HOPR as the transport layer of privacy-aware, GDPR-compliant IoT health systems

Lucas Benedicic, Martin Benedikt Busch, Sebastian Bürgel, Robert Kiel, Rich McDowell

September 2022, v0.3

Abstract

HOPR is a modern, permissionless, GDPR-compliant network for transport of private and confidential data. Implemented as a decentralized mixnet, HOPR provides network metadata protection and communication privacy, as well as compliance with data protection regulatory frameworks like GDPR. This work validates HOPR's compliance by presenting relevant GDPR mandates (e.g., articles and recitals) and evaluating them against HOPR features. As the practical case study for this validation, one of the most strictly regulated environments has been chosen: health data and their applications in an IoT setting. HOPR can provide a readily available, autonomous transport layer for private health data while fulfilling GDPR regulatory requirements, i.e., GDPR's stipulations for handling sensitive data. As such, HOPR positions itself at the intersection of technology, compliance, and healthcare, where new developments are literally lifesaving.

1 Introduction

In 2018, a report by the United Nations defined internet privacy as a fundamental human right [1]: "to which a person is inherently entitled simply because she or he is a human being". In May of the same year, the General Data Protection Regulation (GDPR) [2] came into force. GDPR is a central component of both privacy and human rights laws within the European Union (EU). GDPR applies to all organizations offering goods or services within the EU. GDPR also applies to all organizations processing EU residents' personal data, even if such processing occurs outside of Europe [3]. Organizations not complying with GDPR can be subject to fines of up to EUR 20 million, or four percent of their annual turnover, whichever is greater [4]. Switzerland, not being a member state of the EU, implements the so-called Federal Act on Data Protection (FADP) [5]. FADP is an essentially equivalent framework to GDPR, with few differences [6]. For example, in Switzerland fines are not imposed on legal entities but directly on individuals, who can be found liable up to the amount of CHF 250'000.

During recent years, the general availability of smartphones and autonomous sensors, as well as other Internet-of-Things (IoT) devices, has increased, along with their ability to collect, transmit, and store health data [7, 8]. Similarly, there has been an increase in remote healthcare systems which help doctors diagnose, monitor, and treat diseases by collecting data from indoor sensors and wearable devices in an IoT setting [9]. To be compliant with existing regulatory frameworks (e.g., GDPR, FADP), patient data should be anonymous, aggregated, and in unidentifiable form before being transmitted. Therefore, strong security and privacy mechanisms must be deployed to secure end-to-end channels between IoT devices (e.g., sensors, mobile phones) and the receiving end (e.g., hospital) .

This paper presents a case study of an indoor fall detection system, which is illustrative of many of the privacy challenges medical IoT projects face. The communication platform and network infrastructure used by IoT devices for collecting, sending, and exchanging patient data make GDPR compliance an even more complex task: data should be kept private and confidential not only at rest, but also while in transport [10]. In this context, we introduce a validation framework for HOPR as a readily available, GDPR-compliant transport layer based on relevant aspects of GDPR, the specifics of which are presented as user stories [11].

HOPR provides the technology breakthroughs needed for modern, private, and confidential transport of sensitive data [12] over public, insecure networks. By working at the transport layer, HOPR takes a novel privacy-by-design approach [13] and makes privacy the fundamental building block of today's internet stack.

The rest of this paper is organized as follows: Section 2 introduces the automated fall detection system and its compliance requirements from a GDPR perspective. An overview of the HOPR network and its security goals is presented in Section 3. The features of the HOPR network that fulfil GDPR compliance of the fall detection system while data is in transport, i.e., HOPR packet format, HOPR's implementation of cover traffic and HOPR's adaptable privacy, are described and validated in Sections 4, 5, and 6, respectively. Section 7 introduces a number of attacks and describes how HOPR mitigates them. Lastly, Section 8 summarizes the presented features and how they fulfill the GDPR regulatory requirements of the fall detection system.

2 An automated fall detection system

With a growing elderly population worldwide, falls are recognized as a major cause of physical, psychological, and economic concerns. A report from the World Health Organization indicates that adults over 60 suffer the greatest number of fatal falls [14].

Modern systems for real-time, automated detection of falls fall into two categories: wearable systems and context-aware systems [15]. The former category employs motion sensors that rely on kinematic signals, such as accelerometers and gyroscopes, where the user is required to carry a device. This may be uncomfortable or have a lower usability range. Context-aware systems deploy sensors to track the movement of people in limited environments, e.g., inside a house or a room. These systems are more usable for the elderly

as no dedicated device need be worn and coverage areas can be adjusted by changing the quantity and positioning of sensors. Context-aware systems leverage real-time sensor data to detect whether a user has fallen and request immediate assistance. However, privacy-related questions arise when systems track users in their personal residence, especially when video and audio capture devices are involved [16, 17].

2.1 Requirements for GDPR compliance

Several articles and recitals in GDPR apply to systems that collect personal data. This section presents specific GDPR requirements which either directly or indirectly relate to the communication platform and network infrastructure used by the fall detection system. These requirements are expressed as user stories, making validation of the regulatory context more specific and straightforward since it can be directly applied to features of the system under analysis. This also allows the system to assess and track its level of GDPR compliance, allowing users to take any necessary preventive actions following best risk-management practices.

