DDoS Attacks in 5G Networks: Characteristics and Analysis of Defense Strategies

Ahmad Ahmadov
Azerbaijan Technical University, Baku, Azerbaijan
ahmad.ahmadov.sz@student.aztu.edu.az

***Abstract*— 5G networks introduce transformative technologies like SDN, NFV, Network Slicing, and MEC that enable ultra-low latency and massive IoT connectivity, but simultaneously create new security vulnerabilities. The software-defined control plane, shared infrastructure, and proliferation of insecure IoT devices make DDoS attacks significantly more sophisticated and damaging than in previous mobile generations. This paper analyzes the evolving DDoS threat landscape in 5G environments and evaluates various defense strategies including SDN-based mitigation, machine learning approaches, network slicing isolation, and Zero Trust Architecture, aiming to assess current protective capabilities and highlight areas for future research.**

Keywords*— *5G security; DDoS attacks; network slicing; defense mechanisms; Zero Trust Architecture; mobile network security

I. INTRODUCTİON

The 5G - is not simply an upgraded version of previous technologies, or a faster version of existing 4G technology. The way this technology is designed, built and implemented is completely different to previous generations of mobile telecommunications. Unlike previous generations, which were built and operated on dedicated hardware appliances that operated in a largely static configuration, 5G is completely built off of software. The 5G core network is built on a cloud-native Service Based Architecture (SBA) that uses software microservices and standard HTTP/2 interfaces for communication between the individual functions of the core network (e.g., authentication, session management and routing). In addition, software-defined networking (SDN) decouples the control logic from the actual physical transmission of packets, enabling the creation of centralized and programmable network management capabilities for the entire network infrastructure. Network Function Virtualization (NFV) replaces dedicated hardware appliances with virtualized versions of the same function on commodity servers. Lastly, network slicing enables multiple logical end-to-end networks to coexist on the same physical infrastructure, with each network specifically created for and servicing the requirement of a particular use case (e.g., autonomous vehicles) [1].

The abilities provided by 5G represent a real change. Ultra-Reliable Low-Latency Communications (URLLC) make possible automated manufacturing or working with vehicles in real-time. URLLC demands less than one-millisecond end-to-end latency which is an achievable goal now but would have seemed impossible only 10 years ago. Enhanced Mobile Broadband (eMBB) provides the primary throughput for real-time video, augmented/virtual reality and dense urban areas with many users. As mMTC will support the millions of sensors, meters and actuators that are rapidly becoming the nervous system of infrastructure today, the total cost/benefit of these technologies is extremely high, and 5G deployment is already in process around most of the developed world.

However, with every architectural choice that forms 5G's foundation, there comes a degree of complexity. The flexibility of each 5G software function to dynamically distribute resources causes service degradation due to incorrectly configured or overloaded network functions at scale in ways that were not possible due to the strength of previous generations' hardware.

Every decision that makes 5G powerful also makes 5G a more complex target. The same software that allows 5G to use its resources in a way also means that if 5G network functions are not set up correctly or get overwhelmed they can stop services from working on a big scale. This is something that was hard to do with the hardware that was made to do one thing and do it well. The physical parts that all the 5G network slices use mean that if someone attacks one slice it can hurt another slice if things are not right. The billions of Internet of Things devices that will connect to 5G are often simple devices with very basic software and not much security. This makes them easy targets for bad people to recruit into botnets on a scale. The result is that Distributed Denial-of-Service attacks, which are already a problem, on the internet become even more dangerous in the 5G context.

The problem of DDoS attacks is something we really need to think about. Cloudflare reported that in 2025 there were 47.1 million DDoS attacks. This is then double the number of attacks in 2024 and a big increase from 2023. The number of network-layer attacks went up a lot with one attack in 2025 targeting the internet with 13.5 million attacks over 18 days. A survey by A10 Networks found that most mobile network security professionals think protecting against DDoS attacks is very important for 5G infrastructure [2].

