

Received August 5, 2020, accepted September 20, 2020, date of publication September 28, 2020, date of current version October 9, 2020.

Digital Object Identifier 10.1109/ACCESS.2020.3027150

Reinforced Secure Gossiping Against DoS Attacks in Post-Disaster Scenarios

CHRISTIAN ESPOSITO¹, (Member, IEEE), ZHONGLIANG ZHAO^{2,3}, (Member, IEEE), AND JACEK RAK⁴, (Senior Member, IEEE)

¹Department of Computer Science, University of Salerno, 84084 Fisciano, Italy

²School of Electronic Information Engineering, Beihang University, Beijing 100083, China

³Institute of Computer Science, Universität Bern, 3012 Bern, Switzerland

⁴Faculty of Electronics, Telecommunications, and Informatics, Gdańsk University of Technology (GUT), 80-233 Gdańsk, Poland

Corresponding authors: Zhongliang Zhao (zhaozl@buaa.edu.cn) and Christian Esposito (esposito@unisa.it)

This work was supported by the COST Association (European Cooperation in Science and Technology).

ABSTRACT During and after a disaster, the perceived quality of communication networks often becomes remarkably degraded with an increased ratio of packet losses due to physical damages of the networking equipment, disturbance to the radio frequency signals, continuous reconfiguration of the routing tables, or sudden spikes of the network traffic, e.g., caused by the increased user activity in a post-disaster period. Several techniques have been introduced so far (mainly using data retransmission mechanisms) to tolerate such circumstances. Among them, gossiping has been shown to be efficient in the recovery from message losses. However, a conventional gossiping scheme may exhibit security problems, which can be exploited for further attacks (such as Denial of Service – DoS attack). For instance, the flooding method used by the gossiping can be used to forward the traffic towards many vulnerable nodes to drain their resources and compromise them. Typically, protection against DoS attacks is realized by using cryptographic primitives. However, their scalability limits and costs make them improper for emergency communications after a disaster. In this article, we introduce an approach based on reinforcement learning and game theory to protect the gossiping scheme from DoS attacks without incurring the costs of cryptographic primitives. In our method, nodes properly select which requests to satisfy, which in turn helps other nodes to avoid receiving manipulated gossip messages from malicious and colluded nodes. Additionally, our method operates without exploiting any cryptographic primitives, which prevents excessive energy waste that is undesired in post-disaster resilient networking. Simulation experiments performed in OMNeT++ confirmed the advantages of our approach over the reference schemes in terms of reliability, security, overhead, latency, and power efficiency.

INDEX TERMS Communication system security, disaster-resilience, gossiping algorithm, game theory, reinforcement learning.

I. INTRODUCTION

The pervasive accessibility to the Internet and the availability of cheap and small computing devices are making information and communication technologies (ICT) predominant in our daily life. We have been already living in a digital society thanks to novel concepts such as the Internet of Things, smart cities, or Industry 4.0 [1]. Digital technologies have revolutionized communications among humans, institutional relationships, and the way our society operates. All the main actors from individuals or groups to companies and

institutions strongly rely on computer networks to support their aims or businesses. Typically, the Internet can offer services at a sufficient quality of service (QoS) level. However, numerous massive failure scenarios, often called *disasters*, may compromise the performance of the network to the point that it becomes unavailable or offers services of a remarkably degraded quality [2], [3]. Following [4], disaster events in communication networks can be broadly classified into three categories, namely natural disasters, malicious attacks, and technology-related disasters, shown in Fig. 1.

In particular, *natural disasters* refer to the adverse effects of the forces of nature. They include predictable events such as floods, fires, volcano eruptions, or hurricanes [4]

The associate editor coordinating the review of this manuscript and approving it for publication was Baoping Cai¹.

Major Events Causing Massive Failures in Communication Networks




FIGURE 1. Disaster events affecting communication networks based on [4].

(e.g., hurricane Katrina in 2005 in the US responsible for long-lasting massive failures of network nodes due to power outages for over ten days, on average [5], [6]). Unpredictable disasters, in turn, comprise earthquakes for which the precise location and time cannot be foreseen despite the availability of statistical data on past disasters. For example, the 2011 Greatest Japan Earthquake of 9.0 magnitude resulted in failures of multiple undersea optical links and about 1500 telecom switching offices [7], [8]. As natural disasters commonly occur at specific locations, they lead to the so-called *regional failures* of multiple network elements located in a given area [8]–[11], as shown in Fig. 2.

Malicious attacks are human activities affecting communication networks either directly (see, e.g., electromagnetic pulse attack (EMP), and denial of service (DoS) attack [3], [4]), or indirectly (such as bombing or use of weapons of mass destruction (WMD)). Their impact on network performance may be severe, and their risk is rising [12]. Contrary to natural disasters, failures of network elements following from technology-related disasters and malicious attacks, instead of being confined in a specific region, are commonly spread across the network. In particular, the location of nodes failed due to software issues can be often considered as random. Also, attackers are mainly interested in compromising the network elements not necessarily located only in a given region but playing a significant role in the entire network (such as links of high capacity/nodes switching a large amount of data or acting as information servers).

Massive failures in communication networks can also be a result of *technology-related disasters* due to internal issues such as software errors and hardware faults at multiple locations [13]. External events of power supply faults can, in turn, lead to cascading failures in interdependent networks [14]. For instance, in the scenario of inter-dependence between a power grid and a communication network, a failure of even a single node in one system (e.g., a power grid providing power supply to a communication network) may switch off many nodes in a communication network. If this communication network provides the control functions to the power grid, failures are likely to be propagated back to the power grid, implying even a total collapse of both systems.

Severe disasters can also cause temporary unavailability of links when the damage is not physical and can be automatically recovered through the reconfiguration of routers,




FIGURE 2. Example failure region characteristic to natural disasters.

as shown in [15], possibly through software-based mechanisms and network reconfiguration [13]. Permanent failures are much more challenging to handle, as component substitution is required to revert the full link availability. Moreover, a disaster can, temporarily or permanently, damage also data centres hosting the core services or cloud services. As disasters often affect people, communication networks are viewed as an essential part of the critical infrastructure [16] allowing us to communicate and providing emergency communications between citizens and authorities. Therefore, in a post-disaster period, we can often observe a substantial increase in the network traffic demand volume, which is hard to serve in a network already affected by a disaster event and contributes to network saturation and service unavailability.

In a disaster scenario, communication availability is also essential for rescue team members to coordinate their activities and for people in the affected areas to receive the rescue information. However, it may not be feasible in a network not enhanced with disaster-resilience features [17]. As a concrete example, in a wildfire in Portugal in June 2017, a large number of users were cut off from using fixed-line and cellular communication services [18]. It led to remarkable traffic congestion in the isolated areas of the network, as well as affected the emergency communications among rescue teams, which, in turn, caused a large number of casualties. These results have shown that an efficient and reliable notification mechanism is also necessary to build collaborative crisis management or emergency networks by local authorities and dedicated civil protection organizations [19]. Since disasters are increasing in frequency and scale, and their impact on communication networks is non-negligible [20], it is crucial to design novel mechanisms to sustain the system availability to operate in a post-disaster period. In particular, it is of utmost importance to assure reliable delivery of messages in emergency scenarios despite the possible occurrence of malicious activities in a communication network already affected by a disaster.

As disasters often affect communication networks, dissemination of messages in a post-disaster period can be provided by setting up the alternative communication paths either proactively (i.e., before the disaster occurrence) or in a reactive manner by re-configuring the network after a disaster. In particular, as discussed in this article, a scheme called gossiping (a distributed scheme of message dissemination

where multiple copies of information are forwarded towards the destination via multiple paths) is considered as a proper solution in many disaster scenarios for message dissemination. However, as gossiping exhibits vulnerabilities which can be exploited in further attacks (or to compromise the end-user applications built on top of gossiping), and since the presence of malicious nodes in a post-disaster period is not a rare event [35], [36], a proper functioning of gossiping still requires certain adjustments.

A. MAIN CONTRIBUTIONS

In this article, we focus on the application of the gossiping scheme for reliable communications in a post-disaster period (i.e., when multiple elements of the network become severely affected by a disaster). A particular interest in this article is to protect gossiping against DoS attacks in the network already affected by a disaster. In this context, the achievements of this article are threefold and include:

- (1) The proposal of the gossiping scheme providing protection against DoS attacks without using any digital signatures to guarantee authorship and message integrity. Our solution leverages on game theory to countermeasure DoS attacks by modelling the node and request selection at each gossiping round as a non-cooperative game, where strategy and payoff mapping is implemented by means of reinforcement learning and myopic best response dynamics are used to converge towards an optimum. The novelty is in adapting a solution for congestion-control in gossiping to provide protection against DoS attacks without the need of using the cryptographic primitives with too high costs for emergency networks. The formulated utility function and reinforcement learning allows distinguishing between the traffic caused by the execution of DoS attacks and traffic from legitimate users sending urgent messages simultaneously, thanks to the feedback on the utility of received gossip messages.
- (2) The usability of the proposed solution in post-disaster periods following the occurrence of all major scenarios of disasters. Indeed, as our gossiping scheme does not depend on any centralized authority needed for identification, trust, and certificate management (but makes decisions in a distributed manner in a non-cooperative game), it fits well the post-disaster communications often being data-centric with anonymous data transmissions.
- (3) The performance assessment of our gossiping scheme providing a broad set of scenarios related to different sizes of disasters and scopes of DoS attacks to confirm its high reliability, security, scalability, and efficiency, as well as low overhead, compared to the conventional gossiping scheme.

An essential advantage of our approach is its simplicity in terms of implementation and compatibility with other components either at the network, or at the application level. Our secure gossiping scheme can be widely used to provide

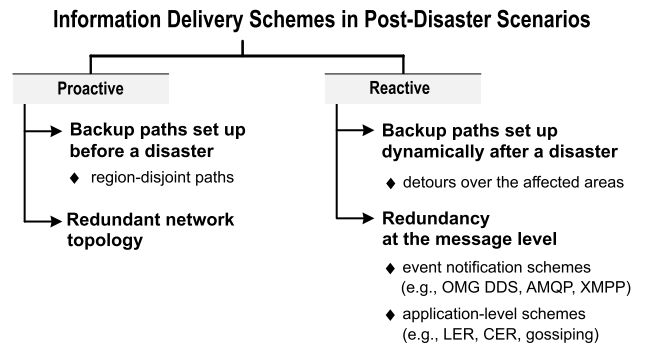


FIGURE 3. Classification of information delivery schemes in a post-disaster scenario.

recovery of messages in a post-disaster period for applications using both the publish-and-subscribe strategy (e.g., for communications between authorities and end-users) and for direct communications among users. The condition for a correct malicious nodes detection is that a DoS attack is conducted only by forging valid identifies to issue more requests than needed on behalf of legitimate nodes. We have left as a future work the protection against more complex attacks where the message content (both to request retransmission and express the gossip utility feedback) can be altered. In these cases, the cryptography primitives are essential, but a proper optimization of their use only when needed is demanding and is planned to be studied by using game theory, as well.

