AquaSonic: Acoustic Manipulation of Underwater Data Center Operations and Resource Management

Jennifer Sheldon*, Weidong Zhu*, Adnan Abdullah*, Sri Hrushikesh Varma Bhupathiraju*,
Takeshi Sugawara§, Kevin R. B. Butler*, Md Jahidul Islam*, Sara Rampazzi*
*University of Florida; §The University of Electro-Communications

***Abstract*—Underwater data centers (UDCs) hold promise as next-generation data storage due to their energy efficiency and environmental sustainability benefits. While the natural cooling properties of water save power, the isolated aquatic environment and long-range sound propagation characteristics in water create unique vulnerabilities which differ from those of on-land data centers. Our research discovers the unique vulnerabilities of fault-tolerant storage devices, resource allocation software, and distributed file systems to acoustic injection attacks in UDCs. With a realistic testbed approximating UDC server operations, we empirically characterize the capabilities of acoustic injection underwater and find that an attacker can reduce fault-tolerant RAID 5 storage system throughput by 17% up to 100%. Our closed-water analyses reveal that an attacker can (i) cause unresponsiveness and automatic node removal in a distributed filesystem with only 2.4 minutes of sustained acoustic injection, (ii) induce a distributed database's latency to increase by up to 92.7% to reduce system reliability, and (iii) induce load-balance managers to redirect up to 74% of resources to a target server to cause overload or force resource colocation. Furthermore, we perform open-water experiments in a lake and find that an attacker can cause controlled throughput degradation at the maximum allowable distance of 6.35 m using a commercial speaker. We also investigate and discuss the effectiveness of standard defenses against acoustic injection attacks. Finally, we formulate a novel machine learning-based detection system that reaches 0% False Positive Rate and 98.2% True Positive Rate trained on our dataset of profiled hard disk drives under 30-second FIO benchmark execution. With this work, we aim to help manufacturers proactively protect UDCs against acoustic injection attacks and ensure the security of subsea computing infrastructures.**

1. Introduction

Data centers play a crucial role in handling and storing vast amounts of data to serve the requirements of different applications owned by private individuals, enterprises, and government institutions. With the increasing interest in environmental sustainability and the surging data center market due to a recent spike in demand for AI [1] and cloud computing [2], companies are actively seeking alternative methods to improve energy efficiency and to reduce operating costs. To this end, Microsoft [3], Subsea Cloud [4], and Offshore Oil Engineering Company [5], among others, have already deployed successful prototypes and released in the market underwater data centers (UDCs). Typical UDCs have submerged structures with metal pressurized vessels containing server racks and filled with nitrogen gas to prevent corrosion [3], [6], which have demonstrated significant advantages due to the natural cooling properties of water, space efficiency, and renewable energy integration [7].

Several attack vectors have been studied for in-land data centers, but they generally involve installing malware [8], [9], [10], target data center networks [11], [12], [13], rely on hardware colocation of virtual machines (VMs), for eavesdropping [14], [15], [16], or require the attacker to gain physical access and tamper with components inside the data center [17], [18], [19]. More recently, Sheldon et al. [20] built upon previous works, which explored in-air Denial-of-Service (DoS) attacks on HDDs [21], [22], by showing that strong sound injection attacks at a resonant frequency of hard disk drives (HDDs) deployed in submerged enclosures can cause throughput reduction and application crashing. While these previous works demonstrated the possibility of DoS attacks on a single disk in air and water environments, it remains unexplored if such acoustic injection attacks can be used to affect critical data center operations and resource management necessary to ensure the reliability and efficiency of such infrastructure.

In this paper, we perform modulated acoustic injection attacks in a controlled testbed and real-world open-water scenarios to characterize and quantify an attacker's ability to manipulate complex operations within UDCs leveraging the capability of acoustic attacks to influence multiple storage devices simultaneously. Specifically, we test the resilience of fault-tolerant storage techniques such as RAID and investigate attackers' ability to gain fine-grained control over geo-distributed database performance and latency (e.g., CockroachDB [23]), distributed filesystem node allocation and replication (e.g., Hadoop Distributed Filesystems (HDFS) [24]), and resource allocation (e.g., OpenNebula [25]). We also show how modulated injection can control the latency of real-world data center workloads such as Microsoft SNIA [26] at different volumes of injected sound. Our evaluation of full-HDD and Hybrid Solid State Drive (SSD) cache - HDD [27], [28], [29] architectures deployed in current data centers shows that, even if the SSD cache is almost immune to acoustic injection, the attack still

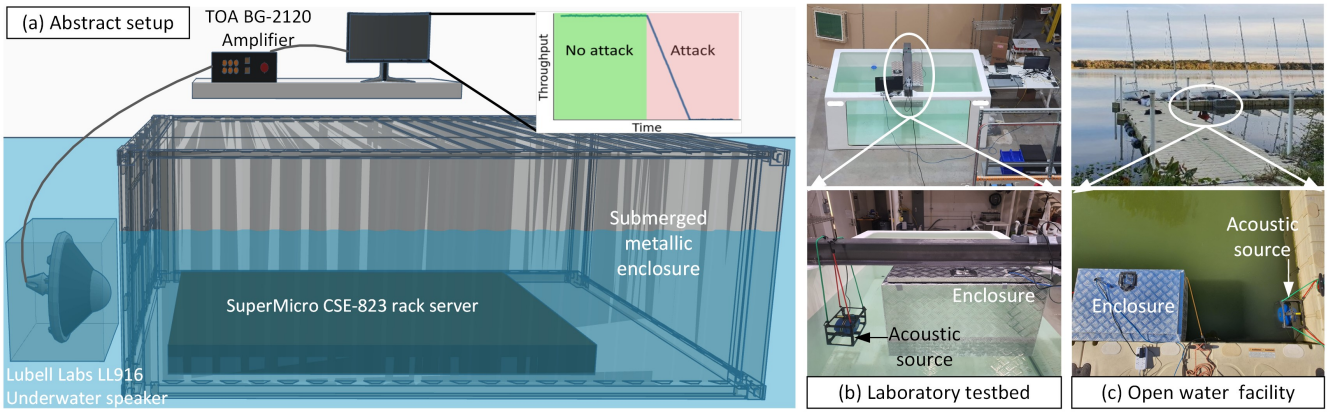

Figure 1: (a) An overview of our experimental setup; a rack server in a RAID 5 configuration is placed in a $0.9 \times 0.6 \times 0.5$ m metal enclosure, while acoustic attacks are carried out using an underwater speaker. (b) The indoor experiments are conducted in a $1.2 \times 3.0 \times 1.5$ m water tank. (c) The studies are extended at an open-water testing facility.

increases latency. For all our tested cache sizes (0.5, 1, 1.5, and 2 GB), we found that the random-write workload latency rose from the 1 to 200 ms range to the 200 to 800 ms range.

In our load manipulation analysis, we demonstrate how acoustic injection induces a resource manager (OpenNebula) to assign a minimum of 58% to a maximum of 74% of VMs to a target server during injection. Such assignment manipulation can be used to force benign user applications to be colocated with malicious tenants in compromised servers or to indirectly overload specific servers. In extreme cases, we also cause virtual machines to run into permanent deadlock after the acoustic injection has stopped, which can stop critical processes and corrupt sensitive data. Finally, we evaluate our attack in open-water scenario and demonstrate that the attack can achieve 61% throughput drop with the speaker placed 6.35 m away from the enclosure[1].

In light of our characterization of the attacker's capabilities, we discuss how simple defenses, such as the use of sound absorbing materials, should be carefully considered to avoid dangerous heating which can impact the performance of the data center servers in submerged enclosures. To address this, we develop and evaluate a novel proof-of-concept machine learning (ML) based defense that detects whether acoustic injection has occurred based on throughput analysis along with information on acoustic vibration patterns.

In summary, this paper includes the following contributions:

- We characterize submerged data center storage devices' vulnerability to acoustic injection. We perform real-world testing in a simplified testbed and open-water scenarios, using a server enclosed in a submerged metal structure. Our evaluation shows remote throughput manipulation more than 6 meters away from the enclosure. We also build a preliminary simulation model for evaluating different subsea UDC structures.
- We demonstrate how an adversarial attacker can affect the performance and reliability of distributed databases and filesystems in UDCs by performing acoustic injection on CockroachDB [23] and a server running the DFSIO [30] file access benchmark on HDFS [24]. We increase the CockroachDB's latency up to 92.7% and block Hadoop from accessing data within 144 seconds of injection.
- Furthermore, we demonstrate how attackers can manipulate data center resource allocators, such as OpenNebula, and force up to 74% resource reassignment while circumventing VM placement policy-based defenses.
- We investigate the effectiveness of standard defenses against acoustic attacks, such as using sound-absorbing material, SSD-hybrid architectures, active noise cancellation, feedback controllers, and sensor fusion.
- We propose and evaluate a novel ML-based defense that models patterns of storage device throughput and identifies attacks based on anomaly estimation across multiple spatially close storage devices. Comprehensive evaluations in our real-world testbed conditions leveraging the FIO sequential write benchmark show that the proposed defense achieves 0% False Positive Rate and 98.2% True Positive Rate in detecting simultaneous degradation caused by the attack.

Overall, our analysis begins to unveil hardware and software vulnerabilities that are unique to submerged environments, while revealing new design flaws in traditional data centers fault tolerance storage devices and resource management systems. This work aims to help manufacturers and designers of UDC infrastructures promptly address those security risks before they become widespread.

2. Background

2.1. Data Center Architectures

Users, companies, and organizations are shifting their data and business to the cloud with unprecedented speed, especially since the COVID-19 pandemic which mandated

1. The testing distance was limited to 6.35 m by the dock used to anchor the enclosure, not by limitations of the setup (See Figure 1).

remote work and study [31]. This surge fueled the ongoing expansion of data center services, which continues to this day. For example, at the time of writing, data center construction is projected to reach $49 billion by 2030 with a power consumption estimated at 35 gigawatts only in the United States [32]. To satisfy the high demand for data storage and computing resources, data centers comprise high-speed processors, servers, network switches and routers, and large-scale storage systems [33].

In addition, data center architectures deploy fault tolerance techniques, resource allocation, load and workload balancing processes, and storage systems management tools to ensure the reliability, performance, and efficiency of the infrastructure when handling data. In this work, we focus on understanding and evaluating how such supporting resources are manipulated by acoustic attacks affecting storage devices in the context of UDC deployment.

Storage Devices in Data Center. Data centers providers, such as Alibaba Pangu [34], Microsoft Azure [35], and Meta [36], use multiple types of storage devices [37], including HDDs and SSDs, to satisfy their demands on storage capacity and performance. Although SSDs have advantages in performance and reliability [38], HDDs remain the main components of current data center storage systems due to the following reasons. First, SSDs suffer of limited lifetime due to finite program/erase (P/E) cycles and high *infant mortality* rate [39]. Second, SSDs rely on garbage collection to deal with underlying NAND flash memory [40], [41] that prohibits in-place updates, degrading the performance [41], [42] and introducing unpredictability [43], [44]. Finally, SSDs are still expensive [29], [45], [46] compared to their equivalent HDDs. Therefore, HDDs and SSDs co-exist in modern data centers, with SSDs mainly used as cache [29], [35], [36] of HDDs to boost the performance of the storage system. In this work, we thus focus primarily on full-HDD and Hybrid SSD cache architectures to evaluate the impact of acoustic attacks on UDCs.

Redundant Array of Independent Disks (RAID). Typically data centers deploy storage devices with RAID technology [47] to enable data processing in multiple hard disks simultaneously while providing fault tolerance. There are six types of configurations in a RAID system. RAID 0 evenly divides storage space and distributes incoming data in chunks, enabling speed up by simultaneous disk access. However, RAID 0 offers no fault tolerance, and the failure of one storage device will lead to data loss. Therefore, RAID 2 to RAID 6 provides fault tolerance by adding redundancy and using erasure coding techniques. Specifically, RAID 1 relies on mirroring data – adding replicas – across multiple storage devices, whereas RAID 2 to RAID 6 use erasure code to provide fault tolerance without significant storage overhead. While RAID 2 and RAID 3 are rarely used in practical applications because they chunk data in byte granularity, making processing I/O requests in parallel difficult and incurring poor performance, RAID 5 and 6 are commonly used in enterprise storage systems. In this work, we consider RAID 5 as the implementation in our experimental setting because widely used in the storage system of data centers [40], [48], [49].