Personas [11] are subjects in the user stories. Personas represent different roles and how they interact with the fall detection system. They are defined as follows:

- A **User** interacts directly with the fall detection system. GDPR stipulates that **Users**' data, as well as their metadata (e.g., IP address), must be treated as private, confidential, and sensitive in the context of medical and health applications [12].
- A **Data Custodian** takes care of aggregation, storage, and usage of data sets, focusing on the technical details of appropriate transport and storage of **Users**' personal data [18].
- A **Data Protection Officer (DPO)** ensures, in an independent manner, that an organization is correctly applying GDPR to protect **Users**' personal data [19].
- A **Supervisory Authority** is an independent public authority provided by each Member State and responsible for monitoring the application of GDPR [20].

Table 1 shows each of the user stories with its tag (column one from the left). These tags are used for reference in the sections that describe various HOPR features and security goals, in order to create a clear relationship between which feature satisfies which GDPR requirement. The second column refers to the GDPR aspect each user story represents, including references to related articles and recitals. The user story itself, highlighting the involved **persona**, appears in column three. The features of the HOPR network and the security goals that implement the GDPR compliance of specific user stories are shown in column four.

Tag	GDPR aspect	User story	HOPR features and Security goals
ust01	Data Protection: pseudonymization [21]	As a **Data Custodian**, the goal is to protect **User** data so that nobody can associate processed personal data with any specific **User**.	Sender anonymity, Receiver anonymity, Sender-Receiver unlinkability, Cover Traffic.
ust02	Data Protection: indirect identification [21]	As a **User**, the goal is to prevent the creation and usage of identifiers that can be associated with individuals when combined with their personal data (e.g., IP addresses).	Metadata protection.
ust03	Data Protection: indirect identification [21]	As a **User**, the goal is to be able to receive customized services from an application without being indirectly identified, so that individuals can protect their privacy.	Sender anonymity, Receiver anonymity, Sender-Receiver unlinkability, Cover Traffic.
ust04	Data protection: sensitive data [22]	As a **User**, the goal is for sensitive data to be protected so that individuals can have privacy and safety of their data.	Sender anonymity, Receiver anonymity, Sender-Receiver unlinkability, Cover Traffic.
ust05	Data protection: confidentiality [23, 10]	As a **Data Custodian**, the goal is to take adequate security measures so that the confidentiality of **Users**' data can be protected.	The use of HOPR as the transport layer of the system protects user's data while in transport.
ust06	Risk management: security risk management [10]	As a **DPO**, the goal is to have security reports so that plans for mitigating security risks can be developed.	Shorter paths, Adaptable privacy.
ust07	Compliance: transfer [24, 25, 26, 27, 28]	As a **User**, the goal is to protect personal data during data transfer processes so that individuals can ensure the safety and security of their data.	The use of HOPR as the transport layer of the system protects user's data while in transport.

Table 1: GDPR requirements presented as user stories of the fall detection system.

2.2 System design

A sample sensor distribution diagram of a context-aware, fall detection system is presented in Figure 1. Each sensor is continuously sending data in order for the system to react as quickly as possible in case of a fall. Therefore, these data must be kept private and confidential within the sensors and the local sensor network even before being

transmitted over the internet. Following a privacy-by-design principle, such measures are already applied at the hardware level by incorporating local processing (pseudonymization to satisfy *ust01*) and privacy thresholds to maintain data anonymity, privacy, and avoid (re-)identification of individuals (satisfies *ust02* and *ust03*). Since sensors are themselves HOPR nodes, these measures provide privacy and confidentiality guarantees for user data even in the case of a home (WiFi) network being compromised by a malicious actor.

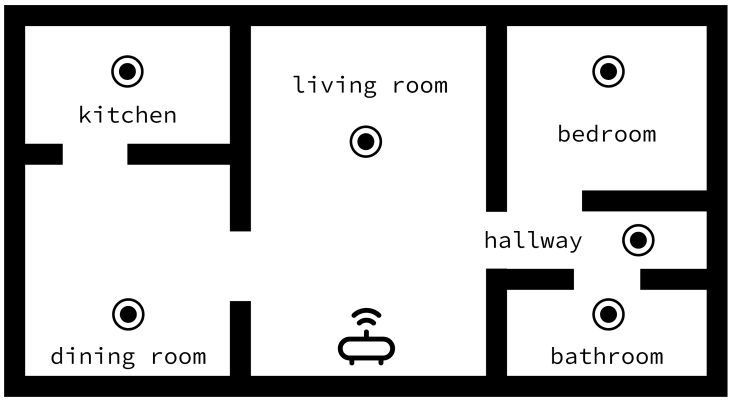

Figure 1: An example of sensor distribution within an apartment covered by the fall detection system.

Figure 2 shows a higher-level overview of the automated fall detection system, where simultaneous gathering of real-time data from several users is depicted. In this case, keeping traffic metadata (e.g., user's IP address) private and confidential is a requirement for the system to be GDPR compliant, thus an anonymous transport layer like HOPR is needed (see Anonymous Network Layer in Figure 2).

The server must also ensure the anonymity and privacy of user data while at rest. Thus, it only participates in the decentralized, anonymous transport layer as either a message sender or a message recipient. It follows that the transport layer of such system plays a fundamental role in the risk management approach as per GDPR, by minimizing the attack surface that could leak private and confidential information even in the case of a compromise (satisfies *ust06*).