With all this the people doing research on DDoS attacks are still in the early stages of finding ways to defend against these attacks that will work for 5G. DDoS attacks are a big concern for 5G. This paper wants to help with that. It does things. First it looks at what makes 5G different from 4G and how that makes it easier for DDoS attacks to happen. It uses what other people have written and the official 5G security rules to figure this out. DDoS attacks are a problem for 5G. It groups the types of DDoS

DOI: 10.25045/forensic.gov.2026.62

attacks that target 5G into categories. It notes that these DDoS attacks are not bigger versions of old attacks but are actually new and different threats. DDoS attacks are a challenge for 5G. Third it compares six ways to defend against these DDoS attacks including using networks, artificial intelligence and trusted security systems. It looks at what's good and bad about each way and which ones will work best in different situations. This analysis is meant to be helpful, for both researchers and people who need to protect their networks from DDoS attacks and keep 5G safe from DDoS attacks.

II. 5G Architecture and the Expanded Attack Surface

A. Core Architectural Elements

The 3GPP TS 23.501 [1] and TS 33.501 [3] define the 5G standalone (SA) architecture as having three integrated technology components. Together, these three technologies are used to define both the capabilities that provide compelling reasons to deploy the 5G network, and the vulnerabilities that create significant challenges for securing the 5G network.

Software-defined networking (SDN) decouples the control plane (the logic for determining how traffic will be routed and forwarded within the network), from the user plane (the actual packet forwarding from source to destination). The separation of these two functions has both operationally powerful and conceptually elegant implications for network operators. With the ability to programmatically enforce policies centrally across an entire network, service providers will be able to utilize a single control point to push uniform security and quality-of-service policies across thousands of packet-forwarding devices. However, while concentration is a great benefit for performance, the centralization of all of the decision-making logic (which was previously distributed across semi-autonomous devices) in the SDN controller transforms the SDN controller into a high-value target. A single successful attack on an SDN controller can create negative consequences that will be quickly propagated across the entire network.

The Virtualized Networks Function, or VNF, virtualizes functions like firewalls, load balancers, and deep packet inspection engines from dedicated hardware provided by different vendors and turns them into software equivalents that run on generic servers or containers. This has the advantage of reducing costs for operators since they can scale their service delivery affordably, introduce new services quickly without having to buy specific hardware dedicated to that service, while also decreasing their dependence on a small number of incumbent vendors. As a side effect of this movement toward virtualized environments, there are also security concerns; throughout the life of the virtualized system, the hypervisor interfaces at every layer, the container runtimes at every layer and the orchestration layers are all potential points of attack from outside the network or virtualized environment [4]. One particular example is CVE-2024-21626, a well-known vulnerability present within the runc container runtime, which has been confirmed to be valid; should this vulnerability be exploited, it would allow an attacker to escape their containerized VNF and inflict significant harm on the 5G Core Environment of the operator [5]. Therefore, there are many ongoing vulnerabilities being identified and patched at this layer regularly.

Network slicing helps separate several logically separate end-to-end networks on one physical shared infrastructure and each with dedicated resource guarantees, security policies, and quality of service parameters. Many will likely argue that the most significant commercial innovation of 5G is slicing. Slicing allows a single operator to offer, e.g., an ultra-reliable low-latency communications (URLLC) slice to support industrial automation (which has strict textual description accuracy requirements) in addition to an enhanced mobile broadband (eMBB) slice to support broadband access for customers and a massive machine-type communications (mMTC) slice to support ubiquitous Internet of Things (IoT) sensors across a city). All of these slices utilize the same physical hardware resources but provide logical separation to the end-users of the slices. However, the same physical resources that make slicing economically viable to consumers also create an opportunity for interference among slices. The isolation of slices from one another in the software does not inherently equate to isolation of slices at the level of CPU cycle usage, memory bandwidth usage, and switching fabric capacity [6] (See Section III for additional discussion.).

Finally, Mobile Edge Computing (MEC) means that compute and storage capacity will be pushed out to the edge of the network and closer to the end-user or device. For latency-sensitive type applications, this capability cannot be overemphasized, as it is the key technology that enables URLLC to become a reality. However, this change in the location of compute and storage resources can have the effect of moving security perimeters from a central data center protected by traditional physical security measures to the outer edges of the network where these types of protections may not always be as readily available. They can also be physically accessed by people you don't know and/or trust (untrusted sources) due to the many edge nodes that have no visible means of preventing access to the nodes. Every edge node will expose a unique attack surface, thus presenting an easily compromised 'security boundary' as opposed to any traditional 'security boundary' previously experienced by operators [7].