Resource Allocation Techniques. Data centers allocate computing and storage resources on demand. Unlike provisioning resources on a single machine, a data center allows a flexible and scalable combination of system resources, such as memory, CPU cores, and storage devices. For example, Infrastructure as a Service (IaaS) [50] allows users to purchase a specific amount of computing power and storage space to fulfill their demands; storage space in the RAID device is assigned upon request and disaggregated on release [48]. In addition, to achieve high-performance utilization of system resources, data centers optimize the use of their resources by balancing the workload between servers [51]. Workloads are distributed to servers based on varying optimization parameters. For example, OpenNebula [25] assigns virtual machines to hosts based on varying optimization parameters such as number of hosts used, available server resources, or custom algorithms [52]. Through our evaluation of OpenNebula in Section 5.3, we demonstrate how attackers can manipulate resource allocation in data centers to overload servers or force colocation with malicious virtual machines to perform potential attacks such as eavesdropping [14], [15], [16].

2.2. Acoustic Injection Attacks

Acoustic injection attacks involve the use of sound to manipulate the behavior of a target system. These attacks usually leverage the *resonant frequency*, meaning the natural frequency at which a solid structure oscillates. Several parameters can generate resonance frequencies, including the elasticity or stiffness of the material from which a system or component is made, the physical dimensions, and the object's mass distribution. Additionally, complex structures may exhibit multiple resonance frequencies corresponding to different vibration modes. Thus acoustic attacks consist of transmitting acoustic waves at a frequency that matches the resonant frequency, to efficiently convert acoustic waves into physical vibrations of the target system [53].

Researchers have explored acoustic injection attacks against a variety of sensors, including (i) cameras [54], [55] to induce obstacle misdetection in autonomous vehicles, (ii) accelerometers [56], gyroscopes [57], [58] and inertial sensors [59], [60] to alter drone locations, and (iii) microphones to stealthily deliver voice commands [61]. In particular, Blue Note [21] used acoustic injection at the resonant frequency of an HDD to vibrate its actuator arm outside the disk track's limits, causing DoS in laptops and security cameras.

More recently, Sheldon et al. [20] in a position work showed a DoS attack against a single HDD located in a closed submerged container made of different materials (plastic and aluminum). The attack is successful in short distances (up to 25 cm from the container), causing operating system crashes and I/O timeout errors. Although this preliminary work pioneered how sound waves in water can induce mechanical vibrations that propagate through solid

structures (e.g., the target hard disk drive and container), it was limited in reproducing Blue Note in underwater settings.

Unlike these previous works, our analysis focuses on leveraging the vulnerability of storage devices to acoustic injection to subtly manipulate complex operations and processes (e.g., resource allocation, fault tolerance techniques) vital for maintaining the reliability of data centers.

2.3. Acoustic Signal Propagation in Water

Sound waves propagate in water at a much higher velocity over long distances than in air. This is due to the higher density of water, which allows for efficient transmission with minimal loss of energy. In general, the speed of sound is around 1,480 m/s, which is about four times faster than in the air [62]. More specifically, sound propagation in water depends on many factors, such as temperature, salinity, depth, seabed relief, currents, and surrounding pressure [63]. In shallow water, sound waves reflect from the surface and the floor of the water body, and scatter from suspending particles and bubbles, resulting in multi-path propagation. In deep water instead, pressure and temperature vary with depth which creates multiple layers in the medium and sound bends while propagating through such layers [64], [65]. High-frequency energy is scattered and absorbed more rapidly by a water medium that allows low-frequency sound to travel longer distances. Additionally, for a given sound frequency, sea water shows typically a higher absorption coefficient than fresh water due to its higher concentration of dissolved minerals [66]. All the experiments in this work are conducted in fresh water scenarios where the acoustic source is kept at a certain depth from the surface and the ambient temperature is maintained.

The sound pressure level (SPL) is the typical measure of the intensity or loudness of a sound wave and is expressed in decibels (dB) as the product of medium density, wave velocity, and particle velocity in a specific medium [67]. Sound pressure levels in water are about 60 dB higher than the equivalent SPL in air for the same sound source. This means that a sound source producing 100 dB SPL in air, is approximately equivalent to 160 dB SPL sound in water. In all the experiments in this work, we determine the sound pressure by using a hydrophone and empirically measure the sound pressure generated by the sound source (e.g., our underwater speaker) at increasing distances (see Figure 2).

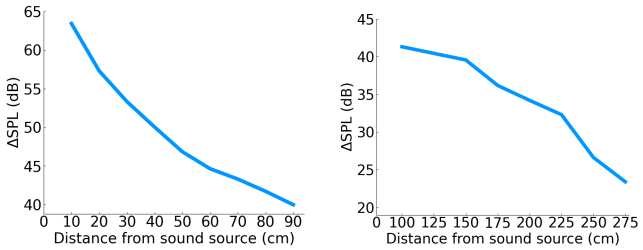

(a) Laboratory testbed (b) Open water

Figure 2: Sound pressure attenuation at increasing distances are shown for closed-water laboratory testbed and open-water scenarios; we find similar trends in both scenarios.

3. Attack Overview and Threat Model

3.1. Threat Model

In this work, we consider an adversary who aims to achieve high-level control over underwater data center infrastructure operations by exploiting the vulnerability of the storage systems to acoustic injection. Unlike the previous works [20], [21], [22] focused on DoS attacks, this work aims to characterize the extent and capability of an attacker to induce at far distances sophisticated and subtle manipulations of data center operations such as redirecting workloads, altering resource allocation, and achieve fine-grade control over latency and throughput of critical distributed applications. We reveal how these manipulations can be pursued without evident and abrupt changes in the process executions, causing storage devices to be automatically removed, forcing automatic node redirection, and dangerous colocation of resources which can overload specific servers.

Attacker Knowledge and Assumptions. We assume that the adversary has no direct access to the underwater infrastructure and cannot tamper with any hardware, software, or network communication. Meanwhile, the adversary can observe the outcome of the acoustic attack by monitoring benign application instances running on the infrastructure. A less sophisticated attacker without direct access to application instances can also analyze effective signal injection frequencies and volumes by performing small-scale evaluations on enterprise HDDs as we show in this work, or simulating the target's structure using physics modeling simulators. The attacker can also perform a simpler attack by emitting sound waves and continuously sweeping a range of frequencies.

We assume that the attacker owns an underwater speaker capable of generating and controlling the volume and frequency of the acoustic waves. The attacker is capable of aiming the speaker at the data center location to perform the attack for the desired time span. This can be achieved by mounting the speaker on a rigid structure connected to a boat or using more sophisticated settings, such as a remotely controlled underwater robot [68]. The attacker can also make use of directional speakers such as speaker arrays or Long Range Acoustic Device (LRAD) devices [69] to transmit focused beams of sounds with a confined range to target specific structure parts.

In order to perform the attack, the adversary should identify the susceptible frequency ranges. This is possible, by observing delays in benign application write requests caused by brief sound injections and different frequencies or by studying similar storage devices and their resonance frequencies as shown in previous work [20], [21].

Attack Scenario. Adversaries can launch the attack several meters away from the underwater infrastructure, depending

on their equipment and the susceptibility of the victim system. We first characterize the vulnerability in a controlled scenario using a laboratory testbed in Section 4. In the open-water scenario described in Section 5.6, we achieve successful attacks at 6.35 meters away using a commercial speaker connected to an amplifier. A well-funded adversary with powerful speakers (e.g., military-grade equipment) can potentially reach further distances as we explored in our simulation scenario in Section 5.7.

3.2. Attack Overview

As shown in Figure 1, the attacker uses an underwater speaker to generate modulated sound waves in the form:

$$s(t) = A \cdot \cos(\omega \cdot t) \tag{1}$$

where A is the amplitude of the sound wave corresponding to the volume, ω is the angular frequency of the transmitted sound, and t is time. We also define the amplitude of the signal as a decibel sound pressure level (SPL) above the noise level using the formula: $\Delta SPL = SPL_m - SPL_n$. Here, SPL_m is the SPL (in dB) measured in the testing environment with a hydrophone, and SPL_n is the environmental noise SPL measured when no sound is being emitted. A higher ΔSPL is associated with a louder volume.

As described in Section 2.2, transmitting acoustic waves at resonant frequencies with sufficient volume can cause mechanical vibrations in the internal components of HDDs (e.g. read/write head and platter), preventing reading and writing operations and consequently causing application crashes. We apply amplitude modulation to induce controlled changes in the behavior of a victim system composed of multiple storage devices in full-HDD and hybrid SDD-HDD architectures as described in Section 2.1.

Indoor Testbed Specifications. To approximate a simplified underwater infrastructure and perform our vulnerability characterization, we use the indoor testbed depicted in Figure 1. This testbed consists of a $1.2 \times 3.0 \times 1.5$ m water tank filled with fresh water. An aluminum metal enclosure of $0.9 \times 0.6 \times 0.5$ m is used to emulate a real-world data center vessel while a rack server (or rack-mounted server) is used as the target system, as UDCs infrastructures are composed of server racks [3], [6].

3.3. Theoretical Analysis

The sound-induced vibrations in submerged enclosures depend on three main physical phenomena: the propagation of sound through fluids, the force applied by fluids on solid boundaries, and the propagation of mechanical vibrations through solids.

In our attack model, sound waves travel through a body of water to reach the enclosure and then propagate to the victim server's solid structure containing the hard disk drives. Sound propagation in fluids is generally modelled using the equations [70]:

$$\begin{cases} \frac{1}{\rho}\nabla^2 p_t(x) - \frac{k^2}{\rho} p_t(x) = 0 \\ k = \frac{\omega}{c} \end{cases} \tag{2}$$

Where ρ is the density of the analyzed fluid, $p_t(x)$ is the total pressure in the fluid at a location x assuming the acoustic source at position 0, ω is the angular frequency of the source sound, and c is the speed of sound, which depends on temperature, salinity, and depth of the fluid.

Sound waves attenuate with distance from the source, and this attenuation can be modeled as:

$$A = A_0 e^{-\alpha x} \tag{3}$$

where A typically is represented by the root mean square amplitude (RMS) of the wave at distance x, A_0 is the RMS when x is zero, and α is the attenuation coefficient of the fluid, which depends on the sound frequency.

The force applied by a sound propagating through a fluid and encountering a solid can be approximated using the equations [71]:

$$\begin{cases} \frac{-\boldsymbol{n}}{\rho}\nabla p_t = -\boldsymbol{n} \cdot (\boldsymbol{u_{tt}}) \\ F_A = p_t \boldsymbol{n} \end{cases} \tag{4}$$

Where ρ is the density of the fluid, $\boldsymbol{n}$ is the surface normal of the solid, $\boldsymbol{u_{tt}}$ is the acceleration vector of the solid, p_t is the total acoustic pressure, and F_A is the fluid load on the solid boundary which depends on the distance of the sound source and its attenuation.

The force applied by the sound pressure induces vibrations in the submerged metal enclosure, which can be expressed as [72]:

$$\rho_s \frac{\delta^2 \boldsymbol{u}}{\delta t^2} = \boldsymbol{F_v} \nabla_X \cdot P \tag{5}$$

Where ρ_s is the density of the solid, $\boldsymbol{u}$ is the displacement vector of each point in the solid material, $\boldsymbol{F_v}$ is the vector force per unit volume applied to the solid. Here, the force is calculated upon F_A from Eq. 4 at the fluid-solid boundary. Finally, the propagation of mechanical vibrations at the interface between solids of different densities, such as the vessel structure, the internal server racks, and hard disk configurations, depends on complex interactions based on the specific materials of each component, boundary conditions, and other factors such as reflection, transmission, and mode conversion [73]. When mechanical vibrations caused by the injection encounter an interface between two different materials, part of the wave is reflected back into the original material, while part of it is transmitted into the second material [74]. The coefficients of reflection and transmission can be calculated using the acoustic impedances of the two materials. This can be represented by the following simplified equations:

$$\begin{cases} R = \frac{Z_2 - Z_1}{Z_2 + Z_1} \\ T = 1 + R = \frac{2Z_2}{Z_2 + Z_1} \end{cases} \tag{6}$$

Where (R) and (T) are the reflection and transmission coefficients, respectively, and (Z_1) and (Z_2) are the acoustic impedances of the first and second materials. In addition to reflection and transmission, mode conversion can occur at the interface. For instance, an incident longitudinal wave can generate reflected longitudinal and transverse waves, and transmitted longitudinal and transverse waves. The same applies to an incident transverse wave based on the angle of incidence of the wave, the acoustic impedances, and the frequency of the wave. At the resonance frequency ranges, meaning where the frequency of the sound wave matches the natural oscillating frequency of the solids in contact with each other, the mechanical impedance will be lower, meaning less force will be needed to propagate the wave at the target object (e.g., the storage system) and cause vibrations at a given velocity and intensity directly proportional to the sound pressure level.