3 What is HOPR?

HOPR is a way for people, companies, and devices to exchange information with complete privacy. People who communicate and transact using HOPR (or apps and services which run on top of the HOPR network) do so without leaking any information about what data is being shared, who is sending or receiving it, or even how much data is being sent. HOPR protects users' transport-level metadata privacy (e.g., IP addresses) by combining effective mechanisms to promote network growth, reliability and resilience.

HOPR is a completely decentralized mixnet [29] that employs privacy-by-design protocols, which makes it trustless, permissionless, and transparent. By never relying on a third party, users are not locked into a service provider or have to give up control of their data.

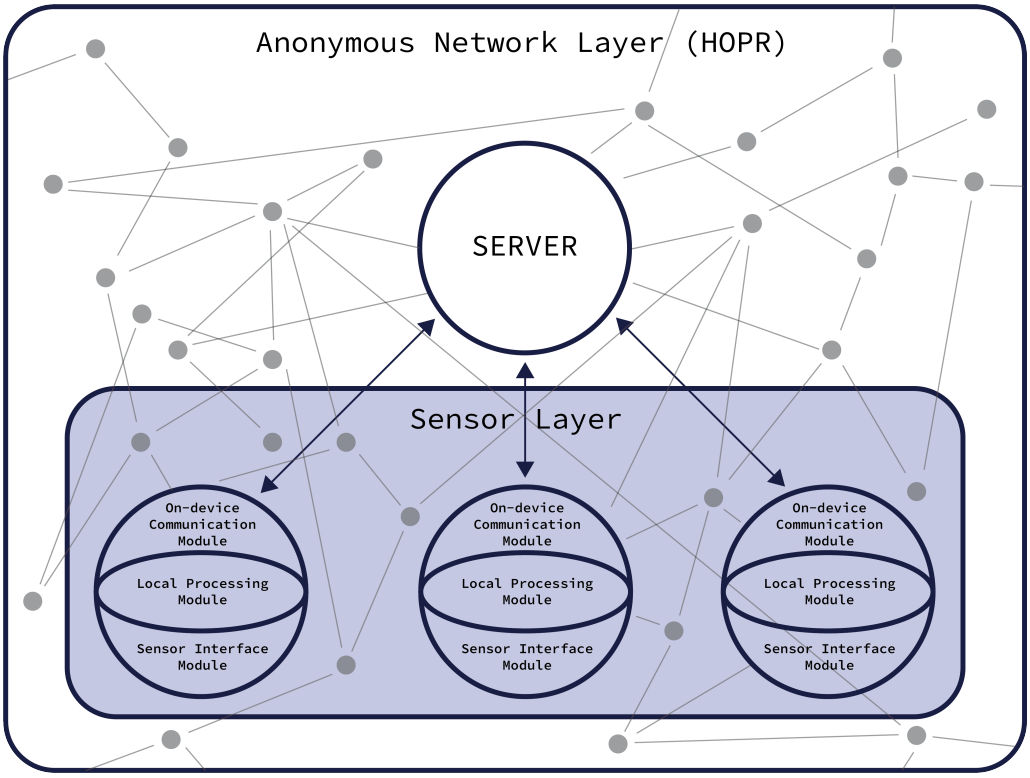

Figure 2: Architecture of a typical IoT health system, where HOPR represents the Anonymous Network Layer.

HOPR is run by users, for users. Anyone can become a member of the HOPR network by relaying data via a HOPR node, either on their computer or as a separate device plugged into their home router (e.g., the sensors of the fall detection system shown in Figure 1).

HOPR is also incentivized, meaning everyone who runs a node and stakes tokens will be rewarded for their efforts. The Ethereum blockchain [30] is used by HOPR to facilitate this incentive framework, specifically to perform probabilistic rewards via payment channels.

3.1 Security Goals

The HOPR protocol hides the fact that any two parties are communicating with each other, as well as the contents of their communication. This information should be hidden from any party external to the network (i.e., a global passive adversary), all intermediaries involved in transferring the data between the two communicating parties, and even the communicating parties themselves. Specifically, the HOPR protocol aims at building a network with distinct anonymity properties for a sender (Alejandro) and a recipient (Betty). Alejandro may encode and send an anonymous message destined for Betty. Betty cannot know that the originator of the message is Alejandro, but he can send another anonymous message using an address included in the message he just received. Betty's reply message is then routed through the network until it reaches Alejandro. In both cases Alejandro benefits from anonymity properties, first as the sender of the message (sender anonymity) and the second time as the anonymous receiver of a message (receiver

anonymity) [31]:

- **Sender anonymity.** In a network with sender anonymity, an observer is not able to tell whether a particular packet was sent by any adversary-selected honest senders A_1 or A_2 to the honest recipient Z. Formally this is denoted as $\{A_1 \rightarrow Z, A_2 \not\rightarrow\}$ or $\{A_1 \not\rightarrow, A_2 \rightarrow Z\}$.
- **Recipient anonymity.** In analogy to sender anonymity, here an observer is not able to tell whether a particular packet was received by any adversary-selected honest recipients Z_1 or Z_2 from the honest sender A. Formally this is denoted as $\{A \rightarrow Z_1, \not\rightarrow Z_2\}$ or $\{\not\rightarrow Z_1, A \rightarrow Z_2\}$.