B. Comparison with 4G LTE

It is worth pausing to ask: how different is the 5G threat landscape, really, from what operators already dealt with in 4G? After all, DDoS attacks have been a concern for mobile operators for well over a decade. The honest answer is that the differences are substantial - not merely quantitative but qualitative. In 4G LTE, the Evolved Packet Core (EPC) ran on dedicated hardware appliances connected by point-to-point interfaces using protocols like Diameter and GTP-U. The control and user planes were logically separated but physically co-located on proprietary equipment. DDoS defenses in 4G were primarily perimeter-based: rate limiting, static firewall rules, traffic scrubbing at well-defined boundaries between the operator's network and the rest of the Internet. The attack surface, while significant, was at least reasonably predictable.

5G will feature Service-Based Architecture (SBA), which offers a more flexible method of addressing 4G's many point-to-point interfaces. Using SBA, all communications between

the different, software-defined network functions will take place via a mesh of separate HTTP/2 API calls. The control plane has evolved from being a monolithic set of hardware devices to being a distributed, software-based system. Like all distributed software systems, the new control plane will be challenged by consistency, coordination and security issues. A UPF (User Plane Function) will perform next hop routing of packets, but it will consume virtualized computer and memory resources and compete for those resources with other co-located network functions. Network slicing (introducing shared physical resources for logical tenants where such sharing previously did not exist) creates significant new attack vectors (each of which did not exist in the previous generation) and demand new methods of defence that do not have direct equivalencies to those used in the 4G operational environment [8].

III. Taxonomy of DDoS Attacks in 5G Networks

A. Signaling Plane Attacks

The network's "brain," which keeps all of us connected through 5G technology, is made up of several layers of interconnected signals to facilitate this. The two most influential network layers are the AMF and SMF. The AMF acts as the "digital bouncer" for our devices, verifying their identities and monitoring their movement on the network. The SMF contains all of the data sessions on the network that allow our devices to utilize the internet.

The AMF and SMF communicate with each other over two different connection protocols. The 4G networks used a protocol called Diameter, an older and less-used type of network protocol; 5G networks utilize HTTP/2, which is the same protocol used for web-based applications today. While this has improved the network's efficiency, it has also opened up new avenues for cybercriminals who have the necessary skills to use HTTP/2 protocols and, therefore, more opportunities to exploit 5G cellular networks based on their experience by exploiting the core of our cellular networks.

Before creation of 5G networks, LTE signalling attack research has established a well-defined attack pattern. By exploiting vulnerabilities in user attach and handover processes, attackers can generate signalling storms with very little effort and create an amplification effect that disrupts the service level at a much greater cost than their investment [9]. A further potential expansion of the attack vector going into 5G SA deployments is the transition from using specialised Diameter stacks for signalling to using general-purpose HTTP/2 service-based interfaces for performing signalling. This means that there will now be many more software endpoints, and that they will all be diverse, operating on an HTTP/2 basis, compared to the relatively few Diameter implementations.

B. User Plane Attacks

The User Plane Function (UPF) in a 5G network does packet routing and forwarding of subscriber data traffic. UPF instances are virtualized workloads running on an industry-standard server, unlike 4G networks that require dedicated hardware. Therefore, when the UPF is struck by a volumetric DDoS attack, the impact of the attack will not only degrade the UPF's performance but will also cause resource contention with other network functions running on the same physical infrastructure, resulting in widespread performance degradation. This creates a domino effect where services unrelated to the original target of the attack are also degraded; this is a significant departure from 4G architecture, where dedicated hardware EPC components provided physical separation between functions, which made cascading failures much less likely. While the transition to 5G offers operational advantages from virtualisation, such as increased flexible scaling, improved resource efficiency, and faster deployment capabilities, it also presents a security challenge in those multiple workloads running on the same physical infrastructure naturally compete for resources during an attack. Developing strategies to address this fundamental challenge remains one of the most critical unresolved issues facing 5G network operations today.

C. Slice-Specific Attacks

As mentioned previously, Network Slicing is one of the most important characteristics of 5G technology. It is also one of the most complex and difficult security challenges faced by 5G. Since physical resources (CPU, memory, bandwidth, and switching fabric) are shared across slices, an attack that consumes all of the allocated resources of one slice can impact the performance of other slices hosted on the same infrastructure in ways that can't be prevented by just isolating the software. This is referred to as a "cascading slice attack" and is a consequence of the shared physical resource model that supports the economics of slicing.