3.4. Sound-induced Vibrations at the Disk

Based on our theoretical analysis, we verify the vibration propagation in a storage device located in our server enclosed in the submerged metal structure in our indoor testbed. To measure the resulting vibrations caused by the sound, we extract the Position Error Signal (PES) as described in previous work [75]. The PES measures the deviation of the read/write head from the center of the track, thus we placed a 500 GB Seagate Barracuda HDD housed in a SuperMicro CSE-823 rack server in our indoor testbed at 6 cm from the edge of the submerged enclosure. To evaluate the deviation of the read/write head, we inject a tone at the HDD resonance frequency (5.1 kHz) at increasing volumes (46 to 64 dB ΔSPL). We leverage the Servo Batch Test [76] in the Seagate terminal command set to get the PES data from the HDD where each test has 296 revolutions. Figure 3 shows that the average displacement ratio increases from 0% to 83% at increasing volumes. This verifies that sound tones at the resonance frequency induce vibrations in the read-write head and platter of the disks by vibration propagation, which is proportional to the acoustic pressure (intensity of the sound, volume) generated by the injection. In this work, we show how an attacker can control the degree of induced effect in applications by varying the injected sound volume.

4. Vulnerability Characterization

In this section, we determine whether an attacker can exploit the resonant frequency of a server composed of multiple storage devices in RAID 5 configuration in a submerged metal enclosure. We also evaluate the attacker's ability to maintain fine-grained control over throughput and latency to perform subtle attacks. Through this evaluation, we then quantify the attacker's limitations and capabilities in attacking a data center high-level operations using acoustic injection in Section 5.

Experimental Setup. We perform our characterization analysis using our indoor testbed described in Section 3.2. The sound source is a Lubell Labs LL916 speaker [77] used

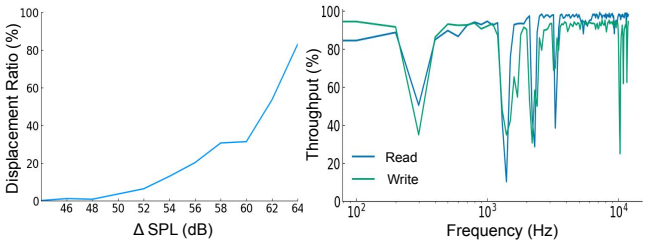

Figure 3: (Left) PES displacement ratio at different $\Delta SPLs$. (Right) RAID 5 throughput as a percentage of the baseline average at increasing frequencies based on a frequency sweep from 100 Hz to 12 kHz at 100 Hz intervals.

for commercial applications (e.g., delivering verbal instructions to divers and swimmers). As target system located in the submerged metal enclosure, we deploy a SuperMicro CSE-823 rack server [78] running Ubuntu 22.04 with 4 Seagate Exos 7E2 1TB SATA enterprise HDDs [79] used in datacenters in a RAID 5 full-HDD configuration and an Intel D3-S4510 Series application SSD [80]. Hybrid storage architectures are explored in Section 5.5.

4.1. Resonant Frequency Identification

As the first step in our attack characterization, we determine whether attackers can find and exploit the resonant frequencies of our target server by observing the response to a simple sound wave at constant volume. To accomplish this goal, we place the running server in the submerged metal enclosure such that the front of the server is 3 cm away from the enclosure. We then inject sound waves at 150 dB SPL (equal to 34 dB ΔSPL) with frequencies ranging from 100 Hz to 12 kHz with 100 Hz increments every 5 seconds. While playing these tones, we run FIO, an I/O tester tool [81], with sequential read and write throughput benchmarks to determine which frequencies, if any, can cause measurable throughput degradation [20]. We run both the read and write sweeps three times and classify an average performance decrease of more than 20%.

As shown in Figure 3, the measured RAID 5 throughput drops at varying frequencies. Per our theoretical analysis, this is likely because the components inside the four HDDs in the RAID configuration and server structure have varying resonant frequencies (around 2.0, 3.7, 5.1–5.3, and 8.9 kHz, see Figure 3) with higher frequencies hitting harmonics of the resonant frequencies. The consistent throughput degradation between 5.1–5.3 kHz is a good attack target, and we use 5.1 kHz for the rest of our experiments in this work.

4.2. Controlled Injection

To demonstrate how an adversary can control the severity of throughput degradation to perform a more subtle attack, we place the sound source speaker at a fixed distance of 6 cm from the outer surface of the submerged enclosure containing the target server. We then play sounds at the

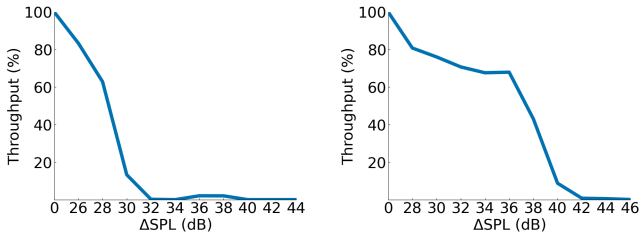

(a) Laboratory testbed (b) Open water

Figure 4: RAID 5 write throughput at increasing sound pressure from the baseline environmental noise (116 dB SPL for the testbed, and 114 dB SPL for the open water scenario) at 5.1 kHz injection frequency. The distance between sound source-enclosure was set to 6 cm for our laboratory testbed and 30 cm for open water scenario.

resonant frequency of 5.1 kHz identified in Section 4.1 with 2 dB SPL increment starting from 26 dB ΔSPL. During these acoustic injections, we record the average RAID 5 throughput of three 30-second runs of the FIO sequential write benchmark. As shown in Figure 4, the attacker can drop the throughput between 17 and 100% by varying the acoustic injection volume between 26 and 32 dB above the noise level (116 dB SPL in our indoor testbed scenario).

Thus, the attacker can subtly control the throughput of storage devices in the target server by varying the injection volume A (See Eq. 1). By associating the volume with sound source-to-target distance using our empirical results to approximate the attenuation constant of water as shown in Figure 2, we also determine that the attacker can control the RAID 5 throughput by varying the distance from which they inject a constant volume signal.

We validate our analysis as well in our open-water scenario. As depicted in Figure 4, the open water scenario required a higher injection volume to reach similar throughput degradation. This is likely because, due to the absence of a fixed anchorage in our open water setup, we weighted the metal enclosure using bags of sand to reach the submerged level required. We believe this procedure may have influenced the resulting vibrations in the enclosure.

4.3. Acoustic Injection Points

As described in Section 3.3, sound propagates through mechanical vibration in the rack server and reaches the storage devices. To understand whether this allows the attacker to successfully degrade RAID 5 performance from different injection positions, we measure the throughput of the RAID 5 configuration when running the FIO sequential write benchmark with the sound source placed in different locations around the enclosure containing the target server (see Figure 5). For each position (described as Locations 1–4 in Figure 5), we use a fixed volume (∼30 dB ΔSPL) and frequency (∼5.1 kHz, as identified in Section 4.1). As vibrations propagate in the entire structure for their nature, we see that the attacker can cause measurable throughput

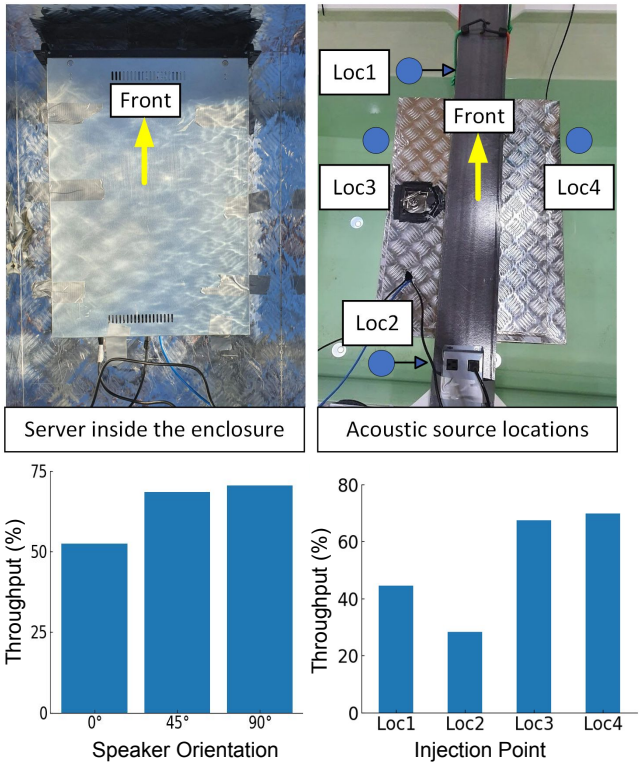

Figure 5: (Top) Four sound injection points relative to the victim enclosure. (bottom-left) FIO throughput with the speaker positioned at different angles with respect to the enclosure. (bottom-right) Normalized RAID 5 throughput during the injection performed at the four injection points.

loss at multiple sound source injection points (including the back of the rack server, far away from the storage device locations in the front). While the drop is more severe in some locations than others, the attacker can compensate for this by raising the volume of the injection (see Section 4.2). This reveals how the attacker is not limited to one injection location to pursue the attack. For the rest of the experiments in this work, we inject sound at Location 1, the front of the rack server, to evaluate the consequences of acoustic injection in a suboptimal location.

4.4. Speaker Orientation

To understand if the orientation of the attacker's speaker affects throughput, we measure the RAID 5 throughput with the speaker turned at different angles with respect to the target enclosure (with 0° representing the speaker aimed towards and 90° representing the speaker oriented parallel to the target enclosure). In our testbed, we place the speaker 30 cm away from the enclosure to allow for rotation and play the 5.1 kHz tone at 40 dB ΔSPL. We find that the attack is approximately 32% less effective at 45° and 34% less effective at the 90° angle than in the direct attack case (See Figure 5). This occurs because the SPL is not uniform at different angles due to the directivity of the speaker,

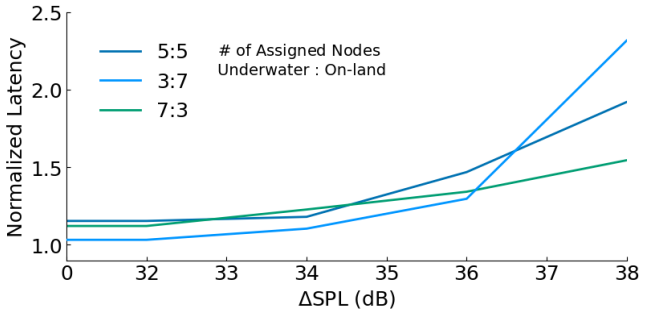

Figure 6: Normalized latency for three node assignment configurations when running TPC-C benchmark [83] on CockroachDB at increasing sound pressure. The latency increases with the amplitude of the sound, even when the majority of the nodes are assigned to the on-land server.

but the lower volume at sub-optimal angles can still cause throughput degradation.

5. Impact on Critical Operations

In light of the findings of our characterization analysis, we evaluate the manipulation capabilities of our acoustic injection on popular data center management software and distributed systems. For this analysis, we use the same setup described in Section 4. In addition, for experiments requiring a second server, we use a PowerEdge R610 rack server [82] placed far from the sound source and submerged enclosure. This "on-land" server acts as an unaffected resource for evaluation of data center management software behavior which uses multiple servers such as distributed databases, distributed filesystems, and resource allocation managers.

5.1. Latency Control on Distributed Databases

Distributed databases are adopted as a widespread solution to address the need for scalability and high availability, such as in streaming services, and also provide fault tolerance in data centers [84], [85]. Compromising distributed databases can lead to service outages [86] and affect the storage of replicas, which in turn can severely degrade the fault tolerance capability of the infrastructure.

We chose CockroachDB [23] as our target to demonstrate the efficacy of our underwater acoustic injection to manipulate real-world distributed database reliability in terms of latency control. CockroachDB is a popular, commercially available, and scalable geo-distributed database for high-performance and data processing that has been adopted by companies such as Netflix [86] and SpaceX [87].

Experimental Setup. We deploy CockroachDB on our aforementioned testbed, which includes two servers, one in the underwater enclosure with RAID 5 configuration while the other on land, outside the influence of underwater acoustic injection. We then consider three different configurations of 10 nodes. In the first configuration, 5 nodes are assigned

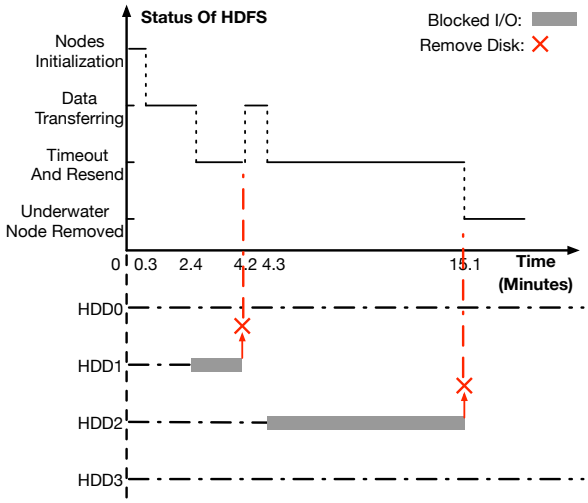

Figure 7: The workflow of HDFS with the corresponding liveness of the four HDDs when running the DFSIO benchmark during an acoustic injection at 36 dB ΔSPL.

to the underwater server and the on-land server, respectively. In the second configuration, only 3 nodes are assigned to the underwater server, while in the third configuration, 7 nodes are assigned to the underwater server.

