," 3rd Generation Partnership Project (3GPP), Technical Specification (TS) 33.360, 12 2024, version 18.4.0.

[73] "Shodan: Search engine for the internet of everything," accessed: 2025-08-02. [Online]. Available: https://www.shodan.io

[74] "Censys," accessed: 2025-08-02. [Online]. Available: https://censys.com/

[75] Broadband Forum, "CPE WAN Management Protocol," Broadband Forum, Technical Report (TR) TR-069 Amendment 6 Corrigendum 1, 6 2020.

[76] C. Kaufman, P. E. Hoffman, Y. Nir, P. Eronen, and T. Kivinen, "Internet Key Exchange Protocol Version 2 (IKEv2)," RFC 7296, Oct. 2014. [Online]. Available: https://www.rfc-editor.org/info/rfc7296

[77] "FOFA," accessed: 2025-08-02. [Online]. Available: https://en.fofa.info/

[78] X. Li, "Xmap," accessed: 2025-08-02. [Online]. Available: https://github.com/idealeer/xmap

[79] "Zgrab 2.0," accessed: 2025-08-02. [Online]. Available: https://github.com/zmap/zgrab2

[80] M. D. Bailey, D. Dittrich, E. Kenneally, and D. Maughan, "The menlo report," *IEEE Secur. Priv.*, vol. 10, no. 2, pp. 71–75, 2012. [Online]. Available: https://doi.org/10.1109/MSP.2012.52

[81] 3GPP, "TS 33.216. Security Assurance Specification (SCAS) for the evolved Node B (eNB) network product class," https://portal.3gpp.org/desktopmodules/Specifications/SpecificationDetails.aspx?specificationId=3129.

[82] ——, "TS 33.511. Security Assurance Specification (SCAS) for the next generation Node B (gNodeB) network product class," https://portal.3gpp.org/desktopmodules/Specifications/SpecificationDetails.aspx?specificationId=3444.

[83] Y. Dong, R. Behnia, A. A. Yavuz, and S. R. Hussain, "Securing 5g bootstrapping: A two-layer ibs authentication protocol," *arXiv preprint arXiv:2502.04915*, 2025.

[84] "Open5GS," https://open5gs.org/, 2025.
[85] "OpenAirInterface," https://openairinterface.org/, 2025.
[86] "Femtocell market size, share, and growth analysis," https://www.skyquestt.com/report/femtocell-market, 2025.
[87] K. Macnish and J. van der Ham, "Ethics in cybersecurity research and practice," *Technology in Society*, vol. 63, p. 101382, 2020. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0160791X19306840
[88] A. Falini, "A review on the selection criteria for the truncated svd in data science applications," *Journal of Computational Mathematics and Data Science*, vol. 5, p. 100064, 2022. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S2772415822000244

APPENDIX

A. Clustering Implementation Details

This appendix provides implementation details of our clustering pipeline. We implemented clustering using the `scikit-learn` and `hdbscan` libraries in Python.

Data preparation. In the protocol-based scan, we collected binary IKEv2 responses, HTTP-based TR-069 responses, HTTP/HTTPS responses, and TLS certificates, as well as SSH and Telnet responses for each device. Our goal was to cluster devices of the same model or vendor. However, since the availability of SSH and Telnet often depends on software configuration rather than firmware alone, we excluded these features from clustering. To prepare the remaining data for feature extraction, we converted IKEv2 fields, HTTP/HTTPS headers, and TLS certificate fields into key-value formats.

Feature extraction. We applied Term Frequency–Inverse Document Frequency (TF-IDF) vectorization, with `ngram_range` set to `(1, 2)` to capture both unigrams and bigrams, and `min_df` set to 2 to exclude terms that appear in only a single response (such as unique nonces in the `WWW-Authenticate` header). For IKEv2 responses, HTTP/HTTPS headers, and TLS certificates, we adopted a key-value tokenization strategy, where each key and each value are treated as individual tokens. For HTTP/HTTPS bodies, we used the TF-IDF vectorizer's default tokenization. We also encoded the open status of HTTP-related service ports (i.e., TR-069 and web-based LMT) as multihot vectors and incorporated them as additional features.

Dimensionality reduction and feature concatenation. We applied Truncated Singular Value Decomposition (Truncated SVD) to reduce the dimensionality of the TF-IDF vectors. The number of components was selected based on a cumulative explained variance ratio threshold of 95%, a commonly used criterion [88]. This resulted in 230 dimensions, achieving a cumulative explained variance ratio of 95.003%. Finally, we concatenated the reduced vector with the L2-normalized multihot vector, yielding a combined 275-dimensional feature vector, and then applied L2 normalization to the final concatenated vector.

Clustering algorithm. To select an appropriate clustering algorithm, we compared the performance of four methods: K-Means, GMM, DBSCAN, and HDBSCAN. For each algorithm, we tuned the key parameters to obtain optimal clustering results. As discussed in Section VI-C, we used the proportion of noise points, *purity*, and *silhouette score* as the evaluation metrics for algorithm selection. For K-Means, we varied the number of clusters; for GMM, we adjusted the number of components (i.e., clusters); for DBSCAN, we tuned the neighborhood radius parameter (ε) and the minimum number of samples required to form a core point (`min_samples`); and for HDBSCAN, we varied the minimum cluster size parameter. Specifically, for K-Means and GMM, the number of clusters was enumerated from 2 to 5,000 with different step sizes: step size 1 from 2 to 50, step size 5 from 50 to 100, step size 20 from 100 to 1,000, and step size 200 from 1,000 to 5,000. For DBSCAN, `min_samples` was enumerated from 2 to 10 in increments of 1, from 10 to 30 in increments of 2, and from 30 to 100 in increments of 5; the neighborhood radius ε was varied from 0.001 to 0.01 in increments of 0.001, from 0.01 to 0.1 in increments of 0.01, and from 0.1 to 0.5 in increments of 0.1. For HDBSCAN, the minimum cluster size was enumerated from 2 to 10 in increments of 1, from 10 to 30 in increments of 2, and from 30 to 100 in increments of 5.

B. Additional Clustering Results

TABLE V: Management port features of devices in Cluster 1. Banners omit the `SSH-2.0-` prefix; country codes use ISO abbreviations.

Port	Banner	IP Owner	Country	Count
10022	dropbear_2012.55	Korea Telecom	KR	73
22	dropbear_2015.67	Korea Telecom	KR	3
22	dropbear_2013.58	Korea Telecom	KR	1

TABLE VI: Representative IP addresses and ports in Cluster 1 where the banner reveals the manufacturer name. The last octet of each IP address is anonymized; country codes use ISO abbreviations.

IP	Port	Protocol	IP Owner	Country
125.150.162.x	80	HTTP	Korea Telecom	KR
14.93.30.x	80	HTTP	Korea Telecom	KR
175.235.12.x	80	HTTP	Korea Telecom	KR
198.166.113.x	443	HTTPS	TELUS-DSL-EDTNABZY	CA
204.191.2.x	443	HTTPS	TELUS Communications Inc.	CA
207.102.184.x	443	HTTPS	TELUS-FIBRE-WHRKBC01	CA
209.89.250.x	443	HTTPS	TELUS-FIBRE-CLGRAB49	CA

This appendix presents additional figures and tables that complement the clustering results discussed in Section VI-C. In particular, Table V summarizes the management port features observed in Cluster 1. Table VI lists representative IP addresses of Cluster 1 and their corresponding ports and protocols.