B. PAPER ORGANIZATION

In the latter part of this article, we present in Sec. II details of the gossiping method and the related security vulnerabilities of the available schemes. Next, Sec. III describes in detail our approach, while Sec. IV shows the evaluation of characteristics for the proposed strategy. The paper is concluded in Sec. V with the lessons learned and future research plans.

II. GOSSIPING IN POST-DISASTER PERIODS: PROBLEM STATEMENT AND RELATED WORKS

As presented in Fig. 3, disaster-resilient routing is typically implemented using proactive or reactive approaches [8], [9]. Under proactive schemes, the network is configured in advance (i.e., before the failure occurrence) with alternate transmission means (e.g., by backup paths or redundant network topology). Damages caused by disasters are then less likely to compromise the connectivity of network nodes and the QoS perceived by the users. The reactive methods are executed after failures and involve either a dynamic setup of backup paths [9] or redundancy at the message level [21] (with message retransmissions triggered after detecting losses).

In general, among schemes based on redundancy at the message level, we can distinguish either event notification or application-level schemes. Standards for event notification schemes, such as Object Management Group Data Distribution Service for Real-Time Systems (OMG DDS), Advanced

Message Queuing Protocol (AMQP), or Extensible Messaging and Presence Protocol (XMPP) [22], suggest using TCP to achieve communications resilience. However, TCP exhibits scalability issues when utilized for reliable multicasting and provides only link-by-link guarantees [22], [23], which are not suitable to assure the end-to-end guarantees.

Several application-level approaches exist within the literature and support the end-to-end guarantees, among which retransmission schemes, such as Lateral Error Recovery (LER) [24], Cooperative Error Recovery (CER) [25], or *gossiping* [26], are preferable due to their high delivery success rate. LER and CER are quite complex to implement as they require global knowledge on the participants in a group and the topology established among them. In particular, CER needs the computation of the minimum-loss correlation groups, while LER involves the segmentation of nodes in distinct planes. Gossiping is simpler and more flexible, as such knowledge is not required, making it more suitable for large-scale infrastructures. Gossiping is a distributed retransmission scheme that can achieve a consistent view of received messages among a group of processes within a given probability. It is thus an example of temporary redundancy at the message level, as opposed, e.g., to spatial redundancy where multiple disjoint paths are exploited. It can offer a high success rate of message delivery with a closed loop of control provided by temporary redundancy by implying the following procedures:

- (a) running the loss detection (which can be done based on the elapsed timeouts on the reception of certain messages or by checking the missing message IDs);
- (b) the resulting recovery using retransmissions.

The gossiping scheme can operate in two modes: pull and push. Its *pull mode* is based on the periodically commenced gossip rounds, where each node randomly selects peers towards which it sends a summary of the latest received message IDs so as the destination nodes can detect losses and request retransmissions. In the *push mode*, the nodes send the last received messages so that destinations can recover lost messages without having to request for any retransmission.

The strength of gossiping is that it exhibits a very high delivery ratio, even in the case of severe loss patterns and node/link unavailabilities. However, this is achieved at the price of the increased network traffic concerning the augmented number of exchanged messages (by implementing the selective flooding). This issue can be solved by combining gossiping with the forward error correction (FEC) code (as in [27]) to reduce the traffic load without compromising the achievable delivery success rate.

The comparison of the service-based publish/subscribe solution using TCP and the one equipped with gossiping presented in [28] proves higher scalability of latency for gossiping than for TCP. With TCP, the latency increases linearly with the number of destinations, the applied loss rate, and the source-destination distance. At the same time, gossiping can only achieve a moderate performance degradation and exhibits a logarithmic trend, instead of a linear one [28].

Since gossiping implements a closed control loop, the respective feedback on the utility of the applied redundancy received by the end nodes may be further used to better tune the redundancy to improve the overall approach quality, as in [29]. In approaches based on spatial redundancy, e.g., exchanging additional data to recover lost packets along a single path or multiple disjoint paths [30], [31], the success rate depends on the optimal setup of the method (including the loss ratio to be tolerated). However, in such cases, the configuration needs to be optimally determined from the beginning, as nodes do not receive any feedback during the operation to adjust it further autonomously. Unfortunately, for disaster events, such a configuration cannot be appropriately determined beforehand as a disaster size is typically not known in advance. It, in turn, leads to either over-estimations (i.e., higher redundancy than needed, implying a waste of networking resources) or under-estimations (when the applied redundancy is not sufficient to cover the experienced losses, and some non-delivered messages are irremediably lost).

As discussed above, resilience schemes based on gossiping are thus well-suited in scenarios of massive failures including natural disasters and disruptions, technology-related massive failures, and malicious activities, where the expected extent of losses is hardly known in advance. In particular, a closed control loop and the possibility of adjusting the applied redundancy degree allow gossiping to achieve a high degree of resilience. Therefore, the publish notification can reach all the interested subscribers with a high probability [27].

However, a conventional gossiping scheme exhibits vulnerabilities which can be exploited in further attacks or to compromise the applications built on top of gossiping. Overcoming this issue is thus of utmost importance, especially if gossiping is applied in critical infrastructures, for instance in blockchain platforms (such as Hyperledger [32]), power grids [33] to achieve disaster-resilient event notification [34].

It is worth noting that the presence of malicious nodes in emergency networks in a post-disaster period is not a rare event [35], [36] and counteracting them is particularly demanding. Among the broad spectrum of possible issues, Denial of Service (DoS) attacks are considered extremely dangerous, and their prevention is of pivotal importance. First of all, unintentional congestion is typically caused by a considerable increase of the user traffic in the first phase after a disaster occurrence (due to activities of humans contacting their dear ones, assessing the situation, or seeking an escape route). Therefore, in a communication network already affected by a disaster, it is simple to intentionally complement the legitimate traffic with fictitious one (to cause network unavailability, degraded QoS or resource starvation).

When networks are severely affected by a disaster, ad-hoc networking may become the only viable communication means [37]. In such a context, a DoS attack has a significant impact on draining device batteries and causing disruptions [38]. Its low complexity, combined with the high effect determines the high risk and importance for prevention.

Previous works have investigated the use of cryptographic primitives, majority consensus, or routing monitoring techniques to protect gossiping from various attacks including DoS, eavesdropping, Sybil or Byzantine node attacks [38]–[43]. DoS attack [38] was found to be simple to be performed for gossiping and very disruptive by compromising the network connectivity and reducing node availability.

Typical approaches to protect gossiping from DoS attacks use cryptographic primitives (*i.e.*, digital signatures) to avoid forgery or data injection [39], [40]. However, as shown in [44] they are affected by space decoupling violations (when the data source cannot be anonymous to trace back its public key or digital certificate), scalability issues for identity claims or certificates, key sharing, and performance worsening, which makes them not fit well in a post-disaster scenario.

Specifically, post-disaster communications in publish/subscribe services is data-centric and anonymous in the sense that subscribers are not aware of the identity of the publishers responsible for being the sources of specific information. It is in contrast to the main cryptographic primitives as the sender identity must be known to obtain its public key for decryption or signature verification. The number of publishers can be massive, as within the affected area all the on-field first responders, the citizens looking for an escape path, and other involved actors are publishers and subscribers of disaster-related data. Therefore, a subscriber may have to store a large number of certificates or to manage a large number of keys in the Public Key Infrastructure (PKI) [45], causing excessive use of memory and time. Last, the majority consensus and routing monitoring are challenging in disaster-related scenarios, as they demand a global knowledge on the network topology, which is not feasible to achieve as mentioned above. It makes safeguarding the gossip-based communications remain an open issue.

A common approach to address the message loss in communication networks is to resend messages that were not received by a destination node. In a retransmission-based scheme, the first phase is to detect the occurrence of message loss, while the second phase is to trigger the retransmissions with a proper command [21]. The destination node can recognize the message loss if the identification of two consecutive messages is not progressive in terms of their numbering. For instance, the destination node identifies the loss of message 2 after receiving two consecutive messages with their IDs equal to 1 and 3, respectively. If a loss is detected, the destination node can send a proper command (such as a negative ACK, containing the sequence numbers of the missing message) to trigger the retransmission.

Alternatively, the destination node has to acknowledge the reception of all the messages. If an ACK is not received before the expiration of a given timeout, the retransmission is commenced. It is the Automatic ReQuest (ARQ) scheme [21] adopted by TCP to achieve reliable communications [46]. However, a disadvantage of this solution is centralization of storing and performing the retransmissions of messages at the data source. As a result, a reliable multicast or

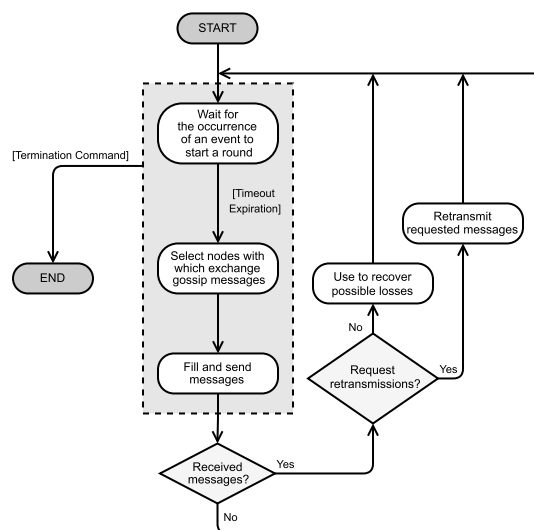


FIGURE 4. Sequence diagram of a generic gossip algorithm.

publish/subscribe service do not scale well with the increase of the number of interacting nodes [22]. The publisher has to keep storing the messages unless all destination nodes have acknowledged their reception, which can quickly saturate its message queue. Therefore, more distributed solutions have been proposed [47], [48], among which gossiping [26] is the most widely used.

Gossiping is a peer-to-peer approach, where all the nodes run the same algorithm, and involves multiple steps as shown in Fig. 4. The starting point of the algorithm is the activation of a node in the overall infrastructure. The algorithm is next repeatedly executed until the node is not terminated. If a termination command is triggered, the eventually on-going task (*i.e.*, the current one in the dashed rectangle in Fig. 4) is interrupted. However, there are many different variants of gossiping implementing each of these steps in a specific manner. The first operation relies on determining when a gossip round should be started (*i.e.*, when the nodes should interact with each other to detect losses and trigger retransmissions). The round can be started, *e.g.*, as soon as a new message is successfully received by one node. However, there is also a possibility to commence the round only when a proper timeout expires. It is worth noting that the initiation of the round is typically characteristic to a given gossiping variant.