To achieve these security goals, the HOPR protocol builds on top of the Sphinx packet format [32] which, among other features, also provides bitwise unlinkability, making it cryptographically difficult to link incoming and outgoing messages. Additionally, the HOPR protocol independently delays messages traversing the network, which makes the timings of packets unlinkable [33]:

- **Sender-recipient unlinkability.** For any pair of senders, A_1 and A_2, communicating with any pair of recipients, Z_1 and Z_2, an adversary must be unable to determine whether two packets travelled from $\{A_1 \rightarrow Z_1, A_2 \rightarrow Z_2\}$ or $\{A_1 \rightarrow Z_2, A_2 \rightarrow Z_1\}$. These properties assume senders and recipients are honest, but they should hold for any dishonest senders and recipients of the adversary's choice.

4 Packet Format

As mentioned in Section 3, HOPR is a decentralized mix network (or mixnet). Mixnets were introduced in 1981 by David Chaum [29] as a technique to achieve anonymous communications. They do so by relaying messages over a sequence of mix nodes, called the path. Each mix node receives a batch of encrypted messages, decrypts them, and sends them forward to follow their path. In this way, a mixnet provides anonymity to the users of the network, since an external actor observing the traffic flowing through a mix node is not able to link incoming and outgoing messages.

HOPR implements an augmented version of the Sphinx packet format, which arguably represents a de-facto standard for modern mixnets. The format determines how mixnet packets are created and transformed, while not leaking path information to any intermediate mix nodes or malicious actors eavesdropping on the communication.

In the context of the fall detection system, Sphinx encapsulates data packets from each of the sensors in Figure 1 (i.e., a HOPR node) to their final recipient (e.g., a HOPR node running as part of an application server, an alarm center or a mobile app), while providing the foundation for sender-recipient unlinkability [32], thus partially satisfying *ust01*, *ust03*, *ust04* and *ust05*.

The next section provides a high-level overview of the internal structure of a Sphinx packet. We then focus on providing the details on how HOPR fully satisfies the mentioned user stories and their GDPR requirements.

4.1 Structure

A Sphinx packet consists of two parts: a header and an onion-encrypted payload.

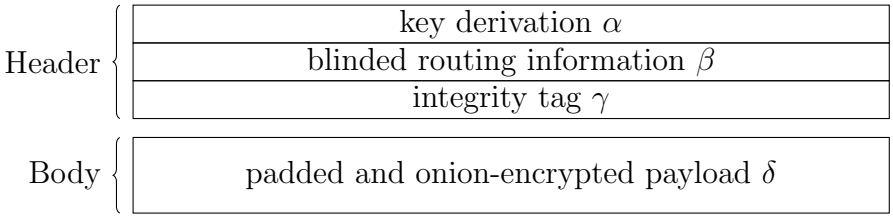

Figure 3: Schematic overview of a Sphinx packet.

A Sphinx packet is initially created by the sender choosing a path and deriving shared keys with each mix node along this path (key derivation α in Figure 3). The shared keys serve as a master secret to derive subkeys, which are used to blind the routing information (blinded routing information β in Figure 3). This is done in such a way that each of the involved mix nodes can solely determine the next downstream mix node.

In addition, the sender encrypts the message (payload δ in Figure 3) with one layer of encryption for each of the mix nodes along the chosen path. Once this is completed, the sender finally sends the packet to the first mix node.

Once a mix node receives the packet, it first derives the key that it has shared with the sender and checks its integrity (integrity tag γ in Figure 3). It then "unblinds" the routing information β to determine the next downstream mix node and removes one layer of encryption from the payload δ. At this point, the mix node is able to decide whether it is the final recipient of the message or whether it should forward the packet to the next mix node in the path by inspecting the blinded routing information β.

Note that this process preserves sender anonymity (satisfying *ust01*, *ust03* and *ust04* from Table 1), as well as preventing the sender's metadata (e.g., IP address) from being exposed when sending a message, thus also satisfying *ust02*.

For additional details about the construction and functioning of the Sphinx packet format, we refer the reader to the work by Danezis and Goldberg in [32].

4.2 Shorter Paths

The term "shorter path" refers to a path where the sender has chosen a lower number of mix nodes than stipulated (see Section 4.1). Creating a header for a shorter path requires less blinded routing information because there are less mix nodes to traverse along the

path. However, since β has a fixed size (see Figure 4), using shorter paths for some packet leaves empty space at the end of β.

Kuhn et al. [34] note that this behavior introduces observable patterns that deviate from normal usage, thus potentially allowing a malicious actor to reconstruct the route of distinguished packets. The HOPR packet format augments the original design of the Sphinx packet by padding β with random data so that no observable patterns emerge if shorter paths are used.

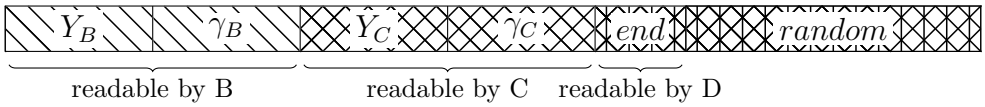

Figure 4: Sender node A has created routing information for intermediate mixnodes B and C, and the final recipient D, filling the empty part of β with random data.