Evaluation Metrics. We evaluate the latency variation of CockroachDB when running the TPC-C benchmark [83], a transaction processing benchmark provided by CockroachDB as an official performance tester. We adopt the normalized latency as the metric to evaluate the amount of overhead produced by the attack at increasing volumes, which is the measured latency divided by the baseline latency without acoustic injection. We then inject the sound at the fixed frequency of 5.1 kHz for the benchmark duration, which persists for ten minutes.

Results and Observations. Figure 6 shows the normalized access latency at increasing injection volumes. The results show that the performance degradation of CockroachDB [23] increases regardless of the number of data nodes. Even if fewer nodes are allocated to the underwater server, acoustic injection can almost linearly decrease the overall performance of the CockroachDB cluster with an average latency increase of 43.7%. In addition, the underwater attack achieves the highest latency increase at three assigned node settings with 38 dB ΔSPL by 92.7% on average. Once the volume is above 38 dB ΔSPL, the underwater nodes enter an out-of-service state and the operations to the CockroachDB cluster cannot be resolved, This causes abnormal termination of the TPC-C benchmark and removal of the underwater node.

5.2. Induced Automatic Node Removal in Distributed Filesystems

Unlike distributed databases which store relational and structured data, distributed filesystems are widely used [24],

[88], [89] in data centers for storing unstructured data. Latency manipulation in distributed filesystems can provoke severe imbalance of the I/O loads in data nodes and decrease data store reliability because fewer nodes remain available for replica storage. Thus, we further evaluate our attack capability to manipulate distributed filesystems. Specifically, we evaluate HDFS [90], a popular distributed filesystem which serves as the data store backend of Hadoop [24] for high-throughput distributed computing in data centers.

Experimental Setup. In this analysis, we adopt the same setup settings used in Section 5.1, where one HDFS data node is allocated to each server. We then apply our acoustic injection attack at 5.1 kHz frequency at increasing volumes.

Evaluation Metrics. We leverage the Hadoop DFSIO benchmark [30] to monitor file accesses to HDFS [90]. We analyze the HDFS logs to extract any changes in workflow and status when accessing the 32 files of 100MB each in the benchmark. We also monitor the liveness of the storage devices in the RAID 5 during the injection.

Results and Observations. At volumes below 152 dB SPL, the acoustic injection does not cause any change in status. At 152 dB SPL ($\sim$38 dB ΔSPL), Figure 7 shows the HDFS workflow and node status while depicting the liveness of each HDD (numbered from 1 to 4) in the RAID 5 configuration. As shown in the graph, the acoustic attack causes abnormal HDD activity which disables the HDFS service after 4.3 minutes. Specifically, after 2.4 minutes of benchmark execution, the underwater node cannot serve incoming I/O requests due to an HDD that became unresponsive because of the sound injection. As a consequence of this effect, the entire RAID 5 became unavailable for data storage. However, the HDFS stays online for another 1.8 minutes (from 2.4 to 4.2 minutes) as the RAID 5 drops the unresponsive HDD1 due to its unavailability. When a second HDD became unresponsive (HDD2) at 4.3 minutes for the prolonged attack, the HDFS remains blocked without serving any other file access operation. Finally, as RAID 5 requires at least three disks to maintain its correct functioning, and the unresponsive HDDs fall below this threshold, the second unresponsive disk HDD2 is dropped at 15.1 minutes causing the RAID failure. Consequently, the HDFS removes the underwater data node.

Such automatic removal of the data nodes from the filesystem due to the attack increases the I/O burden on other data nodes and leaves fewer nodes for storing data replicas dedicated to fault tolerance. This shows how an attacker can automatically induce the removal of selective nodes, maliciously redirecting workload over other data nodes causing overloading and compromising fault tolerance.

5.3. Load Manipulation

In this evaluation, we quantify an attacker's ability to manipulate resource monitoring and allocation applications commonly used in data centers. Various organizations, such as Akamai and Cisco [91], use OpenNebula [25] to monitor and allocate resources based on server resource availability and to ensure load balance between nodes [51]. By manipulating resource allocation, an attacker can force tenants to be assigned to slower servers or an already compromised server to enable the other attacks that require tenant colocation [14], [15], [16]. Although defenses against colocation attacks have been proposed [92], [93], they focus on modifying placement policies and do not consider an attacker that can effectively remove a server from the pool of available resources by reducing storage system functionality.

Experimental Setup. To evaluate the effect of our attack on resource migration and colocation, we use OpenNebula to balance VM assignments between two servers configured as described in Section 4. Both servers are connected to a LAN using a switch to a laptop running OpenNebula as an administrator. The administrator laptop monitors the server states and instantiates VMs to be automatically assigned to each server based on resource availability. Our evaluation is separated into two parts. First, we evaluate how acoustic injection increases latency in individual VMs to determine the affected states. Second, we verify whether an attacker can manipulate the resource manager to force assignment to a particular server (in this case, the attacker forces assignment to the on-land server by blocking the use of the underwater server). This evaluation consists of instantiating 50 VMs and tracking the resource assignments during the acoustic injection at increasing injection volumes.

5.3.1. Effect on VM Status. In this evaluation, we manually assign and instantiate 3 VMs running an Ubuntu OS and writing 1 GB of data to a file using *dd* on the target server during acoustic injection at increasing volumes. For this experiment, the VMs are assigned to one HDD in the underwater server to observe how HDD vulnerability impacts the VM status. We then constantly increase the volume until the VM experiences disk failure.

Evaluation Metrics. For this experiment, we evaluate the average time taken for each individual VM to complete each state during the acoustic injection. Whenever OpenNebula instantiates a VM, the VM passes through various states (INIT, PROLOG, BOOT, and RUNNING [91]). Note that only two of these states (PROLOG and RUNNING) require access to the storage system because the initialization and VM booting run in the application SSD with the server's operating system. The PROLOG state transfers VM files to the host, and, in the RUNNING state, the application runs, meaning that the host server's storage systems will respond to the VM application's I/O requests.

Results and Observations. The evaluation results in Figure 8 shows that PROLOG and RUNNING VM states have increasing latency with increasing injection volume. For the PROLOG state, the average latency increases up to 10%, while the average latency for the RUNNING state increases by a maximum of 280%. The VM fails at 36 dB ΔSPL because the disk becomes unresponsive.

5.3.2. Effect on VM Distributions. For this evaluation, we instantiate 50 VMs and observe which server OpenNebula

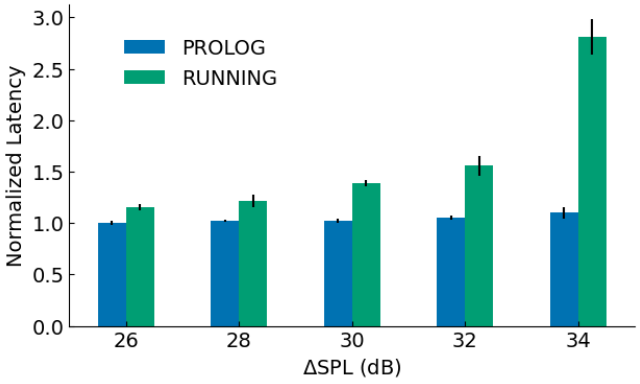

Figure 8: Normalized latency increase during the PROLOG and RUNNING states for individual VMs at increasing injection volumes.

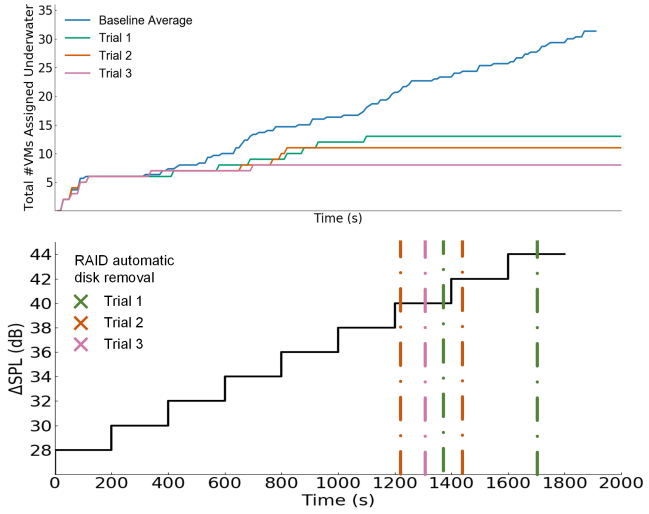

Figure 9: (Top) Average number of VMs assigned to the underwater server over time in the no-attack case (in blue) and with a 5.1 kHz acoustic injection. (Bottom) Corresponding volume increment over time. The graphs show the time when disks are automatically dropped from RAID 5. Note that after the second disk drop, the RAID 5 fails.

automatically assigns each VM to during acoustic injections at increasing injected volumes. In this case, the VMs run on the underwater RAID 5 when assigned to the target server to determine whether the fault tolerance system prevents manipulation of load balancing.

Evaluation Metrics. We evaluate the number of instantiated VMs assigned to each of the two servers. The attack is considered successful by observing a 10% shift in VM assignment from the submerged target to the on-land server. We perform our attack at volumes increasing by 2 dB every 210 seconds, which is 10% of the average time taken for 50 VMs to finish running with no acoustic injection.

Results and Observations. Figure 9 shows the total VM assignment to the underwater server at increasing volumes. From the results, we can observe a maximum of 74% and a minimum of 58% drop in the number of VMs assigned to the underwater server by OpenNebula when reaching up to 44 dB ΔSPL. As in the previous experiments, at high sound levels, the RAID 5 detects the first failure and automatically drops the corresponding disk at approximately 38 dB ΔSPL. However, RAID 5 continues its operation because it still has 3 of the 4 disks. At 44 dB ΔSPL, RAID 5 drops the second disk causing the RAID 5 to fail, and the VMs become permanently blocked in RUNNING state since they cannot interact with their storage disks. Once the RAID fails, the VMs cannot recover even after the injection stops. From these results, we see that the attacker can redirect the VM assignment to the on-land server by decreasing the performance of RAID 5 in the underwater server. We also observe that the overall server performance decreases after each experiment trial, as shown in Figure 9-(top). Both disks are dropped sooner from RAID, and the total number of assigned VMs decreases from 13 to 8. This shows how not only the acoustic attack can manipulate load distribution but also induce a permanent degradation of the storage systems without inducing a complete denial of service.

5.4. Latency Control on Real-World Workloads

To understand how acoustic injection can be used to control the latency of real-world workloads, we run the first 50k requests of three SNIA traces [26] on our submerged server using a RAID 5 device with 4 partitions of 60 GB size each. Among the SNIA traces taken from the operation of real data center workloads for various applications, we select typical data center workloads of a web server (abbreviated as 'web'), a proxy server (abbreviated as 'prxy'), and a media content server (abbreviated as 'mds'). We run the traces while performing acoustic injection at increasing volumes as in the previous evaluation.

Evaluation Metrics. We consider the average number of fulfilled I/O requests at increasing volumes for each benchmark. The benchmark is executed using RAIDmeter [94], a block-level trace replay tool, using a finite, constant timespan for each benchmark based on the time taken to finish sending IO requests (∼38 minutes for mds, ∼3 minutes for prxy, and ∼22 minutes for web). We characterize attack success as the ability to cause a measurable decrease in fulfilled requests, and we associate the decrease in request fulfillment with the attacker's ability to predictably manipulate application performance.

Results and Observations. Figure 10 shows each benchmark's normalized request fulfillment results at increasing volumes ranging from 26 to 38 dB ΔSPL. At 40 dB ΔSPL, RAID 5 fails for all trials of all benchmarks. From these results, we see an approximately linear trend with the number of fulfilled requests decreasing in the range from 26 to 30 dB ΔSPL. We note a spike in request fulfillment at 32 dB ΔSPL, which occurs when the slowest disk, which bottlenecks the RAID configuration, is dropped from the array. These results indicate that an acoustic injection

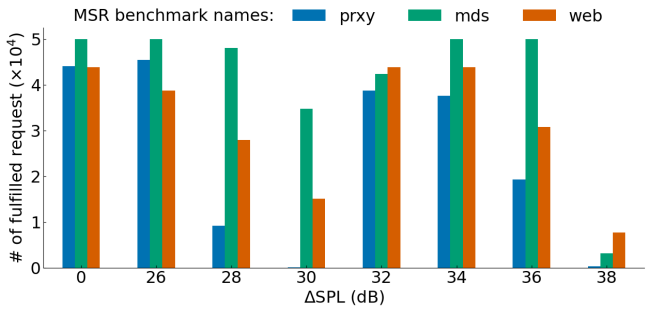

Figure 10: Number of fulfilled requests are shown for three MSR benchmarks at increasing injection volumes.

can measurably alter the performance of real data center workloads by changing the injection volume.