To illustrate this point, Khan et al. generated DoS/DDoS traffic aimed at 5G network slices and then developed a dataset documenting a reduction in the amount of throughput for an adjacent slice that used the same infrastructure [10]. This kind of research is particularly useful because it moves us from the theoretical risk of a problem to an actual demonstration of its impact. The software isolation of slices does not automatically translate to the physical isolation of those slices; therefore, attackers will find an opportunity in the gap between logical and physical isolation of slices. For operators who are hosting mission-critical slices alongside non-mission-critical slices on a shared infrastructure, this is a security risk that cannot be engineered out of the service without sacrificing the efficiency of the shared resources.

D. IoT/mMTC Botnet Attacks

mMTC 5G capabilities will allow for the deployments of low-power, high-volume devices in smart cities, precision agriculture, industrial Internet of Things (IoT), and many other similar uses. Most of these low-resource devices, by design, do not have sufficient processing power, memory, or battery capacity available to run robust security applications. Many mMTC-based devices operate in environments with infrequent or difficult firmware updates (or no firmware updates). Most are designed to have a life span of years to decades so that if any vulnerabilities are discovered post-deployment these devices may never get patched. All of these characteristics combine to make mMTC-connected devices an excellent choice for use in botnets.

This is not a hypothetical threat. The Mirai botnet was first

utilized in 2016 to perform an enormous volumetric attack by compromising hundreds of thousands of IoT devices primarily consisting of home routers and security cameras to create record high volumes of DDoS attacks. As of late 2025, with the creation of the Aisuru-Kimwolf botnet, which consists largely of compromised Android TV devices connected to the Internet via both broadband and mobile 5G networks, the attack surface has grown significantly larger since 2016. In particular, the Aisuru-Kimwolf botnet performed hyper-volumetric HTTP DDoS attacks exceeding 20 million requests per second against Cloudflare's infrastructure in late 2025 [11]. With the advent of 5G mMTC, millions of devices are expected to connect to the Internet in the years to come. This will lead to an exponential increase in the size and attack capabilities of botnets (a group of hijacked computers used together for malicious intent). Therefore, the defense implications for both operators and their customers will be substantial, and researchers are just beginning to tackle these issues somewhat seriously.

E. MEC-Targeted Attacks

MEC (Mobile Edge Computing) nodes are located either within close proximity to radio access networks or at the edge of radio access networks where they process latency-sensitive or latency-critical workloads for end users. Because of their geographic distribution and direct exposure to the traffic generated by radio access networks, MEC nodes are appealing targets for localized DDoS attacks. If a MEC node that supports a factory automation deployment or fleet management vehicle was successfully attacked, it does not only have an impact on consumers who are streaming videos or music, but it can also directly affect physical safety in that environment. In environments where a production line operates with robotic actuators coordinated within milliseconds, the MEC controller must be operational at all times to prevent unsafe operations. Therefore, if the MEC controller is unresponsive and the throughput of data sent from an A-RAN (Autonomous Radio Access Network) is compromised, all of the robotic actuators that are controlled by that MEC controller will stop working altogether or operate in an unsafe manner.

The asymmetrical risk-reward nature of attacking MEC nodes is one of the reasons why targeting MEC nodes with DDoS attacks represents an extremely high-impact threat. Because MEC nodes are deployed in geographically diverse locations, attackers can select their targets based on factors that would not be factors if they were using a traditional centralized data center to mount their DDoS attacks (e.g., the local economic impact of a factory or the political nature of an infrastructure deployment) [7].

IV. COMPARATIVE ANALYSIS OF DEFENSE APPROACHES

The last 5 years have seen a significant increase in DDoS defence research in 5G networks, covering a wide variety of different defence strategies. Rather than viewing those defence strategies as alternative solutions, they should be seen as layers of defence-in-depth working together to provide multiple layers of protection against an increasingly difficult-to-defend cybersecurity threat environment. Each of the defence layers presents different dimensions of the DDoS threat environment; there are different operational trade-offs associated with each layer; and there are situations where a defence layer is much better suited for a specific deployment. The table below summarises this comparison of the various defence strategies according to four factors: advantages, disadvantages, primary application layer and readiness for deployment (maturity). The next subsections will present detailed descriptions of each defence strategy, based on data from recent peer-reviewed literature to provide an empirical basis for the discussion, rather than purely speculative assumptions.