The second operation consists in selecting the nodes with which special messages are exchanged. They are often chosen randomly over the list of available nodes provided by a sampling service [49], [50]. Recently, specific heuristics for the node selection have been proposed to optimize given non-functional properties of the algorithm, such as the performance or the success rate [44]. This operation determines what kind of information such messages convey: the last received message or a list of received identifiers. In the latter case, such a list can be used by the destination node to detect possible missing messages and to ask for retransmission. Alternatively, the list can be used to identify a loss occurred at the sender side to retransmit the missing messages.

A push gossip typically includes sending the last received message as soon as it is successfully received. Such a pushing operation is typically done by the destination node when all the packets underlying a notification are received. In severely damaged networks, transit nodes may also be involved in this procedure. On the contrary, a pull gossip waits a given timeout before sending the list containing the identifiers of the latest received messages. The study in [27] proved that both variants allow achieving high resilience degree with a random node selection. However, the push one is characterized by lower recovery time, while the pull one exhibits a lower overhead. Both approaches are very powerful in the case of multiple (and, in particular, correlated) failures, as the recovery can occur from multiple nodes within the network, with a high probability of circumventing damaged links/nodes. Recovery is thus determined on the fly, without the pre-configured strategies and routes. In this sense, gossiping is a multi-path routing from multiple senders, with message sending occurring at different moments.

Group management is not part of gossiping but depends on applications at the higher abstraction level or the overlay at the lower level. Publish/subscribe services exhibit an intrinsic group definition based on subscription mechanisms. Applications with similar event interests (matching active subscriptions and/or advertisements) are assumed to be in the same group so that nodes are classified based on interest, topic or content-based filtering. Overlays may group nodes according to their proximity or subnet mechanisms.

A. PROBLEM STATEMENT

Studies available in the recent literature (e.g., [51]) show that gossiping, although effective in guaranteeing message delivery, might encounter security-related issues or be exploited to perpetrate the following attacks:

- Hub attack – when some nodes fool the others to become hubs for the overlay established among the gossiping nodes so that when they suddenly leave, the overlay is badly fragmented without any hope for recovery,
- Eavesdropping – denoting stealing the data conveyed by the gossiped messages by honest-but-curious gossipers, or an adversary controlling a part of the network,
- Denial-of-Service (DoS) attack – to influence the nodes to send gossip messages towards some nodes to overload them and affect their performance and availability,
- Sybil attack – with an adversary forging legitimate node identities to inject fabricated data, such as fake retransmission requests, or replaying a valid request,
- Byzantine or Selfish Behaviour – with byzantine nodes acting far from the protocol and not actively participating in supporting the other ones coping with losses.

Due to the crucial role of gossiping in providing resilience in emergency networks [34] and utilized as a foundation block in key ICT platforms such as the blockchain [32] or vehicular networks [52], it is demanding to design its secure implementation by protecting the overall algorithm

from these attacks and avoiding possible misuses. Moreover, emergency networks are characterized by a large scale concerning the number of connected nodes and the amount of generated data. Gossiping can cope with multiple failures within the network and to achieve high data delivery guarantees, as proved in [26], [27]. However, the gossip protocol needs to be enhanced with proper means to counteract the attacks mentioned above and the peculiar requirements of emergency networks. To be efficient, the secure and resilient communication protocol in post-disaster periods also needs to scale well with the increase of the size of the network in particular in the context of:

- (a) low latency (as messages may be time-constrained, and, therefore, they should be received fast),
- (b) low energy consumption profile (as devices may have limited battery and cannot be recharged),
- (c) high delivery success rate, despite the possible message losses caused by network failures,

which is certainly hard to assure for non-anonymous communications with a centralized authority/identity management.

B. ANALYSIS OF STATE-OF-THE-ART LITERATURE

Within the context of emergency networks, DoS attacks are considered extremely dangerous, since they are easy to implement and capable of severe damages to the network by making certain key nodes unavailable. Therefore, protection against them is of pivotal importance [36]. In this section, we provide an overview of the main solutions for DoS protection for gossiping. Also, we describe those means with the primary objective to protect against the other attacks, but which can be useful for DoS attacks as well.

Another notable example is provided in [51], where Sybil attacks are avoided by the Oracle Certificate Authority (CA) service assigning identities to gossiping participants so that malicious adversaries cannot forge identities. Moreover, paper [51] suggests to limit the number of accepted requests per round, to bound the damages applied by DoS attacks. However, it comes at the expense of higher performance and energy consumption and lower scalability guarantees. First, it is mandatory to know the sender identity to verify the signatures attached to the exchanged messages. One way is to contact a trusted service to retrieve the sender certificate from which the public key is obtained to verify signatures. Concerning event notification largely used in emergency networks, spatial decoupling is necessary to achieve flexibility and efficiency, but signature verification demands to know the sender identity by violating anonymity. Moreover, signing and verifying messages has a computational and performance cost, which may be an obstacle to resource-constrained nodes. These existing solutions for DoS protection leveraging on cryptographic primitives [39], [40], [51], [53] are inefficient due to:

- anonymity violations (without knowing the sender public key verified using its identity by PKI, the receiver cannot decrypt the exchanged messages),

- scalability issues to manage the identity claims and certificates (receivers have to store the cryptographic data for all senders, progressively losing memory space),
- performance and energy costs (encryption and decryption introduce communication delay, and imply considerable energy and computing resource consumption).

A different approach is proposed in [54], where nodes consider the received gossiped messages to be valid if many replicas of such a message are delivered from many peers over a given threshold, equal to $t + 1$, where t is half of the overall number of nodes in the system. Despite protecting against false requests and gossip messages used to perform DoS attacks, such a solution increases the algorithm convergence time to reach a consensus (*i.e.*, when all the interested nodes receive a given message). Such a more substantial convergence time causes an increase in communication latency, and the higher number of messages needed for the protocol convergence implies higher energy consumption. Therefore, in the current literature several similar solutions, such as [41], [42] were proposed to reduce the convergence time by lowering the number of nodes needed for obtaining the majority, or by using the contextual information to identify the genuine gossip messages (such as characterizing neighbouring nodes' identities or location information). For this purpose, disjoint paths are used for fast convergence at the expense of larger messages (*e.g.*, any time gossip messages are exchanged, they carry all the respective contextual details).

A different remedy for DoS attacks is presented in [38] by sending the pull gossip messages (*i.e.*, when digests are transmitted) to publicly known ports while retransmitting to randomly selected ports, and randomly picking messages to exchange from digests to prevent from being overwhelmed by bogus messages. Besides DoS protection, there is always the other side of the coin: such a solution increases the convergence time and lowers the gossip reliability degree.

The avoidance of sybil attacks without using cryptographic primitives has been extensively investigated in the current literature. A possible approach is to leverage on trust management, as proposed in [43] to detect malicious nodes and exclude them from gossiping to prevent from DoS attacks. However, such an approach requires a proper certificate or identity management and is characterized by the consequent efficiency, scalability and anonymity limitations.

Routing monitoring is crucial against Hub attacks as it enables excluding the nodes with a high number of outgoing connections. Although it is particularly simple in theory, performing node monitoring by identifying the outgoing connections is far from being easy to implement in real settings. One way is proposed in [49] and involves monitoring the peer sampling, which is a service to provide nodes with the ability to know the identifiers of the other nodes in the network and to measure the structural prestige of nodes in a network. Such a solution is not able to exclude nodes launching DoS attacks, which do not typically have a structural prestige but are composed of a group of nodes that won't stand out. A more distributed approach is proposed in [55] where each

node holds multiple neighbour lists to identify malicious nodes and obtain a secure peer sampling service even able to cope with DoS attacks. In [56], game theory is used to design the "incentive and punishment" scheme to avoid selfish/byzantine behaviour of free-riders. This solution cannot deal with DoS attacks as the adversary nodes do not execute such a scheme and can still overload the other nodes with requests.

The summary of the pros and cons of the existing solutions is presented in Tab. 1. Some of them, despite being not designed especially for DoS attacks, can be adapted/extended to address them. Existing approaches to counteract DoS attacks are characterized by limited scalability. For most of them, the convergence time of the gossiping is increased, implying higher latency and energy consumption, and a lower delivery rate. As discussed in this section, post-disaster emergency networks may be compromised by multiple failures, which limits the routing monitoring and secure peer sampling.

To overcome these problems, we propose here a non-cryptographic approach to overcome DoS attacks by exploiting the game theory to achieve a smarter node selection. The adoption of game theory does not present the drawbacks mentioned above and is used for protection against selfish attacks of free-riders.

It is worth noting that although DoS attacks are typically associated with the intentional congestion of the network, a similar effect can occur unintentionally. Indeed, congestion can also happen due to user traffic spikes combined with the overhead imposed by gossiping (in terms of the additional messages needed to cope with losses). There is a body of literature to make gossiping efficient by reducing the number of required messages to recover from network misbehaviour. In [57], message suppression is presented, where authors statistically modelled the probability of firing messages or silencing them to keep the new information propagation with a lower overhead. In [58], gossiping to members farther away is made less frequently than to nearby nodes, so that the network overhead is kept reduced. These solutions cannot be applied to the intentional starvation of resources or congestion since they are only meant to make the protocol more effective but not to prevent a node (or a group of nodes) from directing too much traffic towards some of the neighbouring nodes. As shown in the remaining part of this article, in our proposal, we have extended them to reduce the overhead of the protocol by avoiding the unnecessary traffic, but also by not automatically replying to an incoming request, where the sender identity may be forged or spoofed. The reply can be decided based on its utility.

III. THE PROPOSED APPROACH

This section presents our approach to protect the gossip protocol against DoS attacks without the need to use the cryptographic primitives. We leverage on a proper node and request selection mechanism to maximize the protocol utility using a game-theoretic formulation and a reinforcement learning

TABLE 1. A summary overview of the existing literature on secure gossiping.