5.5. Evaluation on Hybrid Storage Architectures

Unlike mechanical HDDs, SSDs store data in flash memory on a silicon die and are less likely to be affected by our underwater acoustic injection attack. As discussed in Section 2.1, hybrid storage architecture of modern data centers typically deploys SSDs as cache of HDDs [29], [35], [36], [95]. Therefore, we evaluate how our underwater attack affects such a storage architecture.

Experimental Setup and Metrics. Intel's OpenCAS [96] is a software-level caching tool that allows accelerated access to slow storage devices (e.g., HDDs) by adding a faster device (e.g., SSDs) as a cache. It is a kernel module that allows the creation of a block device to represent the cached HDDs. Thus, we leverage OpenCAS to integrate a SSD cache with a write-back policy for HDDs in a RAID 5 configuration using *mdadm*. As in the previous evaluation, we perform the acoustic injection at 5.1 kHz with a fixed volume of 30 dB ΔSPL based on the previous observations. Then, we use FIO to evaluate the performance of the cached HDDs to evaluate the attacker's capabilities by monitoring the latency and bandwidth of four selected workloads: sequential write (SW), sequential read (SR), random write (RW), and random read (RR). We vary the allocated SSD size of the cache to demonstrate how the cache size impacts the attack efficiency.

Results and Observations. Figure 11 shows the bandwidth degradation when running the four FIO workloads with different cache sizes. Similar to previous experiments in Section 4.1 write operations are more affected by acoustic injections leading to a more noticeable performance drop in write-intensive workloads than in read-intensive workloads.

For workloads with more random access behaviors (i.e., RW and RR), the performance degradation is more significant because random data access to HDDs will incur frequent HDD actuator arm movement, making it easier to be affected by the sound-induced vibrations. Moreover, the hit ratio of RW workload is under 1%, whereas sequential write achieves a higher hit ratio, where its hit ratios are 33.3%, 56.9%, 68.6%, and 76.1% when allocating 0.5 GB, 1 GB, 1.5 GB, and 2 GB cache size respectively. Therefore, the bandwidth degradation provoked by the attack is alleviated in workloads with a high hit ratio.

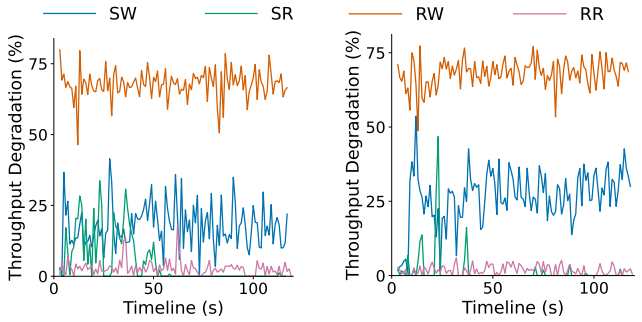

(a) Cache size 0.5 GB. (b) Cache size 1 GB.

(c) Cache size 1.5 GB. (d) Cache size 2 GB.

Figure 11: Bandwidth degradation caused by our attack while running sequential write (SW), sequential read (SR), random write (RW), and random read (RR) FIO workloads.

Figures 12 and13 show the cumulative distribution function (CDF) of access latency at different cache sizes when running RW and SW workloads in our attack. In the RW workload, the access latency under the attack always distributes between 200–800 ms, while the access latency fits within 1–200 ms in the benign case. Even if the cache size increases, the latency increase incurred by our attack is still significant. Since the RW workload presents a low hit ratio – less than 1% – to the cache, our attack can significantly degrade the performance of the cached HDDs because most of the I/O requests are served by the HDDs. In contrast, a large amount of I/O requests are served by the cache in the SW workload, thus our attack has less impact on sequential writes. However, the performance degradation is still nontrivial, as shown in Figure 12 and Figure 13. Therefore, our underwater acoustic injection attack can make the storage system unpredictable, which is critical to provide deterministic latencies [43], [45] desired by data center providers, with overall performance degradation.

5.6. Evaluation for Open-Water Deployments

To evaluate whether such acoustic attacks can be performed in open water and to understand the distance limit for the attacker, we deploy our testbed setup to a lake

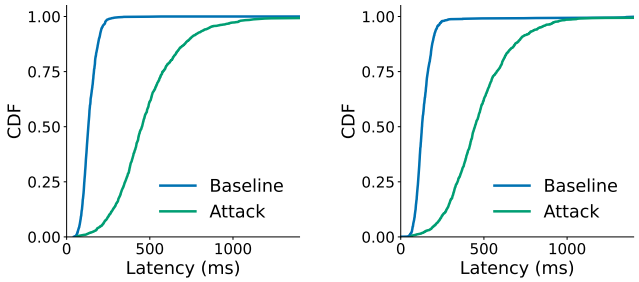

(a) Cache size 0.5 GB. (b) Cache size 1 GB.

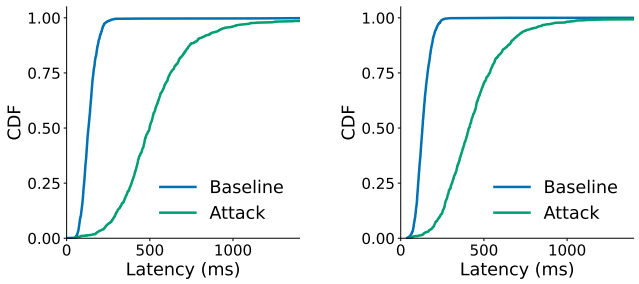

(c) Cache size 1.5 GB. (d) Cache size 2 GB.

Figure 12: CDF of latency when running the random write (RW) workload of FIO with or without underwater injection.

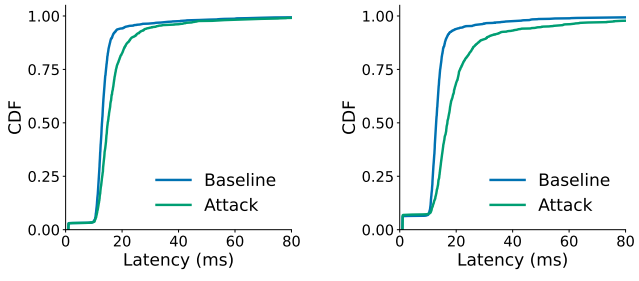

(a) Cache size 0.5 GB. (b) Cache size 1GB.

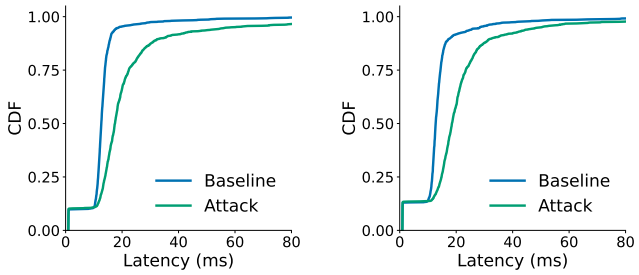

(c) Cache size 1.5 GB. (d) Cache size 2 GB.

Figure 13: CDF of latency when running the sequential write (SW) workload of FIO with or without underwater injection.

(see Figure 1c). For this scenario, we weighed the metal enclosure with bags of sand to reach the required water level. Then, we measure the RAID 5 write throughput at increasing volumes and distances from the sound source.

Evaluation Metrics. We use FIO [81] to record RAID 5 throughput over 30-second spans for 3 consecutive trials. For our volume variation evaluation, we consider a 30 cm distance from the sound source. For our distance evaluation, we consider the maximum achievable distance where the attack can successfully degrade the RAID 5 performance.

Results and Observations. Figure 4 (b) shows the throughput variation at increasing volumes. We observe a similar degradation as in the laboratory testbed scenario (Figure 4 (a)), but a higher volume is required to reach the same amount of degradation. We suspect that the bags of sand altered the propagation properties of the vibrations. Such results also indicate how our laboratory setup, even if limited, can be used to simulate realistic scenarios.

For the maximum achievable distance, we induce an average degradation of 61% at 6.35 meters from the enclosure, which represents the maximum distance available in our lake scenario. This result shows how sophisticated acoustic injection attacks can be performed at far distances with commercially available speakers.

5.7. Finite Element Simulation

At the time of writing, there are no available testing facilities for commercial UDC deployments, thus in our testbed evaluation, we approximate the UDC vessel with an aluminum enclosure. To provide a more realistic preliminary analysis of the attack, we simulate our acoustic injection using a COMSOL Finite Element Method (FEM) [97] model. As used in previous work [21], such modelling allows combining multiple physics phenomena for simulations of real-world scenarios, such as, the sound propagation between two media (seawater and the vessel mechanical structure).

We based our analysis on Microsoft's Project Natick resources [3] and publicly available information on subsea vessel prototypes. We build a 1/100 scaled steel hollow vessel (12.2 m x 1.4 m radius) [98]) with 11.7 mm steel thickness based on thickness recommendations for underwater pressure vessels [99] (See Figure 14). We scale the model to allow for a finer-grained mesh for more accurate simulation results. We account for the 35 m depth below sea level as described for Project Natick, with a salinity level of 35 (reference salinity for seawater [100]). We consider a budgeted attacker with a military-grade speaker which can reach SPL of 220 dB (based on the SPL of sonars [101]) simulated in our model as a sound source generating spherical waves facing the flat surface of the vessel as depicted in Figure 14. As described in our theoretical analysis in Section 3.3, such vibrations propagate to the internal server racks and storage devices through their contact surfaces.

After simulating a frequency sweep of eigenfrequencies to find the structure resonance frequency, we set 6.95 KHz as our injection frequency. We then simulate the injection at 6 cm from the structure as per our capability evaluation, achieving an average total displacement of $\sim$360 nm along the three orientation axis with a maximum of $\sim$718 nm. We then estimate the maximum capability to induce

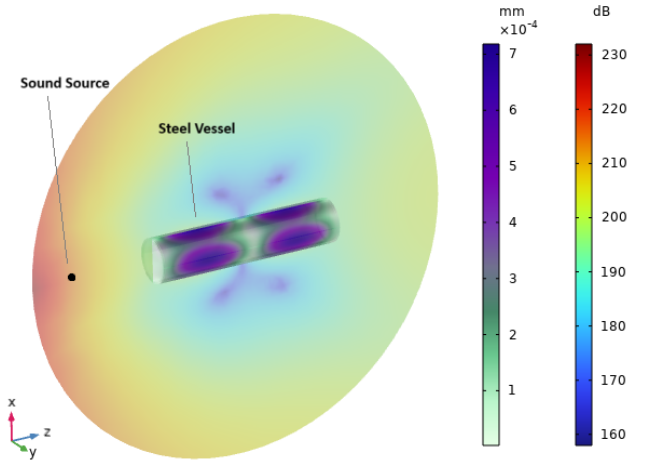

Figure 14: COMSOL simulation of the pressure and displacement of the 1/100 scale model of the Project Natick vessel [98] under a 220 dB SPL sound source in seawater.

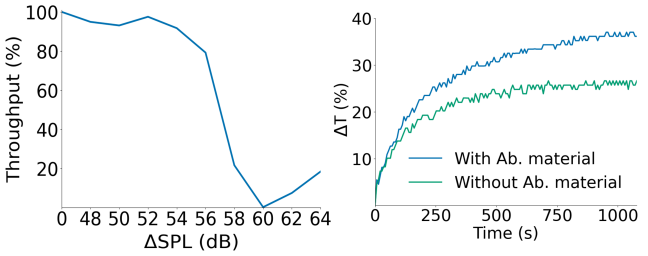

Figure 15: (Left) RAID 5 write throughput at increasing injection volumes at 5.1 kHz injection frequency. In 2 of 3 trials, a disk was automatically removed from RAID 5 at 60 dB ΔSPL above background noise. This causes the increasing throughput at 62 and 64 dB ΔSPL. (Right) Temperature increases in the presence of the absorbing material.

mechanical vibrations on the full-scale vessel, measured in terms of structure displacement along the three orientation axes, reaching an average total displacement of ∼145.5 nm. As a reference, per our PES analysis in Section 3.4 and the literature [21], [75] typically hard disk read and write from the magnetic platters by the read/write head, which floats about ∼5 nm above the disk surface, and, in the case of enterprise-range drives used in data centers, can deviate from the center of the track by no more than ∼7 nm to avoid reading and writing errors. Such simulation results indicate how acoustic injections can potentially generate vibrations strong enough to propagate inside steel vessel structures.