Figure 4 shows a packet header containing the blinded routing information and the integrity tag for several mix nodes. In this case, node A (the sender) has created routing information for B, C (intermediate mixnodes), and D (the receiver). It can be seen that the rightmost portion of the header is left unused and filled with random data to avoid leaking any information.

The longer the path, the higher the privacy at the expense of latency and bandwidth overhead. The use of shorter paths allows applications like the fall detection system to independently adjust the required levels of latency and bandwidth.

This HOPR feature gives developers complete freedom in terms of fulfilling application requirements while keeping privacy considerations under control (partially satisfies *ust06* from Table 1).

4.3 Validation of user stories

To validate how the HOPR packet format (see Section 4) satisfies user stories *ust01*, *ust03*, *ust04*, *ust05* and *ust07*, we focus on the security properties of well-known mix protocols, e.g., onion routing. These security properties have been extensively studied and documented in the related literature [32, 35, 36]. In this context, the relevant security properties are: correctness, integrity, wrap resistance, and security. Together as part of any onion routing protocol they realize ideal functionality in the Universal Composability model [37]. It follows that a malicious actor has no access to the underlying cryptographic implementation of the protocol, but can only observe opaque identifiers for messages. This holds true even when the malicious actor controls some of the mix nodes along the observed path and may change the protocol itself [32].

In terms of HOPR's underlying Sphinx packet format, these four security properties can be verified as follows:

- **Correctness.** By simply inspecting the protocol, it is straightforward to verify

that it works correctly in the absence of an active adversary. The HOPR protocol assumes that its participants behave *rationally*, meaning being *honest* but greedy.

- **Integrity.** Each mixnet packet encodes the path that it is going to take through the mixnet. A malicious actor is neither able to exchange nodes along the path nor change the length of the path, i.e., make the packet travel an additional hop.
- **Wrap resistance.** Given a mixnet packet p, a malicious actor is unable to find a mixnet packet p' whose transformation by an honest mixnet node leads to p.
- **Security and Indistinguishability.** A malicious actor controlling all nodes but one cannot determine, when given a mixnet packet, whether it will travel through their node or not. This is known as the trust assumption [32].

Within the principles of the Universal Composability model [37], the formal proof of all four security properties can be found in [32].

For HOPR as the transport layer of the fall detection system, some additional security properties also apply. Referring to the architecture diagram in Figure 2, it is clear that we are in the presence of a decentralized network, since all involved elements are HOPR nodes themselves (e.g., the sensors, the application server). Thus all exchanged messages are cryptographically indistinguishable at the protocol level and no information is leaked along the path. Note that there are no relay nodes facing the public internet, hence drastically reducing the attack surface.

5 Cover Traffic

Cover traffic (CT) is an essential building block for HOPR and other systems [38, 39] to ensure anonymous communication. CT blends random messages (noise) into the HOPR mixnet to obfuscate *real* communication traffic. Additionally, the fact that both message types have the same packet format (presented in Section 4) makes them indistinguishable. This is how HOPR protects communication traffic from global passive adversaries who observe the mixnet in order to discriminate between communication patterns and compromise anonymity [40]. Especially in the early roll-out stages of the HOPR network, CT is critical to preventing attacks [41, 42, 43] due to the expected low amount of *real* communication traffic.

CT comes at a cost of a trade-off between strong anonymity, low latency, and low bandwidth overhead. This trade-off is known as the *Anonymity Trilemma* stating that only two out of the three properties are achievable at the same time [44]. In general, latency refers to the amount of time it takes a message to reach its receiver, while bandwidth refers to the amount of data a network can transmit in a given timeframe. In our application of HOPR to a real-time fall detection system, where low latency is expected, we only face a trade-off between strong anonymity and low bandwidth overhead. This trade-off has received some treatment in the literature [45, 46]. For example, strong anonymity can be achieved by utilizing all network nodes for CT at all times [45]. However, such strong

anonymity comes at the expense of higher bandwidth overhead (i.e., high cost of running the network) meaning that each *real* message is hidden behind multiple CT messages.

HOPR is a project focusing on privacy, but it should also be able to scale and be accessible to everyone. Therefore, it is important to find a good balance between strong anonymity and the deleterious effect of CT for network bandwidth. In particular, Grube et al. show that the share of nodes emitting CT and the frequency at which nodes relay CT are the two most important parameters affecting the trade-off between strong anonymity and a low bandwidth overhead [45]. They show that efficiency improvements of up to 50% in terms of network bandwidth are possible, while keeping anonymity at a high level. This is an important finding as CT nodes will initially be run by the HOPR Association (or third parties sponsored by the HOPR Association) for the purpose of generating CT. Going forward, the goal is that anyone in the HOPR network will be able emit CT, relay it, and receive a reward for doing so.

The implementation of CT into the mixnet reinforces HOPR's security goals of sender anonymity, receiver anonymity, and sender-receiver unlinkability. This therefore satisfies the GDPR requirements represented by user stories *ust01*, *ust03*, *ust04*, *ust05* and *ust07* - assuring a private transport of health data.