Ref.	Approach	Attacks	DoS Relevance	Scalability	Post-Disaster Efficiency
[38]	Retransmitting to randomly selected ports	DoS	Yes	Limited convergence time	Limited as latency and energy consumption increase
[39]	Encrypted gossip messages	Eavesdropping	No, as a node can be still overloaded by decrypting received messages	Limited due to key management	Limited as latency and energy consumption increase
[40]	Signed gossip messages	Sybil	Yes, as invalid messages can be easily discarded	Limited due to certificate management	Limited as latency and energy consumption increase
[41]	Replicated message reception	Sybil	Yes, as messages with low reception rate are discarded	Slightly increased convergence time	Limited as latency and energy consumption increase
[42]	Replicated message reception	Sybil	Yes, as messages with low reception rate are discarded	Slightly increased convergence time	Limited as latency and energy consumption increase
[43]	Detect and exclude malicious nodes	DoS, Sybil	Yes	Limited due to identity management	Limited as latency and energy consumption increase
[49]	Measure the structural prestige of nodes	Hub	No, as DoS attacks are not performed by nodes with a high structural prestige	Limited as failures disturb the monitoring	Limited as failures disturb routing monitoring
[51]	Assign identities, limit requests to reply	DoS, Sybil	Yes	Limited for identity management	Limited as latency and energy consumption increase
[53]	Signed and encrypted gossip messages	Eavesdropping, Sybil	Yes, as invalid messages can be easily discarded	Limited due to used cryptographic primitives	Limited as latency and energy consumption increase
[54]	Replicated message reception	Sybil	Yes, as messages with low reception rate are discarded	Increased convergence time	Limited as latency and energy consumption increase
[55]	Use secure peer sampling service	Hub	Yes, as malicious senders are excluded	Limited for keeping multiple neighbour lists	Limited as failures disturb the secure peer sampling
[56]	Game-theoretic incentive/punishment scheme	Byzantine/selfish behaviour	No, as malicious nodes may bypass such a scheme	Yes, only local decisions are used	Yes, as no additional overhead is introduced

scheme. We base our work on our previous study on reducing the gossiping overhead [44], extending it towards protection against DoS attacks by optimizing not only the selection of the destinations of the gossip messages but also the decision to reply to or ignore the incoming requests. It is done by integrating simple trust management from feedback (as in [43] without a central authority) with the game-theoretic approach applied not only to gossipers (as in [56]) but also to the gossip receivers. The solution we propose in this article is suitable for post-disaster periods as it does not depend on any centralized authority needed for identification, trust, and certificate management, but exploits the distributed decision making formalized here as a non-cooperative game.

In particular, in Sec. III-A we explain the optimization model for finding the optimal solution to the considered gossiping problem. However, due to its centralized nature and the related scalability problems, we describe it here as a reference model only, while in Sect. III-B, we introduce the respective non-cooperative game to enable the real-time operation of gossiping in a distributed scenario. Section III-C describes the proposal of learning capability introduced to the game to cope with the uncertainty of the payoff functions.

A. GOSSIPING OPTIMIZATION

Gossiping is a stateless protocol with a message flow depicted in Fig. 5(a). A DoS attack attempt can be conducted by a set of malicious nodes faking gossip/request messages on behalf of a target node to push the contacted nodes to flood the target

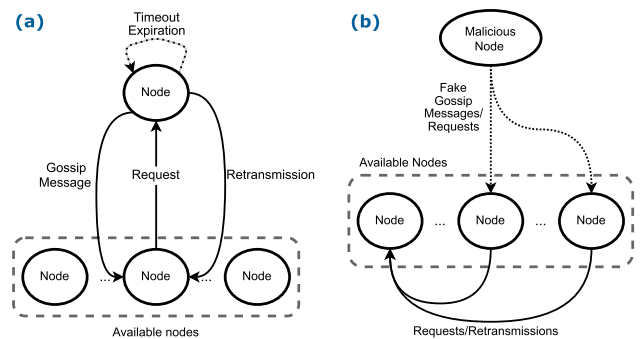


FIGURE 5. The correct message flow (a), and during a DoS attack (b).

with requests or retransmissions, as illustrated in Fig. 5(b). It is possible since the nodes do not remember the sent gossip messages. Introducing a state within the protocol to cope with DoS attacks can cause issues in an unpredictable network environment such as the Internet. On the contrary, by applying the message suppression concept originally introduced for the congestion control [57], we can avoid the necessity to reply to fake messages by measuring a utility function.

A typical solution for the congestion control in gossip protocols consists of defining a utility function related to destination selection, as investigated in our earlier work [44]. Such a function allows a node to decide with which other nodes to gossip, to avoid contacting those nodes which are not in need. To protect against DoS attacks, we also define the utility of replying to an incoming request based on such

a utility. Specifically, in both cases, a node has to make the optimal choice by considering if this step will allow it to recover from a loss and avoid exchanging messages which would be finally recognized as waste. Formally, given a set L of all nodes n within the system, we can define a function named $h_{i,j}^{(k)}$ taking the value of 1 if the message sent by the i th node to the j th node in the k th round was useful to recover a lost message (0, otherwise). In [44], function h was only related to the utility of a sent gossip message. In this work, as presented in formula (1), this function consists of two components: $\hat{h}_{i,j}^{(k)}$ coming from our earlier work [44], and $\bar{h}_{i,j}^{(k)}$ introduced here to measure the utility of replying to the incoming messages.

$$h_{i,j}^{(k)} = \hat{h}_{i,j}^{(k)} + \bar{h}_{i,j}^{(k)} \quad (1)$$

The first function $\hat{h}_{i,j}^{(k)}$ is assigned the value of 1 if the gossip message sent by the i th node to the j th node in the k th round has been useful (i.e., the probability that message m is within the set M_n of received messages of node n after sending the gossip message $m_{i,j}$ is equal to 1):

$$\hat{h}_{i,j}^{(k)} = \begin{cases} 1 & \exists m \in M : P(m \in M_n | m_{i,j}) = 1, n \in L \\ 0 & \nexists m \in M : P(m \in M_n | m_{i,j}) = 1, n \in L \end{cases} \quad (2)$$

The right part ($\bar{h}_{i,j}^{(k)}$) of formula (1) is equal to 1 if a message sent by the i th node to the j th node in the k th round after having received a request (namely $r_{i,j}$) has been useful:

$$\bar{h}_{i,j}^{(k)} = \begin{cases} 1 & \exists m \in M : P(m \in M_n | r_{i,j}) = 1, n \in L \\ 0 & \nexists m \in M : P(m \in M_n | r_{i,j}) = 1, n \in L \end{cases} \quad (3)$$

The utility for a given node i at a given k th round, namely $U_i^{(k)}$, can be expressed as the product of function $h_{i,j}^{(k)}$ from formula (1) with $x_{i,j}^{(k)}$, which is equal to 1 if the i th node has sent a message to node j during the k th round, summed up for all the nodes and the performed rounds, as given in formula (4).

$$U_i^{(k)} = \frac{1}{f_{out}^{(i)} + r_{in}^{(i)}} \sum_{j \neq i} x_{i,j}^{(k)} h_{i,j}^{(k)} \quad (4)$$

where:

- $f_{out}^{(i)}$ indicates the fanout of the protocol, i.e., the number of messages to be sent at each round,
- $r_{in}^{(i)}$ indicates the number of received requests.

The division by the sum of values $f_{out}^{(i)}$ and $r_{in}^{(i)}$ in formula (4) is needed to keep the utility function $U_i^{(k)}$ returning a value within [0, 1] range. To distinguish the utility of gossiping and the one of replying, we introduce two sets of decision variables:

- $\hat{x}_{ij}^{(k)}$ equal to 1 if the i th node sends a gossip message to node $j \in L$ during the k th round (0, otherwise),
- $\bar{x}_{ij}^{(k)}$ equal to 1 if the i th node replies to an incoming request from node $j \in L$ during the k th round (0, otherwise).

By splitting function $h_{i,j}^{(k)}$ into two separate parts related to formulas (2) and (3), accordingly, utility function $U_i^{(k)}$ can be then extended as follows:

$$\begin{aligned} U_i^{(k)} &= \frac{1}{f_{out}^{(i)}} \sum_{j \neq i} \hat{x}_{i,j}^{(k)} \hat{h}_{i,j}^{(k)} + \frac{1}{r_{in}^{(i)}} \sum_{j \neq i} \bar{x}_{i,j}^{(k)} \bar{h}_{i,j}^{(k)} \\ &= \hat{U}_i^{(k)} + \bar{U}_i^{(k)} \end{aligned} \quad (5)$$

where $\hat{U}_i^{(k)}$ indicates the utility for sending the gossiping messages to the selected nodes, while $\bar{U}_i^{(k)}$ is the utility of replying to certain incoming requests, both achieved by the i th node during the k th round.

By summing all such functions over all already started rounds k for all nodes in the system, we obtain the overall gossip utility U as given in formula (6).

$$\begin{aligned} U &= \frac{1}{N} \sum_{i=0}^N \frac{1}{R} \sum_{k=0}^R U_i^{(k)} = \frac{1}{N} \sum_{i=0}^N \frac{1}{R} \sum_{k=0}^R (\hat{U}_i^{(k)} + \bar{U}_i^{(k)}) \\ &= \frac{1}{N} \sum_{i=0}^N \underbrace{\frac{1}{R} \sum_{k=0}^R \hat{U}_i^{(k)}}_{\hat{U}_i} + \frac{1}{N} \sum_{i=0}^N \underbrace{\frac{1}{R} \sum_{k=0}^R \bar{U}_i^{(k)}}_{\bar{U}_i} \end{aligned} \quad (6)$$

where N is the total number of nodes within a large-scale system, while R is the total number of the already started rounds.

Such a utility function has two interpretations. Firstly, \hat{U}_i represents the number of gossip messages that have triggered retransmissions over the total number of messages sent by the node i . Secondly, when a node has to reply to a request, \bar{U}_i can denote the number of retransmitted messages useful to recover a loss over the total number of messages received by the node i . \hat{U} and \bar{U} are respectively the mean of \hat{U}_i and \bar{U}_i over the total number of nodes.

Our previous work [44] aims at only optimizing \hat{U} (not the overall utility function U , as we focus on in this article). Moreover, while the optimization of the \hat{U} part is mainly tailored to the congestion control and the overhead minimization related to the gossip protocol, the maximization of \bar{U} in this article allows us to prevent from DoS attack attempts.

When optimizing the gossip protocol, it is essential to maximize such a utility function given in formula (6) by making the right decisions to achieve a high resilience degree so that the gossip messages are sent towards those nodes which need them. However, this is not the case for a conventional gossiping scheme. In particular, we proved in [59] that the value of U is low in the case of the optimal selection of the nodes (i.e., when a node sends a gossip message only towards the nodes which need it), and further decreases when $f_{out}^{(i)}$ is increased. In addition to strengthening resilience, optimizing the utility U is also necessary for security, and to protect against DoS attacks aimed at sending the retransmission

requests to multiple nodes on behalf of the targeted node. By optimizing the utility U , those contacted nodes can avoid replying to such fake requests (and thus not overload the target nodes, respectively).

The problem of finding the node with which a given node can gossip has been formulated in [44] by the optimization model to maximize the utility function \hat{U} given by formula (6). In particular, to address protection against DoS attacks, the overall problem of determining whether to send gossip messages toward given nodes or to reply to specific incoming requests can be formulated as done in [44], extended with the utility part \bar{U} for the DoS attack protection.