TABLEL I. COMPARATİVE ANALYSİS OF DDOS DEFENSE APPROACHES IN 5G NETWORKS

Defense Approach	Advantages	Limitations	Layer	Maturity
SDN-based Mitigation	Centralized visibility; dynamic re-routing; fast automated response	Controller is a single point of failure; scalability constraints	Core	High
ML/AI-based Detection	High accuracy; adapts to novel patterns; automates pipeline	Compute-intensive; retraining overhead; adversarial evasion	Edge/ Core	Med-High
3GPP Standard Mechanisms	Standardized baseline; vendor-neutral; built into 5G SA	Slow update cycle; limited flexibility; not for volumetric DDoS	All 5G	High
Network Slice Isolation	Contains blast radius; prevents cross-slice propagation	Resource overhead; cannot prevent physical contention	Slice	Medium
CDN / Scrubbing Centers	Proven at scale; absorbs volumetric floods	Adds latency; incompatible with URLLC; weak vs. app-layer	Edge	Very High
Zero Trust Architecture	Limits lateral movement; continuous verification	Immature 5G standards; complex policy management	Cross-layer	Low-Med

A. SDN-Based Mitigation

The SDN separation of control and data planes gives SDN its primary reason to be an ideal platform for DDoS defense. The SDN controller can monitor the entire network in real-time and see how data flows through it. From this global view, the SDN controller will be able to detect any anomalous behavior patterns very early on, and quickly respond to them by programmatically creating and updating flow rules across all relevant forwarding devices all at once and at line speed. This level of responsiveness far exceeds the manual updates of firewall rules used in perimeter protections built on 4G technologies. Responses to DDoS attacks via SDN would also be difficult, if not impossible, to execute via a traditional network. Sheibani et al. created an architecture of defense; a three-layer model for 5G networks (i.e., SDN and NFV). The data plane uses entropy-based traffic classification, the distributed controller layer utilizes self-organizing maps, and the central controller layer makes the final mitigation decision [12]. This architecture produces high detection rates because the workload is distributed across multiple layers, reducing the load on the central controller while performing its traffic classification tasks. Furthermore, Morteza Sheibani, the first author, died before submitting this manuscript; his co-authors

have carried on his work as a tribute to him. The above illustrates that researchers in this area of research are not simply organizations (i.e., faceless); they are also persons who have a passion for addressing important issues affecting society.

The limitation of SDN-based defenses is that they centralize resource-intensive networking intelligence in the SDN controller node. In the case of a distributed denial of service (DDoS) attack against either the SDN controller interface or the southbound application programming interface (API) connecting to data forwarding devices, the organization would lose its SDN-based defense mechanism when it needed it most. Implementing redundancies (essentially clustering) on SDN controllers will help mitigate this risk, however, increasing operational difficulty while also introducing additional consistency concerns. As SDN-based DDoS defenses are deployed globally by providers of managed SDN services, resiliency of the control plane becomes a top priority in the design phase of any SDN-based service, not just an implementation consideration.

B. ML/AI-Based Detection

In recent years, researchers have been studying how machine-learning techniques can be used to detect Distributed Denial of Service (DDoS) attacks. The rationale behind this is simple: DDoS attacks, when significant, typically generate abundant flows of data with statistical properties - for example, the distribution of flow from each source or the rate at which packets from each source arrive - that differ fundamentally from what is considered "normal" for a given network class; thus, a classifier can learn to discriminate between these two classes of traffic. However, the fundamental challenge to successfully implementing this technology in practice is doing so reliably, in real-time, for high throughput and low latency traffic associated with next-generation wireless technologies, such as 5G; data that performs well on a benchmark dataset may not perform similarly when subjected to real-world traffic patterns.

Authors Chen et al. [13] proposed an innovative hybrid model that combines the strengths of longitudinal and spatial analysis through a combination of Convolutional Neural Networks (CNN) and Long Short-Term Memory (LSTM) architectures to detect and mitigate DDoS in real-time within privately deployed 5G environments. Their evaluation of this model utilized a custom-built 5G test environment to demonstrate effective results across multiple data sources. Private 5G deployments (e.g., factories, campuses, hospitals, distribution centers, and so forth) represent an exciting opportunity for the use of this technology since they face the same operational scalability challenges presented by 5G as well as unique operational constraints typically not encountered in larger commercial enterprise environments with limited availability of dedicated security personnel.