Based on Eqs. 3 and 4, we know that the SPL in seawater attenuates exponentially, and the displacement induced in a solid structure is proportional to the applied force given by the intensity of the injected sound. Therefore, we can estimate the maximum distance achievable by our model using a conservative attenuation coefficient α of 10^{-1} Nepers/km [102] (this value is taken at 10 kHz reference frequency while our frequency is lower). We find that such an attacker can theoretically induce an average of 131.2 nm displacement at 1 km from the structure, revealing small decrease in vibration over large distances.

6. Defenses

6.1. Potential Defenses

Passive Attenuation Using Absorptive Material. We repeat the experiment in Section 4.2 at the highest volume that our speaker can achieve (180 dB SPL).

Our results in Figure 15 show that we can cause similar changes in throughput, meaning that an attacker could overpower the vibration absorption by increasing the sound volume. We also evaluate the temperature of the server with and without the absorbing material by running a CPU stress test using the *stress* utility [103] for 20 minutes and logging the average temperature of the server's CPU cores. From the results, we see about a 10% difference in the temperature increase with and without the foam for a single server. Microsoft's Project Natick submerged data center contains 864 servers [3] that generate heat in an enclosed space; this dense configuration implies a significant increase in heat retention, which will be unsustainable for server health. Our results show that the use of sound absorption materials requires a careful redesign of the internal data center structure by considering the tradeoff between cooling efficiency and acoustic attack protection. Design solutions proposed in research for acoustic attenuation typically focus on attenuating internal fan and disk noise [104], [105], [106] rather than mechanical vibrations coming from external sources. Furthermore, they might require the isolation of each individual server rack in materials such as polyurethane acoustic foam (as the one used in our experiments) or custom acoustic metamaterials targeting specific frequency ranges [107]. These solutions might be adapted to target the resonance frequency ranges exploited by the attacker, leveraging the data center's internal server configuration.

Active Noise Cancellation. Another commonly suggested defense against acoustic injection attacks is noise cancellation. Previous works have argued that noise cancellation is an impractical defense because it is difficult to generate a noise signal with the required equal amplitude to the injected signal and which envelops the entire region [21]. Another consideration in the underwater scenario is noise pollution. As described in Section 2.3, sound travels faster in water than in air, so emitting high-volume sound waves surrounding the data center could be detrimental to the environment. Marine life, such as whales [108] and fishes [109], is harmed by noise pollution. As such, this defensive measure would affect the overall environmental sustainability of UDCs. In addition, acoustic emission can interfere with sound-based communication systems underwater.

Sensor Fusion for Detection. Sensor fusion-based detection techniques using hydrophones, accelerometers, and vibration sensors could also be used to detect acoustic injection attacks on data centers. Such defensive measures imply the

deployment of additional hardware and detection control software to the data center. However, it is worth noticing that external sensors such as cameras, accelerometers, and microphones are also vulnerable to the effect of acoustic vibrations at the resonance frequency which can cause measurement errors, false triggering, and DoS, as demonstrated in previous works [54], [56], [110], [111].

Feedback Controller and Firmware Modifications. HDDs have feedback controllers that compensate for vibrations in a narrow band of frequencies to prevent disruption [112]. For instance, Bolton et al. [21] implemented in simulation an augmented feedback controller by updating the storage device firmware. This and other hard disk proprietary firmware modifications can be applied to attenuate vibrations up to a certain displacement level on single hard disk drives.

Physical Security. Our open-water scenario shows how our attack can be deployed more than 6 meters away from the targeted enclosure with the maximum volume achievable by our setup (commercial speaker and amplifier). In UDC deployments, physical surveillance mechanisms such as high-resolution underwater cameras and motion sensors controlled by trained personnel can be placed to prevent attackers from reaching the enclosure's proximity to perform the attack. The detection range depends on their visibility range and resolution which might vary based on the depth and light attenuation [113] reaching a maximum of 20-30 meters (approximately 65-100 feet) in clear water [114]. Our simulation based on the Project Natick [3] prototype deployment shows significant vibration displacement at 1 km from the structure using a 220 dB SPL sound source similar military-grade sonars used in the real world [101]. This preliminary analysis reveals that physical security mechanisms should be designed to take into account the attacker's capabilities, surveillance sensor accuracy, and environmental conditions.

6.2. Proposed Defense

We propose a proof-of-concept detection mechanism that uses a machine learning model to detect multiple simultaneous, low-volume throughput degradations instead of attenuating single disk vibrations. Our defense relies on analyzing the throughput of disk clusters in close physical proximity to differentiate between normal performance degradation and acoustic injection attacks. This approach is based on the idea that sound-induced vibrations affect multiple disks simultaneously because sound radiates, generating a pattern of throughput changes that can be detected. Such a defense can be deployed at the cloud resource management level without requiring access to proprietary HDD firmware or datacenter physical structure redesign.

Evaluation Method and Results. To identify throughput degradation in multiple disks, we consider the full-HDD architecture and generate a profile of each of the four disks in the RAID 5 configuration. Such profiling can be customized based on the system, and for our proof-of-concept analysis, we use the FIO [81] sequential write workload for 30 seconds. We first collect the storage system throughput for 100 trials on a 100 MB partition for each disk without any acoustic injection. We then run 100 30-second trials of the same benchmark during acoustic injection at 26 dB ΔSPL, 28 dB ΔSPL, and 30 dB ΔSPL, the lowest volumes which cause the minimal throughput change in our scenarios. To detect the attack for different injection volumes, we use k-means clustering with Partial Curve Mapping (PCM) [115] metric. PCM uses arch length and area of the throughput data with respect to time to measure the similarity between disk performances. To quantify our defense's ability to differentiate between attack and no-attack cases, we generate 1,000 combinations of all four disks with random benign and attack throughputs. We repeat this for each volume level. We consider an attack if at least three disks show anomalous throughput behavior. Through this evaluation, we achieve a 0% False Positive Rate and 98.2% True Positive Rate. Although we only consider a limited set of hard disk drives and FIO benchmark profiling, this proof-of-concept defense shows how the use of ML techniques can allow the recognition of localized degradation patterns that can reveal the presence of potential acoustic attacks.

Post-detection Defense. Recent replication [116], [117] and erasure coding [118], [119], [120] optimization techniques explored in research can provide selective data redundancy for preventing data loss and ensuring high-quality fast data recovery in cloud settings. Replication techniques generate replicas of data and distribute them to multiple storage nodes located at different places, while erasure coding techniques compute multiple parities for stored data and use the erasure coding calculation with stored parities to recover failed storage nodes. Upon detecting an attack, the resource management system can be configured to leverage these advanced techniques to migrate the I/O requests to specific unaffected nodes outside the realm of the sound-affected areas which house the replicas of the affected data. This is possible since, as shown in our empirical evaluation, the storage regions affected by our attack are physically adjacent to each other due to the nature of the vulnerability.

7. Discussion

Long-Term Disk Degradation. Through our experiments, we note that HDDs suffer from long-term degradation due to acoustic injection attacks. While disks can be re-added to RAID and undetected disks can sometimes be re-detected by rebooting or by physical reconnection, three HDDs became completely undetectable and permanently damaged during our experiments. In underwater infrastructure deployments where the ability to access enclosed vessels to replace damaged storage devices is limited, even short-lived acoustic injections can cause severe performance loss. For instance, the latencies observed before and after the injections for the MSR benchmark of a proxy server showed $\sim$1,350% increase in average request latency. Thus, the attack continues to have an effect after the end of the injection.

RAID 5 Disk Bottleneck. Throughout our evaluations, we found a sudden increase in throughput after a regular decrease in throughput (see Figure 15, Figure 4, and Figure 10). This sudden spike coincided with the first RAID 5 automatically removing the slowest disk from the configuration (we used a 4-disk RAID 5 configuration, and RAID 5 requires a minimum of 3 disks). RAID 5 write requests require that parity is written to all disks in the configuration [121], so dropping the slowest disk can partially alleviate the bottleneck on parity writing. However, this effect rapidly decreases with increasing injection volumes.

Hybrid Storage Architecture. Existing data centers employ SSDs as the cache for HDDs because of limitations in the SSD technology as discussed in Section 2.1. SSDs are immune to performance degradation caused by acoustic vibrations as they contain no mechanical component [21], making them a potential solution to alleviate our underwater attacks. However, the cache cannot guarantee that data will not be evicted from SSDs to HDDs, and SSDs can be overused, which motivates redirect write requests to HDDs [29]. Therefore, acoustic attack remains a security threat to UDCs, as reliance on HDDs remains widespread.

Limitations. In this study, we evaluate an attacker's ability to reduce RAID 5 device performance using acoustic injection. Our laboratory testbed is an imperfect approximation of a UDC, which would be the real-world target of such attackers. While we use a FEM model to simulate a prototype deployment under more realistic attack scenarios, the simulation does not fully capture all the factors of a real-world subsea environment. Furthermore, real-world UDCs might significantly differ in size, type of enclosure, deployment in rough or salty water, and other physical parameters that might considerably impact the results of audio injection. However, our analysis, even if limited to proof-of-concept scenarios, unveils new sophisticated attack vectors that go beyond the simple DoS, influencing the behavior and reliability of traditional fault tolerance and load-balancing storage techniques that cannot withstand acoustic attacks. We also do not consider different RAID configurations or non-cache-based SSD hybrid architectures. Our analysis is limited to one single server deployment underwater and the application analysis focuses on standard benchmarks of realistic data center workflows (e.g., SNIA traces). Formal analysis has been used in previous work [122], [123] to validate the design of cloud storage systems. However, none of the existing works applied formal analysis for a large-scale study of throughput degradation incurred by acoustic attacks in cloud nodes. We leave this analysis as future work.

Acoustic Safety Considerations. All the experiments were conducted in controlled, isolated environments, and participants were wearing appropriate hearing protection.

8. Related Work

Signal Injection Attacks. Signal injection attacks have been performed using a wide range of signal types, including optical [111], [124], acoustic [56], [59], and electromagnetic [110], [125], [126]. By exploiting component sensitivity to these signal types, researchers have performed attacks on various devices, including temperature sensors [125], hard disk drives [20], [21], [22], autonomous vehicles [124], underwater acoustic networks [127], [128], [129], [130], and automatic speech recognition systems [61], [111]. Unlike these works, we investigate acoustic injection in fault-tolerant storage configurations and data center management systems, which are not designed to process acoustic signals.

Underwater Cyber-physical Security. Maritime cyber-physical security generally focuses on defending communication systems [131], [132] and underwater acoustic networks [127], [128], [129], [130]. Researchers studying physical attacks in the underwater domain generally focus on directly tampering with components (e.g., damaging communication cables [132]) or spoofing using in-band signal injection (e.g., spoofing acoustic communication between underwater network devices [127]). While a recent work evaluated acoustic injection on HDDs [20] in underwater scenarios, it was limited to cause DoS on a single consumer-grade HDD and did not include an in-depth analysis of the attacker's capabilities or suggest defensive measures. Our acoustic injection attack, on the other hand, evaluates complex RAID systems consisting of enterprise HDDs and fine-grained control over data center resource allocation and database and distributed filesystem performance. We also evaluate commonly suggested defenses and propose a novel defense against underwater acoustic injection attacks.

9. Conclusion

We evaluate an attacker's ability to use acoustic injection to gain fine-grained control over the performance of fault-tolerant storage systems used in data centers. We deploy a submerged enclosure and test our attack in laboratory and open-water scenarios. Through this evaluation, we identify threats to distributed systems and data center management tools such as distributed filesystems, databases, and data center resource managers. We also describe limitations in commonly suggested defenses against acoustic injection and suggest a novel proof-of-concept ML-based defense which reaches 0% False Positive Rate and 98.2% True Positive Rate in our testbed scenario. Furthermore, we provide a theoretical analysis of the sound-induced vibrations and simulate the effect of our attack in realistic UDC deployment scenarios. We hope that our attack characterization and suggested defense improve the security of UDCs, which have recently emerged as a novel environmentally sustainable cloud computing technology.

Acknowledgements

We thank the anonymous shepherd and reviewers for their valuable comments. This research was supported in part by the National Science Foundation (NSF) under CNS-2055014, the Air Force Office for Scientific Research under FA8650-19-1-1741 and FA8650-19-1-0169, gifts from Meta

and Texas Instruments, and JST CREST JPMJCR23M4. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the NSF.

References

[1] D. Geiger, E. Thomas, and A. Barr, "Data centers are sprouting up as a result of the AI boom, minting fortunes, sucking up energy, and changing rural America," https://www.businessinsider.com/ai-data-energy-centers-water-energy-land-2023-10, 2023.

[2] W. Wong, "AI and cloud workloads drive data center demand," https://www.datacenterknowledge.com/buildconstruction/ai-and-cloud-workloads-drive-data-center-demand, 2023.