Let us consider a set of variables $\rho_i^{(k)}$ indicating the obtained reliability of the message delivery for the i th node, computed at the k th round. The objective is to find during a given k th round the proper values for the decision variables in $\hat{x}_{i,j}^{(k)}$ and $\bar{x}_{i,j}^{(k)}$ to maximize the utility and the achievable reliability at the overall system level (as both are equally important, we sum them up):

$$\begin{aligned} \max & \frac{1}{N} \sum_{i=0}^N (\hat{U}_i + \bar{U}_i) + \frac{1}{NR} \sum_{i=0}^N \sum_{k=0}^R \overbrace{\rho_i^{(k)}}^{\rho_i} \\ & = \frac{\sum_{i=0}^N \sum_{k=0}^R \sum_{j \neq i} (\hat{x}_{i,j}^{(k)} \hat{h}_{i,j}^{(k)} + \bar{x}_{i,j}^{(k)} \bar{h}_{i,j}^{(k)})}{NR (f_{out}^{(i)} + r_{in}^{(i)})} + \frac{1}{N} \sum_{i=0}^N \rho_i, \end{aligned} \quad (7)$$

subject to:

$$\sum_{j \in L} \hat{x}_{i,j}^{(k)} \leq f_{out}^{(i)} \quad (8)$$

$$\sum_{j \in L} \bar{x}_{i,j}^{(k)} \leq r_{in}^{(i)} \quad (9)$$

$$\hat{x}_{i,j}^{(k)}, \bar{x}_{i,j}^{(k)} \in [0, 1] \quad (10)$$

$$\hat{h}_{i,j}^{(k)}, \bar{h}_{i,j}^{(k)} \in [0, 1] \quad (11)$$

$$0 \leq \rho_i^{(k)} \leq 1, \quad \forall i \in L \quad (12)$$

Constraints (8), (9) indicate that a node cannot send more than $f_{out}^{(i)}$ gossip messages and replies to more than $r_{in}^{(i)}$ received requests, accordingly. Constraints (10), (11), (12) express the admissible values for the decision variables and the objective functions.

Here, maximizing the utility function makes evident the added-value of our solution by using a strategy of gossip message suppression not only to optimize the overhead of the protocol (obtained by maximizing function \hat{U}_i) but also to deal with DoS attacks (obtained by maximizing function \bar{U}_i). It allows us to protect the scheme from these attacks without the need for using cryptography while fulfilling at the same time the requirements on fast delivery of messages in post-disaster scenarios.

The overhead of the ‘‘plain’’ gossip protocol and the proposed DoS-resistant one can be kept as minimal as possible, in particular, if the problem is approached in a centralized

manner by a node with a global knowledge on the overall system (by collecting statistics by the respective distributed snapshot algorithm [60]). In this case, the resolution can be made through a Mixed-Integer Program (MIP) formulated above, or by relaxing the integral constraints and solving the corresponding pure Linear Program (LP). Such a solution is simple to implement and allows us to obtain the optimal decisions. However, since this approach requires the global knowledge of the lost messages for each node to be able to make the best choice, it is unfeasible in a large-scale scenario.

Moreover, even if assuming the possibility to acquire global knowledge, the model still exhibits severe scalability limitations prejudicing its usage in systems consisting of a large number of nodes. The time necessary for collecting the loss statistics by a given central decision node increases exponentially with the number of nodes within the multicast tree, as the system is asynchronous [61]. Also, the memory required to store all the received data and the load resulting from resolving the optimization problem may overwhelm the resources of the central decision node. As a result, it may become infeasible to make any decision at all.

A better approach is to distribute the resolution duties among all the nodes, as considered in the following part of this section, by relying on the local decisions based on the local knowledge acquired by the neighbouring nodes. Such an approach has the strength of overcoming the scalability issues for large systems. However, as discussed later in this article, this advantage comes at the price of obtaining a solution that is relatively far from the optimal one.

B. FORMULATION OF A NON-COOPERATIVE GAME

Among the possible means for a distributed optimization, we have adopted a resolution method based on game theory [62]. The reason for this choice is that game theory inverts the typical resolution approach into a distributed optimization [63]. It does not require to divide the overall optimization task into independent sub-problems, one per each node, and to define a proper supervision mechanism to make all the sub-problems converging towards the global optimum. On the contrary, the game theory allows reaching the optimal solution, even without imposing any complex coordination and supervision protocols among the nodes. It makes the proposed approach suitable for post-disaster scenarios as the latency and overhead of the gossiping scheme is kept minimal, and massive failures in the network can only slightly influence the approach. As the optimization is conducted using local decision and local knowledge on the overall problem without any supervision and coordination, the method is meant to scale well even for a high number of nodes. In the rest of this subsection, the overall cooperative game will be introduced and analyzed.

In our approach, we have a set $P = c_1, c_2, \dots, c_N$ of players of a finite size $N \geq 2$, each player associated with each node of the communication system. A node cannot host more than one player. The players simultaneously decide whether to reply or not to requests received from the other players,

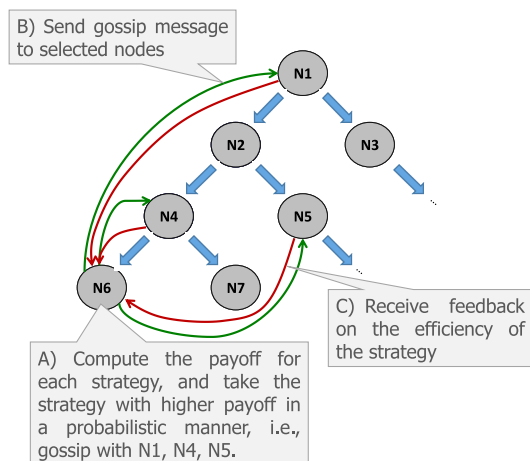


FIGURE 6. A general idea of the proposed approach.

to recover the maximum number of messages dropped at the requesting nodes and to minimize the relative overhead. Formally, the set of strategies for each player is defined as S . A strategy of any player is the selection of a set of the other players (which can even be null) with whom it is allowed to communicate by sending a gossip message and/or replying to their requests. Based on the adopted strategy, indicated as s_j , the proper action is conducted by the j th player, namely o_j among the set of allowed actions $O_j(s_j)$, i.e., sending a certain message according to the gossip algorithm. All the players have the same set of strategies, and there are no differences in their strategy sets.

In our game, we consider the nodes being structured in a tree, as many content-based publish/subscribe services establish a tree-based overlay [64], where each node is placed in a given layer, as shown in Fig. 6. For example, the root node is in layer 0, the children of the root node are in layer 1, and so on. The strategy for a player is given by an integer number from $[1, \lambda]$ interval, where λ is the highest layer within the tree (i.e., with the lowest number). The output of the player after the k th round is a random selection of $f_{out}^{(k)}$ reachable nodes belonging to the layer indicated by the strategy to which this player sends gossip messages or replies to the received requests. Combining the strategy sets of all the players, namely $S = S^{c_1} \times S^{c_2} \times \dots \times S^{c_p}$, a strategy profile $s \in S$ implies certain payoffs to each player c , namely $\Phi^c(s)$, which are aggregated in the so-called profiles of payoffs denoted as π . The payoff is the gain for the player to gossip with a set of nodes. Specifically, in our game, player c receives gain for having one of its messages being useful for a node during the previous k th round, i.e., being able to recover a lost message over the total number of sent messages:

$$\Phi^c(s) = \alpha \left(\frac{\sum_{i \in L} h_{i,j}^{(k)}}{\sum_{i \in L} x_{i,j}^{(k)}} + \rho_i^{(k)} \right) \quad (13)$$

where α is a weight within the $(0, 1)$ interval, s contains the vector of decision variables $x_{i,j}$, being the sum of $\hat{x}_{i,j}^{(k)}$ and $\bar{x}_{i,j}^{(k)}$, while $h_{i,j}^{(k)}$ is the sum of $\hat{h}_{i,j}^{(k)}$ and $\bar{h}_{i,j}^{(k)}$.

The scope of the game is to determine the best strategy profile s^* that implies the maximum payoff for all the players, expressed by formula (7) distributed among the players.

$$s^* \in \arg \max_{s \in S} \frac{\sum_{i \in L} \sum_{j \neq i} (\hat{x}_{i,j}^{(k)} \hat{h}_{i,j}^{(k)} + \bar{x}_{i,j}^{(k)} \bar{h}_{i,j}^{(k)})}{N (f_{out}^{(i)} + r_{in}^{(i)})} + \frac{1}{N} \sum_{i \in L} \rho_i^{(k)} \quad (14)$$

By having each player to maximize their payoff function in formula (13), we have a distributed approach for the optimization of formula (7), where the players are myopic, in the sense that they update their gossiping strategy based on instantaneous parameters and feedback, while ignoring the future implications of their actions. It is suitable in a post-disaster scenario as no centralized resolution is used. Compared to the strategy from [44], here we do not only consider the utility when sending the gossip messages but also when replying to the incoming requests.

In our game, players are selfish, i.e., there is no direct communications between the players, and each one only cares to maximize its own profit or to minimize its own costs without considering the state of the other players (with the eventuality of unintentionally damaging them). Then, the game is defined non-cooperative, and its standard form is given by $\Gamma = (P, S, \pi)$ to maximize the payoff for all the players.

C. STRATEGIC LEARNING FOR PAYOFF CHARACTERIZATION

In a typical application of game theory, the payoff functions of the players are assumed to be well known and externally given, as a part of the game formalization. In our case, however, the payoff functions for each player are not given in advance, as it is not possible to predict the impact and the effect of disasters to the network and its constituents. Also, disaster-related failures may activate at different time instants. Therefore, it is crucial to have a learning scheme for post-disaster resilience by having the payoff functions to be dynamically computed based on the feedback provided by the contacted nodes. To address this challenge, we propose to utilize the concept of a distributed strategic learning [66], where each player can learn from the received feedback to create a payoff value and to determine if a given strategy is the best response, as shown in Fig. 6. Specifically, following our earlier work [44], here we also apply the COmbined fully DIstributed PAYoff and Strategy-RL (CODIPAS-RL) [66], which is a learning scheme derived from the strategy and payoff (Q-learning) Reinforcement Learning.