6 Adaptable Privacy

Analogous to the compromise introduced in Section 5, anonymous communication is set as a lower bound in terms of latency and bandwidth overhead due to the consequences of the Anonymity Trilemma [44]. The degree of anonymity depends on specific application requirements and it varies from case to case, even within the same application. For example, in case of an emergency triggered by the fall detection system (e.g., the detection of a severe fall), a message may have to traverse the network as quickly as possible with lower privacy guarantees since faster delivery of messages can save lives. On the other hand, regular status updates broadcast by various sensors (i.e., HOPR nodes) may have more flexible latency requirements, but require higher privacy in order to remain compliant with GDPR and avoid leaking users' personal patterns during daily system use.

HOPR provides a framework that gives the application developer the means to individually adjust the privacy level of messages, hence trading provable anonymity for lower latency or reduced bandwidth overhead. Within the HOPR network, this is achieved by specifying a custom implementation for two mechanisms that directly affect latency and bandwidth: the path selection algorithms, and the strategy module that defines where to stake HOPR tokens.

To monitor its performance and continuously enhance its privacy, the HOPR network will publish an anonymity score that relates to the achieved privacy of HOPR nodes. This empowers application developers to enforce different requirements in terms of privacy demands to various aspects of an application. The adaptability available at the protocol level covers several aspects, namely:

- **Number of hops:** Higher numbers of hops in a path, provide a higher degree of privacy. Note, however, that adding more hops increases the package header overhead for all mixnodes in the network.
- **Added latency per hop:** Similar to the previous point, higher latency results in higher levels of privacy.
- **Amount of CT:** As mentioned in Section 5, added noise will increase the achieved privacy level, but it will leave less bandwidth available for real traffic.

Note that this feature satisfies *ust06* from Table 1, since the anonymity score provides the **DPO** with a continuous security report that supports the creation and adaptability of risk management policies.

7 Threat Model

In this section, we introduce potential threats to the HOPR network. Since threats can create gaps in GDPR compliance (see Table 1), we show how the features of the HOPR network mitigate them. We assume a threat model with a global passive adversary (GPA) with the ability to either observe all network traffic and launch network attacks or to inject, drop, or delay packets.

7.1 Global passive adversaries

A GPA is an attacker who can observe the entirety of network traffic passing between nodes. A GPA is considered passive because their attacks are based on observation alone. A theoretical GPA is arbitrarily powerful, and their full abilities are unlikely to be manifest in any real-world attacker. Nonetheless, building a network which is GPA-resistant introduces extra security through redundancy and future proofing.

Due to the properties of the HOPR mixnet, the HOPR packet format (see Section 4), and CT (see Section 5), HOPR is able to defend against GPAs.

7.2 Intersection attacks

In an intersection attack, an attacker observes the network and repeatedly gathers information about which nodes participate in message communication or exhibit distinguishable communication behavior. The attacker then *intersects* the sets of participating nodes, collected at different points in time, to determine communication relations between nodes [38, 40, 43].

HOPR mitigates the risk of intersection attacks through the use of the HOPR packet format and CT, making it indistinguishable from real communication traffic (see Section 4 and 5, respectively). In essence, CT increases the set of nodes that participate in

communication at any given time. Hence, it is nearly impossible for a GPA to reduce the set of participating nodes and determine communication relations by simply intersecting sets of nodes. This again highlights that CT is critical in the early rollout stages of the HOPR network, where *real* communication traffic might be low, to specifically prevent intersection attacks.

7.3 Sybil attacks

In a Sybil attack, an attacker forges multiple identities in the network, thereby introducing network redundancy and reducing system security. The attacker can potentially de-anonymize packet traffic and thus link the sender's and recipient's identities [47].

HOPR mitigates Sybil attacks via the trust assumption of the Sphinx packet format (see Section 4 and the Integrity property in Section 4.3). The trust assumption states that only a single honest mix node is needed to ensure integrity of the entire transmission path [32]. Since communication traffic is source routed, message senders can choose routes themselves to ensure that this minimal requirement of one honest mix node is met.

7.4 Eclipse attacks

Many decentralized networks suffer from a general misunderstanding of their topology, hence nodes cannot determine whether their respective local views are complete or accurate. As an attacker, it might be attractive to flood a victim with inaccurate information about collaborating nodes while withholding information about honest nodes.

HOPR mitigates this issue by announcing entry nodes to the network using a medium that is append-only and permissionless: a smart contract on a blockchain. This feature is known within HOPR as DEADR (Decentralized Entry Advertisement and Distributed Relaying). Among other benefits, DEADR facilitates network address translation (NAT) for relaying traffic that cannot be sent directly between adjacent nodes, the most common example of which are computers behind a home router and without a public internet-facing internet protocol address (IP address).

Since a HOPR node only requires access to one honest DEADR node, it is significantly cheaper to check whether a DEADR node is honest than to conduct an eclipse attack on the network, which inevitably requires a malicious actor to yield multiple on-chain transactions. As such, HOPR inherits the security guarantees from the Ethereum blockchain and the specific Ethereum node that a HOPR node is connected to.