To resolve the constraints of privacy inherent in training a model for detecting Distributed Denial of Service (DDoS) attacks across multiple organizations, Pourahmadi et al. [14] have proposed the use of federated learning. In this system, models are trained locally on the data of each participant, and then participants share only parameters of the model (as opposed to the raw traffic data) within a federation. By using this federated learning system, the authors demonstrated competitive accuracy in detecting DDoS attacks while avoiding the centralized exposure of data. This allows for the use of federated learning in scenarios where multi-organization or multi-tenant agreements for sharing data would not be feasible.

One of the main limitations with machine learning (ML) based detection methods is the cost of computation. An open-source engineering problem is how to create inferences about high-volume 5G traffic, which will have limitations on MEC nodes due to limited resources. An additional problem is related to adversarial robustness; an intelligent adversary who knows that ML-based detection is in place could attempt to design attack traffic with characteristics defined as "normal" by the rules on the model in order to circumvent detection.[15] This problem of adversarial robustness is the subject of focus in additional research, but has not yet reached a high level of comprehensive resolution.

C. 3GPP Standardized Security Mechanisms

The 3GPP TS 33.501 [3] security architecture provides a complete set of mechanisms for 5G SA deployments [16]. This includes SUCI's (Subscriber Concealed Identifier) method of concealing a subscriber's identity, which mitigates the longstanding IMSI-catcher vulnerability present in prior generations of mobile technology; mutual authentication between end user devices and the network using 5G-AKA and EAP-AKA', and OAuth 2.0 for authentications performing service calls inter-network functions; and mandatory use of TLS 1.2/1.3 to protect service-based interfaces. Signaling traffic between public networks and mobile networks are protected by Security Edge Protection Proxy (SEPP), which ensures confidentiality and integrity.

These mechanisms provide a necessary and non-negotiable minimum baseline. No operator can reasonably say that they have properly secured their 5G deployment without using them. However, they were primarily designed to provide authentication, authorization and confidentiality not to mitigate resource exhaustion. Therefore, a volumetric DDoS attack can flood an AMF with millions of syntactically correct registration requests and successfully complete all of the authentication checks detailed in TS 33.501 but still cause the AMF to exceed its processing capacity. The established mechanisms define a lower limit to DDoS mitigation in 5G and provide no upper limit; to rely solely on them would be a significant operational area of weakness for a carrier.

D. Network Slice Isolation

The next best thing to preventing an attack from reaching its target slice is preventing it from spreading to adjacent slices. Strictly isolating slices through dedicated resource pools, separate orchestration domains, or hard bandwidth guarantees will contain a DDoS attack's blast radius, allowing high-value services to remain safeguarded even if lower-value slices are attacked. If you're an operator with mission-critical slices running alongside consumer broadband slices, then isolation discipline cannot be an option for you; it represents either a localized service disruption or a total system failure.

Cost, however, is the main barrier to complete isolation. To assign physical resources completely to one slice means you no

longer have the resource pooling benefits that make slicing an economically viable option. Literature has addressed many of the different ways to mitigate this problem, including how dynamic resource reallocation triggered by anomaly detection will help; the quarantining process of an attacked slice so that it cannot consume as many of the shared resources as possible; and utilizing trust frameworks to vary inter-slice access based on the historical behavior of slices. However, none have yet produced a completely satisfying solution demonstrating a practical balance between isolation and resource efficiency [6, 10]. This remains an active area of research, and there is no consensus on the best approach to take.

E. CDN and Scrubbing Centers

Dedicated scrubbing centers and content delivery networks are the most advanced ways to mitigate volumetric DDoS attacks, and this has been shown to be true on a large scale for more than ten years. A scrubbing center receives traffic from an attack, removes the attacking packets using signature matching, behavioral analysis, and rate limiting. After doing that, the scrubbing center forwards the remaining non-attacking (cleaned) packets on to the destination. In early 2023, Cloudflare stated that they successfully mitigated a 900.1 Gbps volumetric DDoS attack using their scrubbing infrastructure of 32 locations around the world. Additionally, this architecture has evolved from a hardware-based model toward a fully software-defined model that can be deployed and scaled based on customer needs [11].

The main problem with implementing 5G is low latency. By sending traffic to a scrubbing center and then back again, round trip delays are added that don't work with URLLC service needs. For instance, if an industrial automation slice has a 1ms end-to-end latency budget, it can't go through a scrubbing center that is far away, no matter how well the scrubbing works. Both CDN and scrubbing options are still useful for eMBB and mMTC use cases, where latency is less of an issue, and for keeping operators' Internet-facing surfaces safe. However, they can't be relied on as a one-size-fits-all solution for all 5G service classes.