[3] J. Roach, "Microsoft finds underwater datacenters are reliable, practical and use energy sustainably," https://news.microsoft.com/source/features/sustainability/project-natick-underwater-datacenter/, 2020.

[4] P. Judge, "Subsea Cloud announces three underwater data center projects," https://www.datacenterdynamics.com/en/news/subsea-cloud-announces-three-underwater-data-center-projects/, 2022.

[5] ——, "Work begins on Chinese underwater data center," https://www.datacenterdynamics.com/en/news/work-begins-on-chinese-underwater-data-center/, 2022.

[6] M. Shaw and M. Goldstein, "Open CloudServer JBOD specification," https://www.opencompute.org/documents/microsoft-ocs-v1-jbod-blade, 2015.

[7] B. Cutler, S. Fowers, J. Kramer, and E. Peterson, "Dunking the data center," *IEEE Spectrum*, 2017.

[8] K. Chung, Z. T. Kalbarczyk, and R. K. Iyer, "Availability attacks on computing systems through alteration of environmental control: smart malware approach," in *ACM/IEEE ICCPS*, 2019, pp. 1–12.

[9] A. Libri, A. Bartolini, and L. Benini, "pAElla: Edge AI-based real-time malware detection in data centers," *IEEE IoT-J*, 2020.

[10] M. Abdelsalam, R. Krishnan, Y. Huang, and R. Sandhu, "Malware detection in cloud infrastructures using convolutional neural networks," in *IEEE CLOUD*, 2018.

[11] H. Liu, "A new form of dos attack in a cloud and its avoidance mechanism," in *ACM CCSW*, 2010.

[12] Z. Anwar and A. W. Malik, "Can a ddos attack meltdown my data center? a simulation study and defense strategies," *IEEE Communications Letters*, 2014.

[13] J. Chen, X. Zheng, H.-X. Duan, J. Liang, J. Jiang, K. Li, T. Wan, and V. Paxson, "Forwarding-loop attacks in content delivery networks." in *NDSS*, 2016.

[14] Y. Zhang, A. Juels, M. K. Reiter, and T. Ristenpart, "Cross-VM side channels and their use to extract private keys," in *ACM CCS*, 2012.

[15] M. S. Inci, B. Gulmezoglu, G. Irazoqui, T. Eisenbarth, and B. Sunar, "Seriously, get off my cloud! cross-vm rsa key recovery in a public cloud," *Cryptology ePrint Archive*, 2015.

[16] T. Ristenpart, E. Tromer, H. Shacham, and S. Savage, "Hey, you, get off of my cloud: exploring information leakage in third-party compute clouds," in *ACM CCS*, 2009.

[17] M. A. Islam, L. Yang, K. Ranganath, and S. Ren, "Why some like it loud: Timing power attacks in multi-tenant data centers using an acoustic side channel," *ACM POMACS*, 2018.

[18] M. A. Islam, S. Ren, and A. Wierman, "Exploiting a thermal side channel for power attacks in multi-tenant data centers," in *ACM CCS*, 2017.

[19] M. A. Islam and S. Ren, "Ohm's law in data centers: A voltage side channel for timing power attacks," in *ACM CCS*, 2018.

[20] J. Sheldon, W. Zhu, A. Abdullah, K. Butler, M. J. Islam, and S. Rampazzi, "Deep note: Can acoustic interference damage the availability of hard disk storage in underwater data centers?" in *ACM HotStorage*, 2023.

[21] C. Bolton, S. Rampazzi, C. Li, A. Kwong, W. Xu, and K. Fu, "Blue note: How intentional acoustic interference damages availability and integrity in hard disk drives and operating systems," in *IEEE S&P*, 2018.

[22] M. Shahrad, A. Mosenia, L. Song, M. Chiang, D. Wentzlaff, and P. Mittal, "Acoustic denial of service attacks on hard disk drives," in *ACM ASHES*, 2018.

[23] R. Taft, I. Sharif, A. Matei, N. VanBenschoten, J. Lewis, T. Grieger, K. Niemi, A. Woods, A. Birzin, R. Poss, P. Bardea, A. Ranade, B. Darnell, B. Gruneir, J. Jaffray, L. Zhang, and P. Mattis, "CockroachDB: The resilient geo-distributed sql database," in *ACM SIGMOD MOD*, 2020.

[24] Facebook, "Apache Hadoop," https://hadoop.apache.org/, 2023.

[25] D. Milojičić, I. M. Llorente, and R. S. Montero, "OpenNebula: A cloud management tool," *IEEE Internet Computing*, 2011.

[26] D. Narayanan, A. Donnelly, and A. Rowstron, "MSR Cambridge traces (SNIA IOTTA trace set 388)," in *SNIA IOTTA Trace Repository*, G. Kuenning, Ed. Storage Networking Industry Association, Mar. 2007. [Online]. Available: http://iotta.snia.org/traces/block-io?only=388

[27] Microchip, "Hybrid RAID solutions," https://www.microsemi.com/product-directory/storage-innovations/4061-hybrid-raid.

[28] J. Niu, J. Xu, and L. Xie, "Hybrid storage systems: A survey of architectures and algorithms," *IEEE Access*, 2018.

[29] S. Wang, Z. Lu, Q. Cao, H. Jiang, J. Yao, Y. Dong, and P. Yang, "BCW: Buffer-Controlled writes to HDDs for SSD-HDD hybrid storage server," in *USENIX FAST*, 2020.

[30] Hadoop, "Benchmarking HDFS using DFSIO," https://subscription.packtpub.com/book/data/9781783285471/1/ch01lvl1sec17/benchmarking-hdfs-using-dfsio, 2023.

[31] Forbes, "Microsoft Teams has seen a 775% rise in users in Italy because of COVID-19," https://www.forbes.com/sites/martingiles/2020/03/30/microsoft-cloud-service-775-percent-rise-covid-19/?sh=60f293086862, 2023.

[32] McKinsey & Company, "Investing in the rising data center economy," https://www.mckinsey.com/industries/technology-media-and-telecommunications/our-insights/investing-in-the-rising-data-center-economy, 2023.

[33] Sunbird, "Data Center Components," https://www.sunbirddcim.com/glossary/data-center-components.

[34] Alibaba, "Pangu – the high performance distributed file system by Alibaba cloud," https://www.alibabacloud.com/blog/pangu_the_high_performance_distributed_file_system_by_alibaba_cloud_594059, 2018.

[35] B. Calder, J. Wang, A. Ogus, N. Nilakantan, A. Skjolsvold, S. McKelvie, Y. Xu, S. Srivastav, J. Wu, H. Simitci, J. Haridas, C. Uddaraju, H. Khatri, A. Edwards, V. Bedekar, S. Mainali, R. Abbasi, A. Agarwal, M. F. u. Haq, M. I. u. Haq, D. Bhardwaj, S. Dayanand, A. Adusumilli, M. McNett, S. Sankaran, K. Manivannan, and L. Rigas, "Windows Azure Storage: A highly available cloud storage service with strong consistency," in *ACM SOSP*, 2011.

[36] S. Muralidhar, W. Lloyd, S. Roy, C. Hill, E. Lin, W. Liu, S. Pan, S. Shankar, V. Sivakumar, L. Tang, and S. Kumar, "f4: Facebook's warm BLOB storage system," in *USENIX OSDI*, 2014.

[37] FDCServers, "How to choose between SSD and HDD storage within your data center?" https://www.fdcservers.net/blog/how-to-choose-between-ssd-and-hdd-storage-within-your-data-center, 2021.

[38] B. Schroeder, A. Merchant, and R. Lagisetty, "Reliability of NAND-based SSDs: What field studies tell us," *Proceedings of the IEEE*, 2017.

[39] C. Min, K. Kim, H. Cho, S.-W. Lee, and Y. I. Eom, "SFS: random write considered harmful in solid state drives." in *USENIX FAST*, 2012.

[40] S. Im and D. Shin, "Flash-aware RAID techniques for dependable and high-performance flash memory SSD," *IEEE ToC*, 2011.

[41] S. Yan, H. Li, M. Hao, M. H. Tong, S. Sundararaman, A. A. Chien, and H. S. Gunawi, "Tiny-Tail Flash: Near-perfect elimination of garbage collection tail latencies in NAND SSDs," *ACM TOS*, 2017.

[42] F. Wu, J. Zhou, S. Wang, Y. Du, C. Yang, and C. Xie, "FastGC: Accelerate garbage collection via an efficient copyback-based data migration in SSDs," in *DAC*, 2018.

[43] M. Hao, L. Toksoz, N. Li, E. E. Halim, H. Hoffmann, and H. S. Gunawi, "LinnOS: Predictability on unpredictable flash storage with a light neural network," in *USENIX OSDI*, 2020.

[44] H. Li, M. L. Putra, R. Shi, X. Lin, G. R. Ganger, and H. S. Gunawi, "IODA: A host/device co-design for strong predictability contract on modern flash storage," in *ACM SOSP*, 2021.

[45] H. Li, Y. Zhang, D. Li, Z. Zhang, S. Liu, P. Huang, Z. Qin, K. Chen, and Y. Xiong, "URSA: Hybrid block storage for cloud-scale virtual disks," in *EuroSys*, 2019.

[46] A. Ramirez, "SSD vs. HDD: Choosing the right storage for you," https://reviewed.usatoday.com/laptops/features/ssd-vs-hdd, 2023.

[47] P. M. Chen, E. K. Lee, G. A. Gibson, R. H. Katz, and D. A. Patterson, "RAID: High-performance, reliable secondary storage," *ACM Compututer Survey*, 1994.

[48] J. Shu, R. Zhu, Y. Ma, G. Huang, H. Mei, X. Liu, and X. Jin, "Disaggregated raid storage in modern datacenters," in *ACM ASPLOS*, 2023.

[49] B. Mao, H. Jiang, S. Wu, L. Tian, D. Feng, J. Chen, and L. Zeng, "Hpda: A hybrid parity-based disk array for enhanced performance and reliability," *ACM TOS*, 2012.

[50] Microsoft Azure, "What is IaaS?" https://azure.microsoft.com/en-us/resources/cloud-computing-dictionary/what-is-iaas.

[51] K. Yasar and A. Irei, "What is load balancing?" https://www.techtarget.com/searchnetworking/definition/load-balancing, 2023.

[52] OpenNebula Systems, "OpenNebula 6.6 Documentation," https://docs.opennebula.io/6.6/index.html, 2023.

[53] D. Halliday, R. Resnick, and J. Walker, *Fundamentals of physics*. John Wiley & Sons, 2013.

[54] X. Ji, Y. Cheng, Y. Zhang, K. Wang, C. Yan, W. Xu, and K. Fu, "Poltergeist: Acoustic adversarial machine learning against cameras and computer vision," in *IEEE SP*, 2021.

[55] Y. Cheng, X. Ji, W. Zhu, S. Zhang, K. Fu, and W. Xu, "Adversarial computer vision via acoustic manipulation of camera sensors," *IEEE TDSC*, 2023.

[56] T. Trippel, O. Weisse, W. Xu, P. Honeyman, and K. Fu, "Walnut: Waging doubt on the integrity of mems accelerometers with acoustic injection attacks," in *IEEE EuroS&P*, 2017.

[57] Y. Son, H. Shin, D. Kim, Y. Park, J. Noh, K. Choi, J. Choi, and Y. Kim, "Rocking drones with intentional sound noise on gyroscopic sensors," in *USENIX Security*, 2015.

[58] S. Khazaaleh, G. Korres, M. Eid, M. Rasras, and M. F. Daqaq, "Vulnerability of MEMS gyroscopes to targeted acoustic attacks," *IEEE Access*, vol. 7, pp. 89 534–89 543, 2019.

[59] Y. Tu, Z. Lin, I. Lee, and X. Hei, "Injected and delivered: Fabricating implicit control over actuation systems by spoofing inertial sensors," in *USENIX Security*, 2018.

[60] M. Gao, L. Zhang, L. Shen, X. Zou, J. Han, F. Lin, and K. Ren, "Exploring practical acoustic transduction attacks on inertial sensors in MDOF systems," *IEEE TMC*, 2023.

[61] G. Zhang, C. Yan, X. Ji, T. Zhang, T. Zhang, and W. Xu, "Dolphinattack: Inaudible voice commands," in *ACM CCS*, 2017.

[62] X. Lurton, *An introduction to underwater acoustics: principles and applications*, 2002.

[63] E. Kozaczka and G. Grelowska, "Theoretical model of acoustic wave propagation in shallow water," *Polish maritime research*, 2017.

[64] Y. Kim, "The underwater propagation of sound and its applications," *Dartmouth undergraduate journal of science*, 2012.

[65] C. Erbe, A. Duncan, and K. J. Vigness-Raposa, "Introduction to sound propagation under water," *Exploring Animal Behavior Through Sound: Volume*, 2022.