8 Conclusion

This paper introduced HOPR: a modern, readily available, GDPR-compliant mixnet which provides privacy and confidentiality of user data while in transport. HOPR independently manages the complexity and challenges of keeping data in transport private and confidential, while fulfilling GDPR regulatory requirements. The case study of an automated fall detection system serves as a reference to illustrate the strict compliance framework GDPR imposes to typical IoT projects dealing with health data. In this context, several features and security goals of HOPR (sender anonymity, receiver anonymity, sender-receiver unlinkability, packet format, cover traffic, and adaptable privacy) were introduced. GDPR compliance was validated against several GDPR articles, which were articulated as user stories and personas within the reference fall detection system.

As an autonomous network, HOPR enhances the privacy of typical IoT projects without introducing additional processes. Being a closed network, it reduces the attack surface that external malicious actors can exploit, while not significantly increasing overall system complexity. By implementing features like cover traffic, HOPR is even capable of hiding whether it ever communicated or sent any meaningful message other that cover traffic itself. Among other privacy-enhancing capabilities, HOPR also provides metadata protection so that the IP addresses of the communicating parties are kept private.

Since security is generally about finding a compromise between privacy, latency, and bandwidth, HOPR provides features like shorter paths and adaptable privacy. Their use in the context of health applications like the fall detection system allow application developers to independently adjust the levels of latency and bandwidth overhead, thus fulfilling application requirements while keeping privacy considerations under control. This is achieved with virtually zero changes to the project's code base and in full compliance with GDPR requirements.

References

[1] UN. The right to privacy in the digital age : report of the united nations high commissioner for human rights 17 p. (2018). http://digitallibrary.un.org/record/1640588.

[2] Wikipedia. General Data Protection Regulation — Wikipedia, the free encyclopedia. http://en.wikipedia.org/w/index.php?title=General%20Data%20Protection%20Regulation&oldid=1096269637 (2022). [Online; accessed 10-July-2022].

[3] European Commission. GDPR, Article 3: Territorial Scope. https://gdpr-info.eu/art-3-gdpr/ (2016). [Online; accessed 10-July-2022].

[4] European Commission. GDPR, Article 83: Right to compensation and liability. https://gdpr-info.eu/art-83-gdpr/ (2016). [Online; accessed 10-July-2022].

[5] Swiss Federal Council. New Federal Act on Data Protection (nFADP). https://www.kmu.admin.ch/kmu/en/home/facts-and-trends/digitization/data-protection/new-federal-act-on-data-protection-nfadp.html (2020). [Online; accessed 10-July-2022].

[6] Register, G. New Swiss Data Protection Act completed: comparison with GDPR. https://www.gdprregister.eu/news/swiss-data-protection-act-2020/. [Online; accessed 10-July-2022].

[7] Shahzad, A. & Kim, K. Falldroid: An automated smart-phone-based fall detection system using multiple kernel learning. *IEEE Transactions on Industrial Informatics* **15**, 35–44 (2018).

[8] Moosavi, S. R. *et al.* End-to-end security scheme for mobility enabled healthcare internet of things. *Future Generation Computer Systems* **64**, 108–124 (2016).

[9] Marin, E., Mustafa, M. A., Singelée, D. & Preneel, B. A privacy-preserving remote healthcare system offering end-to-end security. In *International Conference on Ad-Hoc Networks and Wireless*, 237–250 (Springer, 2016).

[10] European Commission. GDPR, Article 32: Security of processing. https://gdpr-info.eu/art-32-gdpr/ (2016). [Online; accessed 10-July-2022].

[11] Ahola, J. *et al.* Handbook of the secure agile software development life cycle (2014).

[12] European Commission. Sensitive data. https://ec.europa.eu/info/law/law-topic/data-protection/reform/rules-business-and-organisations/legal-grounds-processing-data/sensitive-data_en (2018). [Online; accessed 10-July-2022].

[13] Wikipedia. Privacy by design — Wikipedia, the free encyclopedia. https://en.wikipedia.org/wiki/Privacy_by_design (2022). [Online; accessed 10-July-2022].

[14] World Health Organization. Falls. https://www.who.int/news-room/fact-sheets/detail/falls (2021). [Online; accessed 10-July-2022].

[15] Igual, R., Medrano, C. & Plaza, I. Challenges, issues and trends in fall detection systems. *Biomedical engineering online* **12**, 1–24 (2013).

[16] Asif, U. *et al.* Privacy preserving human fall detection using video data. In *Machine Learning for Health Workshop*, 39–51 (PMLR, 2020).

[17] Edgcomb, A. & Vahid, F. Automated fall detection on privacy-enhanced video. In *2012 Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, 252–255 (IEEE, 2012).

[18] Wikipedia. Data custodian — Wikipedia, the free encyclopedia. https://en.wikipedia.org/wiki/Data_custodian (2022). [Online; accessed 10-July-2022].

[19] Wikipedia. Data Protection Officer — Wikipedia, the free encyclopedia. https://en.wikipedia.org/wiki/Data_protection_officer (2022). [Online; accessed 10-July-2022].

[20] European Commission. GDPR, Article 51: Supervisory Authority. https://gdpr-info.eu/art-51-gdpr/ (2016). [Online; accessed 10-July-2022].

[21] European Commission. GDPR, Article 4: Definitions. https://gdpr-info.eu/art-4-gdpr/ (2016). [Online; accessed 10-July-2022].

[22] European Commission. GDPR, Article 9: Processing of special categories of personal data. https://gdpr-info.eu/art-9-gdpr/ (2016). [Online; accessed 10-July-2022].