F. Zero Trust Architecture

The core idea behind Zero Trust Architecture is one that is both straightforward and expansive in scope: "Trust Nothing & Verify Everything." This principle applies to any and all requests, no matter their origin or the network from which they originate. In the 5G ecosystem, this translates to authentication on a per-request basis for API calls between network functions, micro segmentation to ensure that the compromise of one network function does not provide free reign to communicate with all other functions, and real-time monitoring of unusual access behavior.

The analysis of Alnaim on applying the ZTA concept to the 5G virtual networks shows that it is a promising idea that is not ready for practical deployment [6]. Currently, there are no 3GPP specifications to define the security reference architecture based on ZTA for 5G networks. Ericsson has already performed some analyses regarding the ability of existing 5G standards to meet the requirements of the Zero Trust architecture. However, these are preliminary works that cannot be used as recommendations yet. Implementing ZTA would require performing per-request verification in high volumes, which poses non-trivial challenges in terms of the number of computational resources and the tooling needed to perform such actions within cloud-based environments.

V. Open Challenges

A brief analysis of the defensive landscape indicates multiple unresolved challenges. Technically, the most demanding task would likely be inspecting encrypted traffic. TLS version 1.3, required for 5G service-based interfaces according to the 3GPP standard [16], guarantees end-to-end encryption, rendering impossible any attempts by an intermediary network component to inspect the payload of packets. Consequently, the use of machine learning models for detection purposes is limited to metadata and flow statistics – data types significantly inferior to actual packets in their discriminative power. Creating detection solutions that are effective while respecting widespread use of TLS encryption represents an essential research objective within this domain.

The second challenge involves the need for sub-millisecond latency in URLLC use cases. Any end-to-end pipeline for detection and mitigation introducing significant processing delays is unacceptable in these scenarios. The methods discussed in this paper were developed and tested in environments in which minimizing latency was not the primary goal. Modifying these techniques for URLLC applications, or even creating entirely new approaches tailored specifically for URLLC environments, will require substantial research work still to be done.

The third technical issue is the multi-tenancy aspect of network slicing, which introduces challenges of attribution and accountability that transcend the purely technical sphere. In case a DDoS attack is executed via traffic originating within the shared network - say, through a botnet with its command-and-control channel running through an mMTC slice - finding out who to blame and taking measures would require the participation of several parties whose agendas might conflict with each other. This problem has less to do with technologies than with governance and contracting, and none of the methods listed above provide an answer to it.

Lastly, the appearance of the O-RAN architecture allows xApps and rApps from non-operator vendors to be integrated into the RIC layer - another source of potential vulnerabilities since the security attributes of such third-party software products would be hard to validate. O-RAN offers numerous advantages over the proprietary model, but also poses a completely new set of security risks whose interactions with DDoS defenses have hardly been explored at all yet [5]. It appears that this topic will keep researchers busy for some time to come.

Conclusion

In conclusion, this paper has analyzed DDoS threats in 5G networks from both architectural and analytical perspectives. The overarching thesis has been that precisely those unique characteristics which define 5G technology – software-defined

control plane, virtualized network functions, shared slicing capabilities, large-scale IoT connectivity – also set apart DDoS attacks in 5G from their predecessors in 4G, rendering them more complex and potentially more dangerous. A clear understanding of the nature of the threat is imperative if a response to 5G-specific DDoS risks is to be effective.

The comparison of six selected mitigation techniques indicates that none of them is comprehensive enough to provide adequate protection against all 5G attack scenarios and applications. SDN-based mitigations and machine learning/ artificial intelligence-based detections demonstrate the highest degree of flexibility and versatility, but suffer from the lack of research into issues related to their deployment. Standard 3GPP procedures are necessary, yet insufficient for providing any meaningful protection; at best, they create an absolute minimum standard of security. Network slicing allows limiting the impact of any attacks, yet the contradiction between high security and high performance cannot be easily resolved. CDNs and scrubbing centers work but are impossible to deploy for URLLC traffic types.