Let us denote by $p_{j,t}(s_j)$ the probabilities of the j th player to choose strategy s_j at time t , and $p_{j,t} = [p_{j,t}(s_j)]_{s_j \in A_j} \in P_j$ be a mixed strategy of the j th player. Moreover, we indicate by $\hat{g}_{j,t}$ the perceived payoff at time t and $g_{j,t+1}$ is the collected feedback at time $t + 1$ for the selected strategy s_j . CODIPAS-RL works as follows. At time-slot $t = 0$, each

player chooses strategy s and derives from it action o to be performed. Then, it receives feedback for its action and builds a numerical value of its payoff. The payoff is properly initialized to $g_{j,0}$. In time-slot $t > 0$, each player having an estimation of its payoffs, namely $g_{j,t}$, chooses strategy $p_{j,t+1}$ for the next time-slot, which is a function of only the previous strategy $p_{j,t}$, the estimated payoff $g_{j,t}$ and the target value for the payoff function. The game moves to time-slot $t+1$. Such a scheme is combined with a proper payoff and learning strategy leading to CODIPAS-RL with the Boltzmann-Gibbs distribution as a strategy mapping, formulated as follows:

$$\begin{cases} p_{j,t+1}(s_j) = (1 - \lambda_{j,t})p_{j,t} + \lambda_{j,t}\tilde{\beta}_{j,\epsilon_j}(\hat{g}_{j,t}, s_j) \\ \hat{g}_{j,t+1}(s_j) = \hat{g}_{j,t}(s_j) + v_{j,t}\alpha_{\{o_{j,t+1} \in O_j(s_j)\}}\delta_g \end{cases}, \quad (15)$$

$$j \in [1, N], t \geq 0, o_{j,t} \in O_j(s_j),$$

$$\delta_g = g_{j,t+1} - \hat{g}_{j,t}(s_j);$$

where $\tilde{\beta}_{j,\epsilon_j}$ is the strategy mapping expressed as a softmax function (*i.e.*, a function turning a vector of a given number K of real values into a vector of the same number K of real values that sum up to 1):

$$\tilde{\beta}_{j,\epsilon_j}(\hat{g}_{j,t}, s_j) = \frac{e^{\frac{1}{\epsilon_j}\hat{g}_{j,t}(s_j)}}}{\sum_{s'_j \in S^j} e^{\frac{1}{\epsilon_j}\hat{g}_{j,t}(s'_j)}}, \quad (16)$$

$$s_j \in S^j, j \in [1, N]$$

where ϵ_j is a parameter assigned to the j th node, which is either identical or different for all the players, and represents the player rationality. When ϵ is closer to zero, the mapping returns the strategy offering the highest payoff. For ϵ closer to 1, the mapping makes a fully random strategy selection. As we aim at having mixed strategies, we assume all the players have the same rationality level: $\forall i \neq j : \epsilon_i = \epsilon_j = 0.1$.

The upper part in formula (15) referring to $p_{j,t+1}$ defines the selection of the strategy to be used, while the lower part (*i.e.*, $\hat{g}_{j,t+1}(s_j)$) updates the payoff based on the received feedback. $\lambda_{j,t}$ is the strategy learning rate that may vary from player to player and/or during the learning process. The active strategy of the player is indicated by $\alpha_{\{o_{j,t+1} \in O_j(s_j)\}}$ taking the value of 1, if action $o_{j,t}$ has been played by the j th player at time t ; 0 otherwise. This implies that only the component corresponding to the action that has been played is updated. The stable solution for such a formula can be assumed as the equilibrium for a modified game, where the game payoff is perturbed with an extra entropy part (to indicate its dependence on the loss pattern applied by the network dynamics and the actions of the other players).

D. GAME ANALYSIS

The scope of this subsection is to provide a formal analysis of the proposed approach to find a solution to the optimization problem expressed in Eq. 7. In particular, we first start by studying the existence of Nash Equilibria representing the widely known solution concept for non-cooperative games

as the one described in Subsection III-B. Then, we motivate the need for the application of the reinforcement learning approach, as a case of a mixed non-cooperative game repeated over the time, as mentioned in Subsection III-C, and verify if such an approach is able to provide a solution for Eq. 7. The applied learning scheme can result in Price of Anarchy (PoA) (*i.e.*, the distance, in terms of the objective functions to be optimised, of the Nash Equilibria with the Pareto front) being lower than the one achieved with the naive non-cooperative game formulation [81].

Given a particular strategy $s \in S$, it is not profitable for a player to select a different node than the one in the current strategy profile since moving to a neighbour node will not change (or even reduce) the achievable payoff. In other words, such a player thus has no incentive to change its strategy. The demonstration of the existence of Nash Equilibria is typically resolved using theorems by making proper assumptions of the characteristics of certain elements of the game, *e.g.*, as presented in [65]. A game can have a deterministic or probabilistic nature: it is possible to take pure strategies, *i.e.*, to pick a single action and play it, or mixed ones, *i.e.*, to have strategies with a random selection over the set of available actions according to some probability distribution. Specifically, the existence of at least one pure Nash Equilibrium is proved if function (13) is continuous and non-positive in the strategy set. This can be proved if the second derivative is non-positive. It is not too complex to compute that

$$\frac{\partial^2 \Phi^c(s)}{\partial s^2} = \frac{2\alpha H}{s^3} \quad (17)$$

as α and $H = \sum_{i \in L} h_{i,j}^{(k)}$ are both non-negative, the second derivative is always positive for s being positive, which is true by definition. Therefore, the Nash Equilibrium is not guaranteed to exist. However, mixed-strategy games, such as the one investigated in this article, always have at least one Nash equilibrium [62], so our gossiping approach considers a mixed-strategy formulation, which is a set of probability distributions over the actions linked to the strategy profile s expressed by $p_{j,t+1}(s_j)$ in Eq. 15. Randomization is mainly needed in our approach as players are uncertain about the other players' action and the predictability of the loss rate applied by the network.

In the proposed work, the mixed-strategy game is repeated over the time, with reinforcement learning used to update the payoff estimation and to determine the best response. Based on the results of the work in [80], it has been proved that the strategic learning applied in this work converges to the so-called rest point (*i.e.*, a case in which players do not change their applied strategy despite the provided feedback) in a limited number of iterations. The Folk Theorem adapted to the evolutionary games indicates such a rest point as the Nash Equilibrium of the expected game [82]. Therefore, the demonstration by using the Banach-Picard fixed point theorem contained in [66] shows that at the convergence of the adopted CODIPAS-RL, the reached rest point is a Wardrop




FIGURE 7. Implementation of the proposed solution.

Equilibrium, *i.e.*, the equilibrium for a modified game with the payoff in Eq. 13 perturbed with an extra entropy term, which helps to overcome the local optimal points and reaching a global one maximising Eq. 7.

The convergence towards the global optimum in a given limited number of iterations depends on the quality of the used approximation function for the payoff and strategy mapping. The used Boltzmann-Gibbs distribution is a smooth best response function and has proved to be an excellent function for this scope [66]. In fact, in the reinforcement learning literature, it is known that the softmax decision process, as Eq. 16, is beneficial [83] in situations where the players have to determine if applying a known but possibly sub-optimal action (*i.e.*, playing the same action of the previous stage assuming the network conditions have not been changed considerably) or to explore a risky but more rewarding one (*i.e.*, playing a different action betting that the network conditions have considerably changed).

IV. ASSESSMENT

This section has the scope of presenting the experimental results showing the achievable quality of our solution in terms of:

- scalability measured as the increase of the end-to-end latency due to the increase of the network and/or disaster scale,
- reliability measured as the success rate, *i.e.*, the percentage of correctly delivered messages over the sent ones,
- security defined as the overall utility of the approach being not compromised by a DoS attack attempt,
- overhead, *i.e.*, the additional traffic generated by the approach in terms of the number of packets per link,
- efficiency measured as the power consumed to run the protection means.

In this section, after highlighting details of the testbed and the experimental setup in Section IV-A, the analysis of characteristics of our gossiping approach is divided into two parts. The first one (Section IV-B) presents the performance of our scheme in the post-disaster period for the three considered scenarios of disasters (*i.e.*, natural disasters, attacks and technology-related disasters) assuming no DoS attacks in that

period. Section IV-C, in turn, is to verify the usability of our scheme assuming that the network already suffering from the disaster is additionally affected in the post-disaster period by a DoS attack.

A. TESTBED AND SETUP DESCRIPTION

We implemented our scheme in OMNET++, which is an event-based simulator for networks and distributed systems [67]. Specifically, we deployed our approach on top of the existing OMNET++ modules implementing Pastry [68] and Scribe [64]. The first one realizes a Distributed Hash Table (DHT) overlay with routing implemented based on a circular hash table's key-space. In contrast, the second one realizes a topic-based publish-subscribe service on top of Pastry. Losses of messages are implemented over Scribe using our gossiping module optimized by the proposed approach to discipline the node selection and request replies. We have emulated a data-centric application for post-disaster recovery, periodically sending data of 23 kB. It resembles the behaviour of rescue teams and/or city administrators regularly broadcasting data to help citizens escape or get information on the current conditions in the affected area. The publication rate has been set to 1 message per second. We are not considering the mobility of the nodes in our simulations because studies, such as the one in [79], proved that mobility and the consequent topological changes are not an obstacle to gossiping and help speeding up the convergence of the message delivery to all interested nodes. Moreover, the infrastructure nodes, such as base stations and routers, are fixed and not mobile by definition. At the same time, the mobility of end users is constrained due to the disaster itself, for instance, by a fire or a volcano eruption. In fact, users in such cases tend to move along the escape and rescue routes to the assembly points.

The links among the routers in the emulated network were assumed to be characterized by a mean delay value of 50ms. In such a simulation, the root generates the data to be distributed along the tree. A pull-based gossip has been implemented, where after the expiration of two seconds, a gossip request is sent to those nodes optimally determined by maximizing the utility function from Eq. 6. We have modified the router's implementation provided by the INET framework in OMNET++ by inserting a piece of software emulating the message losses. It can be tuned by using two parameters: the Packet Loss Rate (PLR) being the probability of losing a packet, and the Average Burst Length (ABL) denoting the mean number of consecutive messages lost by the network. The link loss process has been modelled regarding the *Gilbert-Elliott model* [69], depicted in Fig. 8. It consists of the 1-st order Markov chain model having only two states: a "Good" state and a "Bad" one, each characterized by a given probability to lose packets. In our simulations, the probability of losing packets when the model is in the first (*i.e.*, "Good") state is 0, while such a probability is equal to 1 if the model is in the second (*i.e.*, "Bad") state. The model is characterized by four transition probabilities to change a state from the current one:




FIGURE 8. Schematic overview of the Gilbert-Elliott model.

- (i) P is the probability of shifting to the “Bad” state when the current state is “Good”,
- (ii) $1-P$ is the probability of remaining in the “Good” state,
- (iii) Q is the probability of shifting to the “Good” state when the current state is “Bad”; and last,
- (iv) $1-Q$ is the probability of remaining in the “Bad” state.

Given PLR and ABL , it is possible to compute P and Q as follows [70]:

$$P = \frac{PLR \cdot Q}{1 - PLR} \quad Q = ABL^{-1}. \quad (18)$$

The three disaster cases, all having the same values for PLR and ABL , were simulated as follows:

- **Case A** to model the natural disasters causing failures confined to a given area, which can be emulated by enabling the message loss only for those routers within a given cluster. This is performed by selecting a router that fails together with all its neighbouring ones until the total number of faulty routers is reached.
- **Case B** to represent malicious attacks in the communication network modelled by enabling the message loss for the routers within the core of the network topology having a high number of connections, and leaving those at the edge without the message loss enabled.
- **Case C** referring to the technology-related disasters which can be modelled by tuning the message loss within all the routers at random. In such scenarios, a DoS attack attempt is modelled by having a certain percentage of nodes sending requests for random messages to a given node, and the destinations of such fake requests are randomly selected among the available nodes in the tree.