[66] J. M. Hovem, "Underwater acoustics: Propagation, devices and systems," *Journal of Electroceramics*, 2007.

[67] A. D. Pierce, *Acoustics: an introduction to its physical principles and applications*, 2019.

[68] S. Gu, S. Guo, and L. Zheng, "A highly stable and efficient spherical underwater robot with hybrid propulsion devices," *Autonomous Robots*, 2020.

[69] Genasys, "Long Range Acoustic Devices," https://genasys.com/lrad-products/.

[70] F. B. Jensen, W. A. Kuperman, M. B. Porter, H. Schmidt, and A. Tolstoy, *Computational ocean acoustics*, 2011.

[71] I. Harari, K. Grosh, T. Hughes, M. Malhotra, P. Pinsky, J. Stewart, and L. Thompson, "Recent developments in finite element methods for structural acoustics," *Archives of Computational Methods in Engineering*, 1996.

[72] M. S. Howe, *Acoustics of fluid-structure interactions*, 1998.

[73] Y. Feng and W. Gao, "On the strain energy distribution of two elastic solids under smooth contact," *Powder Technology*, 2021.

[74] U. Ingard, "Influence of fluid motion past a plane boundary on sound reflection, absorption, and transmission," *The Journal of the Acoustical Society of America*, 1959.

[75] A. Kwong, W. Xu, and K. Fu, "Hard drive of hearing: Disks that eavesdrop with a synthesized microphone," in *IEEE SP)*, 2019.

[76] "Seagate ST2000DL001 F3 commands," https://gist.github.com/rigred/d424673a113addf45805ad715b4b86a0.

[77] Lubell Labs, "Lubell Labs LL916, LL916H, and LL916C Pistonic Piezoelectric Underwater Speakers," https://www.lubell.com/LL916.html.

[78] SuperMicro, "2U Chassis," https://www.supermicro.com/products/chassis/2u/?chs=823.

[79] Seagate, "Exos 7E2 Datasheet," https://www.seagate.com/www-content/datasheets/pdfs/exos-7-e2DS1956-1-1709US-en_US.pdf.

[80] Solidigm, "D3-S4510," https://www.solidigm.com/products/data-center/d3/s4510.html.

[81] J. Axboe, "Flexible I/O Tester." https://github.com/axboe/fio.

[82] Dell, "Dell PowerEdge R610," https://i.dell.com/sites/csdocuments/Shared-Content_data-Sheets_Documents/en/R610-SpecSheet.pdf.

[83] CockroachDB, "Performance Benchmarking with TPC-C," https://www.cockroachlabs.com/docs/stable/performance-benchmarking-with-tpcc-large, 2023.

[84] TiDB, "Tidb: The advanced distributed sql database," https://www.pingcap.com/tidb/.

[85] A. Verbitski, A. Gupta, D. Saha, M. Brahmadesam, K. Gupta, R. Mittal, S. Krishnamurthy, S. Maurice, T. Kharatishvili, and X. Bao, "Amazon aurora: Design considerations for high throughput cloud-native relational databases," in *Proceedings of the 2017 ACM International Conference on Management of Data*, 2017.

[86] CockroachDB, "The history of databases at netflix: From cassandra to cockroachdb," https://www.cockroachlabs.com/blog/netflix-at-cockroachdb/, 2023.

[87] ——, "The new stack: Meet cockroachdb, the resilient sql database," https://www.cockroachlabs.com/blog/the-new-stack-meet-cockroachdb-the-resilient-sql-database/#:~:text=With%20companies%20like%20SpaceX%2C%20and,%2C%20resilience%2C%20and%20data%20locality., 2015.

[88] B. Zhu, Y. Chen, Q. Wang, Y. Lu, and J. Shu, "Octopus+: An rdma-enabled distributed persistent memory file system," *ACM Transactions on Storage*, 2021.

[89] S. Ghemawat, H. Gobioff, and S.-T. Leung, "The google file system," *Proceedings of the nineteenth ACM symposium on Operating systems principles.*, 2003.

[90] Hadoop, "HDFS Architecture Guide," https://hadoop.apache.org/docs/r1.2.1/hdfs_design.html, 2023.

[91] OpenNebula Systems, "OpenNebula Users," https://opennebula.io/users/, 2023.

[92] V. D. Long and T. N. B. Duong, "Group instance: Flexible co-location resistant virtual machine placement in IaaS clouds," in *2020 IEEE 29th international conference on enabling technologies: infrastructure for collaborative enterprises (WETICE)*, 2020.

[93] A. Agarwal and T. N. B. Duong, "Co-location resistant virtual machine placement in cloud data centers," in *IEEE ICPADS*, 2018.

[94] L. Tian, D. Feng, H. Jiang, K. Zhou, L. Zeng, J. Chen, Z. Wang, and Z. Song, "PRO: A popularity-based multi-threaded reconstruction optimization for RAID-structured storage systems." 2007.

[95] B. Mao, H. Jiang, D. Feng, S. Wu, J. Chen, L. Zeng, and L. Tian, "HPDA: A hybrid parity-based disk array for enhanced performance and reliability," in *IEEE IPDPS*, 2010.

[96] Intel, "Open Cache Acceleration Software," https://open-cas.github.io/.

[97] C. Multiphisics, "Comsol multiphisics," 2014.

[98] Microsoft, "Project Natick," https://natick.research.microsoft.com/.

[99] David Chen, "Pressure Vessels," https://processdesign.mccormick.northwestern.edu/index.php/Pressure_Vessels.

[100] I. O. Commission *et al.*, "The international thermodynamic equation of seawater, 2010: calculation and use of thermodynamic properties," 2010.

[101] A.R. Collins, "Underwater sound pressure levels," https://www.arc.id.au/SoundLevels.html.

[102] M. Ainslie and M. A. Ainslie, "Propagation of underwater sound," *Principles of sonar performance modelling*, 2010.

[103] "stress(1) - Linux man page," https://linux.die.net/man/1/stress.

[104] J. Killeen, I. Davis, J. Wang, and G. J. Bennett, "Fan-noise reduction of data centre telecommunications' server racks, with an acoustic metamaterial broadband, low-frequency sound-absorbing liner," *Applied Acoustics*, 2023.

[105] S. Wasala, L. Stevens, R. Sosseh, and T. Persoons, "Acoustic noise insulation for air-cooled data centre hard disk drive enclosures: Effect on thermal management," in *THERMINIC*, 2022.

[106] Y. Joshi and P. Kumar, "Introduction to data center energy flow and thermal management," in *Energy Efficient Thermal Management of Data Centers*, 2012.

[107] S. Ramamoorthy and S. Krishnan, "Towards thermal-acoustic co-design of noise-reducing heat sinks," *IEEE transactions on components, packaging and manufacturing technology*, 2018.

[108] L. S. Weilgart, "The impacts of anthropogenic ocean noise on cetaceans and implications for management," *Canadian journal of zoology*, 2007.

[109] L. Weilgart, "The impact of ocean noise pollution on fish and invertebrates," *Report for OceanCare, Switzerland*, 2018.

[110] D. F. Kune, J. Backes, S. S. Clark, D. Kramer, M. Reynolds, K. Fu, Y. Kim, and W. Xu, "Ghost talk: Mitigating emi signal injection attacks against analog sensors," in *IEEE SP*, 2013.

[111] T. Sugawara, B. Cyr, S. Rampazzi, D. Genkin, and K. Fu, "Light commands: Laser-based audio injection attacks on voice-controllable systems," in *USENIX Security*, 2020.

[112] J. Teoh, C. Du, G. Guo, and L. Xie, "Rejecting high frequency disturbances with disturbance observer and phase stabilized control," *Mechatronics*, 2008.

[113] J. Zhou, Q. Liu, Q. Jiang, W. Ren, K.-M. Lam, and W. Zhang, "Underwater camera: Improving visual perception via adaptive dark pixel prior and color correction," *IJCV*, 2023.

[114] A. Kulshreshtha and P. Shanmugam, "Estimation of underwater visibility in coastal and inland waters using remote sensing data," *Environmental monitoring and assessment*, 2017.

[115] K. Witowski and N. Stander, "Parameter identification of hysteretic models using partial curve mapping," in *AIAA ATIO and AIAA/ISSMO MA&O*, 2012.

[116] A. Cidon, R. Escriva, S. Katti, M. Rosenblum, and E. G. Sirer, "Tiered replication: A cost-effective alternative to full cluster geo-replication," in *USENIX ATC*, 2015.

[117] J. Liu, H. Shen, H. Chi, H. S. Narman, Y. Yang, L. Cheng, and W. Chung, "A low-cost multi-failure resilient replication scheme for high-data availability in cloud storage," *IEEE/ACM ToN*, 2021.

[118] C. Huang, H. Simitci, Y. Xu, A. Ogus, B. Calder, P. Gopalan, J. Li, and S. Yekhanin, "Erasure coding in windows azure storage," in *USENIX ATC*, 2012.

[119] Y. Hu, L. Cheng, Q. Yao, P. P. C. Lee, W. Wang, and W. Chen, "Exploiting combined locality for Wide-Stripe erasure coding in distributed storage," in *USENIX FAST*, 2021.

[120] J. C. W. Chan, Q. Ding, P. P. C. Lee, and H. H. W. Chan, "Parity logging with reserved space: Towards Efficient updates and recovery in erasure-coded clustered storage," in *USENIX FAST 14*, 2014.

[121] Microchip, "Choosing the right RAID configurations," https://www.microsemi.com/product-directory/raid-controllers/4047-raid-levels#12.

[122] R. Bobba, J. Grov, I. Gupta, S. Liu, J. Meseguer, P. C. Ölveczky, and S. Skeirik, "Survivability: design, formal modeling, and validation of cloud storage systems using maude," *Assured cloud computing*, 2018.

[123] S. Liu, J. Meseguer, P. C. Ölveczky, M. Zhang, and D. Basin, "Bridging the semantic gap between qualitative and quantitative models of distributed systems," *PACMPL*, 2022.

[124] Y. Cao, S. H. Bhupathiraju, P. Naghavi, T. Sugawara, Z. M. Mao, and S. Rampazzi, "You can't see me: Physical removal attacks on LiDAR-based autonomous vehicles driving frameworks," in *USENIX Security*, 2023.

[125] Y. Tu, S. Rampazzi, B. Hao, A. Rodriguez, K. Fu, and X. Hei, "Trick or heat? manipulating critical temperature-based control systems using rectification attacks," in *ACM CCS*, 2019.

[126] K. Wang, R. Mitev, C. Yan, X. Ji, A.-R. Sadeghi, and W. Xu, "GhostTouch: Targeted attacks on touchscreens without physical touch," in *USENIX Security*, 2022.

[127] M. Zuba, Z. Shi, Z. Peng, and J.-H. Cui, "Launching denial-of-service jamming attacks in underwater sensor networks," in *WUWNet*, 2011.

[128] M. Zuba, Z. Shi, Z. Peng, J.-H. Cui, and S. Zhou, "Vulnerabilities of underwater acoustic networks to denial-of-service jamming attacks," *ACM SACN*, 2015.

[129] Y. Dong, H. Dong, and G. Zhang, "Study on denial of service against underwater acoustic networks." *J. Commun.*, 2014.

[130] P. Xiao, M. Kowalski, D. McCulley, and M. Zuba, "An experimental study of jamming attacks in underwater acoustic communication," in *WUWNet*, 2015.

[131] T. Davenport, "Submarine cables, cybersecurity and international law: An intersectional analysis," *Cath. UJL & Tech*, 2015.

[132] R. Beckman, "Protecting submarine cables from intentional damage—the security gap," in *Submarine Cables*, 2014.

Appendix A.
Meta-Review

A.1. Summary

This paper analyzes the impact of modulated acoustic injection attacks against a submerged enclosure in a controlled lab setting and open-water scenarios. The paper conducts an extensive analysis of ultrasonic injection attacks on UDCs.

A.2. Scientific Contributions

- Independent Confirmation of Important Results with Limited Prior Research
- Identifies an Impactful Vulnerabilities

A.3. Reasons for Acceptance

1) The attack is of actual importance as more and more underwater data centers are deployed.
2) The authors follow a stringent scientific approach and have thought through their experiments. This becomes apparent in the precise representations and the structured execution of the experiments (e.g., volume, distance, orientation).
3) The lessons learned from this study will prompt better engineering of underwater structures and other countermeasures before large-scale deployment, which has not occurred yet.

A.4. Noteworthy Concerns

1) The paper studies ultrasonic injection in a valuable and realistic experimental setup but has not yet applied it to a real-world underwater data center. Real-world UDC might significantly differ in size, type of enclosure, deployment in rough or salty water, and other physical factors that might considerably impact the results of audio injection.
2) In a real scenario, the attacker might not have a feedback channel to tune the frequency and volume of the attack for a given distance.