This means that the defense against 5G DDoS attacks would have to be multi-faceted and tailored to the particular application and latency requirements of each slice. Future research directions include: light-weight machine learning inference engines appropriate for MEC, isolation methods that do not undermine the efficiency gains from slicing, standardized ZTA security architecture for 5G Standalone systems, and O-RAN security solutions mindful of the software supply chain issues in 5G. These are all challenging problems, but they are all tractable and becoming more and more relevant with increasing 5G adoption across the globe.

REREFENCES

[1] 3GPP, "System Architecture for the 5G System (5GS)," Technical Specification TS 23.501, Release 17, 3rd Generation Partnership Project, 2022.

[2] Cloudflare, "DDoS Threat Report for 2025 Q4," Cloudflare Radar, Feb. 2026. https://radar.cloudflare.com/reports/ddos-2025-q4.

[3] 3GPP, "Security Architecture and Procedures for 5G System," Technical Specification TS 33.501, Release 17, 3rd Generation Partnership Project, 2022.

[4] I. Ahmad, S. Shahabuddin, T. Kumar, J. Okwuibe, A. Gurtov, and M. Ylianttila, "Security for 5G and beyond," IEEE Communications Surveys & Tutorials, vol. 21, no. 4, pp. 3682–3722, Fourthquarter 2019. doi: 10.1109/COMST.2019.2916180.

[5] "Investigating Security Vulnerabilities in 5G Control and User Planes: Attack Patterns and Protection Strategies," Journal of Cybersecurity and Privacy, vol. 6, no. 1, p. 37, Feb. 2026. doi: 10.3390/jcp6010037.

[6] A. K. Alnaim, "Securing 5G virtual networks: a critical analysis of SDN, NFV, and network slicing security," International Journal of Information Security, vol. 23, pp. 1–21, 2024. doi: 10.1007/s10207-024-00900-5.

[7] T. Taleb, K. Samdanis, B. Mada, H. Flinck, S. Dutta, and D. Sabella, "On multi-access edge computing: A survey of the emerging 5G network edge cloud architecture and orchestration," IEEE Communications Surveys & Tutorials, vol. 19, no. 3, pp. 1657–1681, Thirdquarter 2017. doi: 10.1109/COMST.2017.2705720.

[8] M. A. Setitra, M. Fan, I. Benkhaddra, and Z. E. A. Bensalem, "DoS/DDoS attacks in software defined networks: Current situation, challenges and future directions," Computer Communications, vol. 213, 2024. doi: 10.1016/j.comcom.2023.11.012.

[9] "Modern DDoS Threats and Countermeasures: Insights into Emerging Attacks and Detection Strategies," arXiv preprint arXiv:2502.19996, Feb. 2025.

[10] M. S. Khan, B. Farzaneh, N. Shahriar, and M. M. Hasan, "DoS/DDoS Attack Dataset of 5G Network Slicing," IEEE DataPort, Sep. 2023. doi: 10.21227/32k1-dr12.

[11] "5G Network Slicing: Security Challenges, Attack Vectors, and Mitigation Approaches," NCBI/PMC, 2025. [Online]. Available: https://pmc.ncbi.nlm.nih.gov/articles/PMC12251764/.

[12] M. Sheibani, S. Konur, I. Awan, and A. Qureshi, "A multi-layered defence strategy against DDoS attacks in SDN/NFV-based 5G mobile networks," Electronics, vol. 13, no. 8, p. 1515, Apr. 2024. doi: 10.3390/electronics13081515.

[13] C. L. Chen, Y. J. Rao, and C. H. Sung, "Mitigating DDoS attacks in private 5G networks using deep learning," Mobile Networks and Applications, 2025. doi: 10.1007/s11036-025-02475-4.

[14] V. Pourahmadi, H. A. Alameddine, M. A. Salahuddin, and R. Boutaba, "Spotting anomalies at the edge: Outlier exposure-based cross-silo federated learning for DDoS detection," IEEE Transactions on Dependable and Secure Computing, vol. 20, no. 5, pp. 4002–4015, Sep.–Oct. 2023. doi: 10.1109/TDSC.2022.3197118.

[15] A. Apostu, S. Gheorghe, A. Hîji, N. Cleju, A. Pătraşcu, C. Rusu, R. Ionescu, and P. Irofti, "Detecting and Mitigating DDoS Attacks with AI: A Survey," arXiv preprint arXiv:2503.17867, Mar. 2025.

[16] A10 Networks, "DDoS Weapons Report," A10 Networks, 2024. https://www.a10networks.com/resources/reports/ddos-weapons-report/.