By using such a module, we have simulated the three envisioned disaster scenarios by randomly selecting PLR and ABL higher than 0 for the nodes belonging to specific portions of the network, depending on the post-disaster case to be simulated. Specifically, we have configured for the faulty routers PLR being around 0.2, and ABL consisting of 3 packets, respectively, which are the values assumed for faulty networks also in other similar disaster-related evaluations (see, e.g., [71], [72]). On the contrary, for healthy routers we have left a 0 value for both PLR and ABL .

DoS attacks were modelled by allowing a percentage of application nodes to simulate the fake replies by waiting for a specific time to send a fake request and trigger the unneeded




FIGURE 9. Communication latency for Case A when varying the number of nodes for Success Rate = 1 when using gossiping.

retransmissions. The simulated network was composed of 1024 nodes connected by 100 routers. The topology of the routers and application nodes has been designed by using the ReaSE framework,¹ which offers a graphical user interface for generation of Network Description (NED) files including a realistic topology to be used in OMNET++ simulations. Specifically, among the available Internet topology models, the one selected for our experiments has been the Heuristically Optimal Topology (HOT) model from [73].

B. EVALUATION WITHOUT DoS ATTACKS

In the first set of experiments, we analyzed the scalability of the classic gossiping protocol (i.e., without having our solution applied), by presenting the results related to the end-to-end latency of the used publish/subscribe service without gossip (using only Scribe) and with gossip. We performed our experiments for multiple cases defined by varying the number of nodes and the number of faulty routers, to show the scalability of the communication protocol considering the horizontal scale of the system. In this scenario, all the routers were assumed to obtain random values of PLR and ABL .

Results presented in Fig. 9 for Case A show that both methods encounter an increasing trend when augmenting the number of nodes. However, the overhead added by the gossiping is slightly dependent on the number of nodes (i.e., the additional latency is 30% on average). This slight dependence on the scale of the system is due to the decentralized approach of gossiping to loss recovery (i.e., with the increase of the number of nodes, more retransmission sources are available which are closer to the requesting nodes, which increases the probability to recover a message and its consequent overhead). As our approach adds a limited additional latency, Fig. 9 shows its usability in post-disaster periods.

Next, we evaluated the reliability of our gossiping scheme by focusing on the achievable success rate for the three major disaster cases A-C described in Section IV-A. Fig. 10 shows the achievable success rate (considered here as the measure

¹<https://omnetpp.org/download-items/ReaSE.html>




FIGURE 10. Success rate over the different cases A-C as a function of the number of contacted nodes to gossip with (i.e., fanout).




FIGURE 11. Success rate for Case A for different percentages of faulty routers as a function of the fanout.

of reliability) over three different runs of our simulation with 10% of faulty routers in the three different identified cases of disaster scenarios by increasing the applied fanout (i.e., the number of contacted nodes to gossip with). We can notice that the case when the core elements are attacked, i.e., Case B is the most difficult to address as such failures can lead to the partitioning of the overall network infrastructure. In such cases, the gossip is not able to recover all the lost messages, as a small portion of them remains undelivered even with a high fanout (due to the high connectivity of the faulty routers able to partially and temporarily partition the overall infrastructure). Also, gossiping is made not possible to overcome such fragmentation: if a loss occurs in a given part of the infrastructure, the nodes in the other parts are not always able to support the recovery from such losses. The other two cases (i.e., A and C) are simpler to recover from as the gossip can circumvent the failed routers and resend the lost messages.

To study the impact of a disaster scale on the performance of our gossiping scheme, we focused on Case A to show the success rate values as a function of the percentage of failed routers in the cluster. Results presented in Fig. 11 show the impact of the increasing scale of a disaster on the increase of the latency due to the higher redundancy degree applied to tolerate all the losses imposed by the network.




FIGURE 12. Latency for Case A as a function of the percentage of faulty routers, for success rate = 1.




FIGURE 13. Success rate difference as a function of the fanout over the different cases A-C without and with DoS attack being applied, by having 3 colluded malicious nodes.

Fig. 12 shows the latency achievable for Case A when increasing the number of faulty routers due to a disaster at the success rate equal to 1 obtained with a fanout of 5, 7 and 12 respectively for 5%, 10%, and 20% of faulty nodes. It is evident that the size of a disaster impacts gossiping by increasing the mean latency due to the higher fanout needed for full delivery of the messages.

C. DoS PROTECTION EVALUATION

Experiments described in this subsection were performed assuming the occurrence of a DoS attack in the post-disaster period. The executed gossiping protocol is augmented by the solutions described in the previous Sec. III. As depicted in Fig. 13, the difference in the success rate between the cases with and without a DoS attack is visible. It is since under a DoS attack, messages are used to compromise the availability of the target node rather than being useful in letting nodes recover from their losses. Also, the congestion within the network results in additional messages being lost. The difference between the impacts of the three cases of disaster scenarios A-C on the efficiency of gossiping is evident also in the context of the effects introduced by the DoS attack in the post-disaster period, with Case B being the one with the highest difference.




FIGURE 14. The gossip utility over an increasing number of nodes involved in DoS attempts and related to Case C with fanout 3.




FIGURE 16. Analysis of the overhead over an increasing number of nodes involved in a DoS attempt, related to Case C with fanout 3.




FIGURE 15. Analysis of the success rate over an increasing number of nodes involved in a DoS attempt, related to Case C with fanout 3.




FIGURE 17. Analysis of the power consumption with TLS and our approach being applied with and without DoS attacks, related to Case C with fanout 3.

Fig. 14 presents results for the utility of the gossip messages without and with our extensions, referred to as “without protection”, and “with protection” cases, accordingly. Results related to DoS attacks analyzed for different numbers of colluded nodes sending fake messages (whose number follows the fanout) are presented in Fig. 14 for the disaster Case C.

Generally speaking, gossiping does not exhibit a utility equal to 1, as reported in [44]. However, in the post-disaster scenario, the appearance of a DoS attack can decrease the utility even more. From Fig. 14, it is evident that the utility is higher for our gossiping approach than in the case of no protection. It proves that our approach helps in preventing DoS attack attempts. Similar trends were noticed for the other disaster cases A and B. However, we can notice that the utility decreases slightly, as there are cases when fake feedback messages are considered valid.

Fig. 15 presents the comparison of the success rate values for our gossiping scheme (“with protection”) with the basic gossiping scheme (“no protection”) and the reference success rate of 1.0 achievable under no attacks for Case C. It shows that the reduction in the gossip utility due to DoS attacks correlates with a decrease in the attainable success rate, as already presented before. Our approach can obtain the success rate closer to the one exhibited by the case without

any attacks. The two trends (i.e., orange and grey curves) do not follow the same pattern due to the message suppression mechanism of our approach. It, in turn, shows the successful feedback deception.

Last, Fig. 16 presents the measured overhead in terms of additional messages exchanged over the network for Case C. From Fig. 16, it is clear that in the case of our approach, the overhead is considerably reduced, which proves a protection degree against the adverse effects of such attacks.

To show the power efficiency of our approach, we analyzed the power consumption by considering the cost of receiving a gossip message of about 1 kB taken from [74] for a representative resource-constrained device. We compared it with the power consumption of the handshake scheme of Transport Layer Security (TLS) [75] – a standard protection strategy against DoS attacks to protect communications and authenticate the message sender and cryptographic primitives to guarantee message authentication and integrity. The port randomization does not imply any worsening in energy consumption. The solution based on identity management and removing those nodes detected as malicious does not have a direct worsening of energy consumption when sending messages, but when maintaining the correct view of identities within the system. We are aware that TLS has more

features than our solution, so that the comparison may look unfair. However, as protection against DoS attacks requires to prevent eavesdropping, tampering and message forgery, TLS is a perfect candidate to protect gossip against DoS attacks. Concerning this issue, we are interested in their comparison and it will be the base of our future work for the protection against message forgery by optimizing the use of TLS and similar cryptographic primitives.

Five protocols, shown in Fig. 17, were analyzed related to Case C, representing technology-related disasters: conventional gossiping (“no protection”), classic gossiping with TLS, and our gossiping scheme – both analyzed for two scenarios of DoS attacks and no attacks being applied. From Fig. 17, we can notice that TLS introduces a fixed overhead due to the handshake performed between the sender and receiver of the gossip message, even if there are no attacks. It is worth noting that this is not the case for our approach. When a DoS attack is introduced, the overhead of TLS grows and is visibly higher than the one obtained for our message suppression scheme.

V. CONCLUSION

This article addressed the problem of the resilience of a retransmission-based recovery scheme called gossip, in particular in the case when challenged by DoS attacks in the post-disaster scenario. The attack model for this distributed algorithm was presented, and the related literature on the available protection means was discussed. We noticed that countermeasures for DoS attacks rely mainly on cryptographic primitives or trust management, which do not scale well and are inefficient in post-disaster scenarios. We proposed a solution based on game theory and showed that it has the merit to reduce the occurrence of overloading due to such attacks. It is clear that there is room for improvement, and further work is needed to have a fully secure approach against DoS attacks, such as adequately integrating cryptographic primitives to support the secure gossiping. As mentioned above in the paper, a possible future direction of our work consists in properly integrating our approach with cryptographic primitives and to trigger opportunistically the use of TLS or similar cryptographic primitives when our approach is starting to become compromised. We will investigate the use of signaling games [76] for this aim.

The open problem for our approach is that the adversary may issue not only fake requests for retransmissions but also the feedback, which may compromise the decision making. Moreover, the attacker may be able to alter the content of the exchanged messages using cooperative transmission protocols, making our approach compromised. As in this article, we did not rely on cryptographic primitives, to avoid forging the fake feedback and malicious modification of valid messages, we can further introduce heuristics to assure the received response are valid. One possible solution that we consider is based on the reception time: a node monitors the latency needed by another node to send feedback and assumes the received feedback messages to be valid only if

received within the estimated latency. It is thus possible to have a more sophisticated approach by modelling the interaction among the nodes as a dynamic Bayesian signalling game [76] and accepting/refusing incoming messages based on the node assumptions of the sender being honest or malicious, similarly as applied in [77] or [78] in the context of the positioning systems and trust management. We leave this more sophisticated approach for future work.

ACKNOWLEDGMENT

This article is based upon work from COST Action CA15127 (Resilient communication services protecting end-user applications from disaster-based failures–RECODIS).