[23] European Commission. GDPR, Article 25: Data protection by design and by default. https://gdpr-info.eu/art-25-gdpr/ (2016). [Online; accessed 10-July-2022].

[24] European Commission. GDPR, Article 44: General principle for transfers. https://gdpr-info.eu/art-44-gdpr/ (2016). [Online; accessed 10-July-2022].

[25] European Commission. GDPR, Article 45: Transfers on the basis of an adequacy decision. https://gdpr-info.eu/art-45-gdpr/ (2016). [Online; accessed 10-July-2022].

[26] European Commission. GDPR, Article 46: Transfers subject to appropriate safeguards. https://gdpr-info.eu/art-46-gdpr/ (2016). [Online; accessed 10-July-2022].

[27] European Commission. GDPR, Article 47: Binding corporate rules. https://gdpr-info.eu/art-47-gdpr/ (2016). [Online; accessed 10-July-2022].

[28] European Commission. GDPR, Article 49: Derogations for specific situations. https://gdpr-info.eu/art-49-gdpr/ (2016). [Online; accessed 10-July-2022].

[29] Chaum, D. Untraceable electronic mail, return addresses, and digital pseudonyms. *Commun. ACM* **24**, 84–88 (1981). http://doi.acm.org/10.1145/358549.358563.

[30] Wood, G. Ethereum: A secure decentralised generalised transaction ledger (istanbul version) (2021). https://ethereum.github.io/yellowpaper/paper.pdf.

[31] Danezis, G., Dingledine, R. & Mathewson, N. Mixminion: Design of a type iii anonymous remailer protocol. In *2003 Symposium on Security and Privacy, 2003.*, 2–15 (IEEE, 2003).

[32] Danezis, G. & Goldberg, I. Sphinx: A compact and provably secure mix format. In *30th IEEE Symposium on Security and Privacy (S&P 2009), 17-20 May 2009, Oakland, California, USA*, 269–282 (IEEE Computer Society, 2009). https://doi.org/10.1109/SP.2009.15.

[33] Piotrowska, A. M., Hayes, J., Elahi, T., Meiser, S. & Danezis, G. The loopix anonymity system. *CoRR* **abs/1703.00536** (2017). http://arxiv.org/abs/1703.00536, 1703.00536.

[34] Kuhn, C., Beck, M. & Strufe, T. Breaking and (partially) fixing provably secure onion routing. *CoRR* **abs/1910.13772** (2019). URL http://arxiv.org/abs/1910.13772. 1910.13772.

[35] Camenisch, J. & Lysyanskaya, A. A formal treatment of onion routing. In *Annual International Cryptology Conference*, 169–187 (Springer, 2005).

[36] Feigenbaum, J., Johnson, A. & Syverson, P. A model of onion routing with provable anonymity. In *International Conference on Financial Cryptography and Data Security*, 57–71 (Springer, 2007).

[37] Canetti, R. Universally composable security. *Journal of the ACM (JACM)* **67**, 1–94 (2020).

[38] Berthold, O., Federrath, H. & Köhntopp, M. Project "anonymity and unobservability in the internet". In *Proceedings of the tenth conference on Computers, freedom and privacy: challenging the assumptions*, 57–65 (2000).

[39] Van Den Hooff, J., Lazar, D., Zaharia, M. & Zeldovich, N. Vuvuzela: Scalable private messaging resistant to traffic analysis. In *Proceedings of the 25th Symposium on Operating Systems Principles*, 137–152 (2015).

[40] Raymond, J.-F. Traffic analysis: Protocols, attacks, design issues, and open problems. In *Designing privacy enhancing technologies*, 10–29 (Springer, 2001).

[41] Mallesh, N. & Wright, M. Countering statistical disclosure with receiver-bound cover traffic. In *European Symposium On Research In Computer Security*, 547–562 (Springer, 2007).

[42] Levine, B. N., Reiter, M. K., Wang, C. & Wright, M. Timing attacks in low-latency mix systems. In *International Conference on Financial Cryptography*, 251–265 (Springer, 2004).

[43] Berthold, O. & Langos, H. Dummy traffic against long term intersection attacks. In *International Workshop on Privacy Enhancing Technologies*, 110–128 (Springer, 2002).

[44] Das, D., Meiser, S., Mohammadi, E. & Kate, A. Anonymity trilemma: Strong anonymity, low bandwidth overhead, low latency - choose two. In *2018 IEEE Symposium on Security and Privacy, SP 2018, Proceedings, 21-23 May 2018, San Francisco, California, USA*, 108–126 (IEEE Computer Society, 2018). https://doi.org/10.1109/SP.2018.00011.

[45] Grube, T., Thummerer, M., Daubert, J. & Mühlhäuser, M. Cover traffic: A trade of anonymity and efficiency. In *International Workshop on Security and Trust Management*, 213–223 (Springer, 2017).

[46] Grube, T., Daubert, J. & Muhlhauser, M. Asymmetric dcnets for effective and efficient sender anonymity. In *2018 IEEE Global Communications Conference (GLOBECOM)*, 1–7 (IEEE, 2018).

[47] Wikipedia. Sybil attack — Wikipedia, the free encyclopedia. https://en.wikipedia.org/wiki/Sybil_attacks (2022). [Online; accessed 10-July-2022].