REFERENCES

- [1] M. Wollschlaeger, T. Sauter, and J. Jasperneite, “The future of industrial communication: Automation networks in the era of the Internet of Things and industry 4.0,” *IEEE Ind. Electron. Mag.*, vol. 11, no. 1, pp. 17–27, Mar. 2017.
- [2] A. Mauthe, D. Hutchison, E. K. Cetinkaya, I. Ganchev, J. Rak, J. P. G. Sterbenz, M. Gunkel, P. Smith, and T. Gomes, “Disaster-resilient communication networks: Principles and best practices,” in *Proc. 8th Int. Workshop Resilient Netw. Design Modeling (RNDM)*, Sep. 2016, pp. 1–10.
- [3] M. Furdek, L. Wosinska, R. Gosciencin, K. Manousakis, M. Aibin, K. Walkowiak, S. Ristov, M. Gushev, and J. L. Marzo, “An overview of security challenges in communication networks,” in *Proc. 8th Int. Workshop Resilient Netw. Design Modeling (RNDM)*, Sep. 2016, pp. 1–8.
- [4] J. Rak, D. Hutchison, E. Calle, T. Gomes, M. Gunkel, P. Smith, J. Tapolcai, S. Verbrugge, and L. Wosinska, “RECODIS: Resilient communication services protecting end-user applications from disaster-based failures,” in *Proc. 18th Int. Conf. Transparent Opt. Netw. (ICTON)*, Jul. 2016, pp. 1–4.
- [5] A. Kwasinski, W. W. Weaver, P. I. Chapman, and P. T. Krein, “Telecommunications power plant damage assessment for hurricane Katrina—site survey and follow-up results,” *IEEE Syst. J.*, vol. 3, no. 3, pp. 277–287, Sep. 2009.
- [6] R. Gosciencin, K. Walkowiak, M. Klinkowski, and J. Rak, “Protection in elastic optical networks,” *IEEE Netw.*, vol. 29, no. 6, pp. 88–96, Nov./Dec. 2015.
- [7] F. Dikbiyik, M. Tornatore, and B. Mukherjee, “Minimizing the risk from disaster failures in optical backbone networks,” *J. Lightw. Technol.*, vol. 32, no. 18, pp. 3175–3183, Sep. 15, 2014.
- [8] M. F. Habib, M. Tornatore, F. Dikbiyik, and B. Mukherjee, “Disaster survivability in optical communication networks,” *Comput. Commun.*, vol. 36, no. 6, pp. 630–644, Mar. 2013.
- [9] J. Rak, *Resilient Routing in Communication Networks*. Cham, Switzerland: Springer, 2015.
- [10] X. Long, D. Tipper, and T. Gomes, “Measuring the survivability of networks to geographic correlated failures,” *Opt. Switching Netw.*, vol. 14, pp. 117–133, Aug. 2014.
- [11] A. F. Hansen, A. Kvalbein, T. Cicic, and S. Gjessing, “Resilient routing layers for network disaster planning,” in *Proc. Int. Conf. Netw. (ICN)*, in Lecture Notes in Computer Science, vol. 3421. Berlin, Germany: Springer, 2005, pp. 1097–1105.
- [12] Reuters. *Experts Warn of Substantial Risk of WMD Attack*. Accessed: Apr. 2020. [Online]. Available: <http://research.lifeboat.com/lugar.htm>
- [13] C. M. Machuca, S. Secci, P. Vizarreta, F. Kuipers, A. Gouglidis, D. Hutchison, S. Jouet, D. Pezaros, A. Elmokashfi, P. Heegaard, S. Ristov, and M. Gusev, “Technology-related disasters: A survey towards disaster-resilient software defined networks,” in *Proc. 8th Int. Workshop Resilient Netw. Design Modeling (RNDM)*, Sep. 2016, pp. 1–8.
- [14] L. Martins, R. Girao-Silva, L. Jorge, T. Gomes, F. Musumeci, and J. Rak, “Interdependence between power grids and communication networks: A resilience perspective,” in *Proc. 13th Int. Conf. Design Reliable Commun. Netw.*, Mar. 2017, pp. 1–7.
- [15] F. Palmieri, U. Fiore, A. Castiglione, F.-Y. Leu, and A. de Santis, “Analyzing the Internet stability in presence of disasters,” in *Proc. CD-ARES Workshops*, 2013, pp. 253–268.
- [16] A. Di Pietro, S. Panzieri, and A. Gasparri, “Situational awareness using distributed data fusion with evidence discounting,” in *Critical Infrastructure Protection IX*. Cham, Switzerland: Springer, 2015, pp. 281–296.

- [17] M. Kobayashi, "Experience of infrastructure damage caused by the great east japan earthquake and countermeasures against future disasters," *IEEE Commun. Mag.*, vol. 52, no. 3, pp. 23–29, Mar. 2014.
- [18] TPN/Lusa. (Oct. 2017). *Communications Failure Partly to Blame for June Fire Deaths—Report, The Portugal New*. [Online]. Available: <http://www.theportugalnews.com/news/communications-failure-partly-to-blame-for-june-fire-deaths-report/43524>
- [19] M. Cinque, D. Cotroneo, C. Esposito, and M. Fiorentino, "Secure crisis information sharing through an interoperability framework among first responders: The SECTOR practical experience," in *Proc. IEEE 13th Int. Conf. Wireless Mobile Comput., Netw. Commun. (WiMob)*, Oct. 2017, pp. 316–323.
- [20] T. Gomes, J. Tapolcai, C. Esposito, D. Hutchison, F. Kuipers, J. Rak, A. de Sousa, A. Iossifides, R. Travanca, J. Andre, L. Jorge, L. Martins, P. O. Ugalde, A. Pasic, D. Pezaros, S. Jouet, S. Secci, and M. Tornatore, "A survey of strategies for communication networks to protect against large-scale natural disasters," in *Proc. 8th Int. Workshop Resilient Netw. Design Modeling (RNDM)*, Sep. 2016, pp. 1–12.
- [21] S. Lin, D. J. Costello, and M. J. Miller, "Automatic-repeat-request error-control schemes," *IEEE Commun. Mag.*, vol. 22, no. 12, pp. 5–17, Dec. 1984.
- [22] C. Esposito, D. Cotroneo, and S. Russo, "On reliability in publish/subscribe services," *Comput. Netw.*, vol. 57, no. 5, pp. 1318–1343, Apr. 2013.
- [23] F. Baccelli, A. Chaintreau, Z. Lin, A. Riabov, and S. Sahu, "Scalability of reliable group communication using overlays," in *Proc. IEEE INFOCOM*, Mar. 2004, p. 430.
- [24] W.-P. K. Yiu, K.-F. S. Wong, S.-H. G. Chan, W.-C. Wong, Q. Zhang, W.-W. Zhu, and Y.-Q. Zhang, "Lateral error recovery for media streaming in application-level multicast," *IEEE/ACM Trans. Multimedia*, vol. 8, no. 2, pp. 219–232, Apr. 2006.
- [25] G. Tan and S. A. Jarvis, "Improving the fault resilience of overlay multicast for media streaming," *IEEE Trans. Parallel Distrib. Syst.*, vol. 18, no. 6, pp. 721–734, Jun. 2007.
- [26] P. Costa, M. Migliavacca, G. P. Picco, and G. Cugola, "Epidemic algorithms for reliable content-based publish-subscribe: An evaluation," in *Proc. 24th Int. Conf. Distrib. Comput. Syst.*, Mar. 2004, pp. 552–561.
- [27] C. Esposito, M. Platania, and R. Beraldi, "Reliable and timely event notification for publish/subscribe services over the Internet," *IEEE/ACM Trans. Netw.*, vol. 22, no. 1, pp. 230–243, Feb. 2014.
- [28] F. Campos and J. Pereira, "Improving the scalability of DPWS-based networked infrastructures," *CoRR*, vol. abs/1407.8546, pp. 1–28, Jul. 2014.
- [29] C. Esposito, A. Castiglione, F. Palmieri, and M. Ficco, "Distributed strategic learning for effective gossiping in wireless networks," in *Proc. IEEE Conf. Comput. Commun. Workshops (INFOCOM WKSHPs)*, Apr. 2016, pp. 509–514.
- [30] J. W. Byers, M. Luby, and M. Mitzenmacher, "A digital fountain approach to asynchronous reliable multicast," *IEEE J. Sel. Areas Commun.*, vol. 20, no. 8, pp. 1528–1540, Oct. 2002.
- [31] M. Ghaderi, D. Towsley, and J. Kurose, "Reliability gain of network coding in lossy wireless networks," in *Proc. IEEE INFOCOM-27th Conf. Comput. Commun.*, Apr. 2008, pp. 2171–2179.
- [32] E. Androutaki, A. Barger, V. Bortnikov, C. Cachin, K. Christidis, A. D. Caro, D. Enyeart, C. Ferris, G. Laventman, Y. Manevich, S. Muralidharan, C. Murthy, B. Nguyen, M. Sethi, G. Singh, K. Smith, A. Sorniotti, C. Stathakopoulou, M. Vukolić, S. W. Cocco, J. Yellick, "Hyperledger fabric: A distributed operating system for permissioned blockchains," in *Proc. 13th EuroSys Conf.*, Apr. 2018, pp. 1–15.
- [33] S. Bolognani and S. Zampieri, "A gossip-like distributed optimization algorithm for reactive power flow control," *IFAC Proc. Volumes*, vol. 44, no. 1, pp. 5700–5705, Jan. 2011.
- [34] V. F. S. Mota, D. F. Macedo, and J. M. S. Nogueira, "An hierarchical routing protocol for opportunistic emergency networks," in *Proc. 7th Latin Amer. Netw. Conf. (LANC)*, 2012, pp. 36–43.
- [35] M. Radenkovic, A. Walker, and L. Bai, "Towards better understanding the challenges of reliable and trust-aware critical communications in the aftermath of disaster," in *Proc. 14th Int. Wireless Commun. Mobile Comput. Conf. (IWCMC)*, Jun. 2018, pp. 648–653.
- [36] A. Seba, N. Nouali-Taboudjemat, N. Badache, and H. Seba, "A review on security challenges of wireless communications in disaster emergency response and crisis management situations," *J. Netw. Comput. Appl.*, vol. 126, pp. 150–161, Jan. 2019.
- [37] P. Asuquo, H. Cruickshank, Z. Sun, and G. Chandrasekaran, "Analysis of DoS attacks in delay tolerant networks for emergency evacuation," in *Proc. 9th Int. Conf. Next Gener. Mobile Appl., Services Technol.*, Sep. 2015, pp. 228–233.
- [38] G. Badishi, I. Keidar, and A. Sasson, "Exposing and eliminating vulnerabilities to denial of service attacks in secure gossip-based multicast," in *Proc. Int. Conf. Dependable Syst. Netw.*, Jun./Jul. 2004, pp. 223–232.
- [39] J. Decouchant, S. B. Mokhtar, A. Petit, and V. Quema, "PAG: Private and accountable gossip," in *Proc. IEEE 36th Int. Conf. Distrib. Comput. Syst. (ICDCS)*, Jun. 2016, pp. 35–44.
- [40] S. Durr-e-Zehra Naqvi and M. H. Islam, "Incorporating data integrity in gossip based overlay networks," in *Proc. 4th Int. Conf. Emerg. Technol.*, Oct. 2008, pp. 120–125.
- [41] D. Malkhi, M. K. Reiter, O. Rodeh, and Y. Sella, "Efficient update diffusion in byzantine environments," in *Proc. 20th IEEE Symp. Reliable Distrib. Syst.*, Oct. 2001, pp. 90–98.
- [42] Y. M. Minsky and F. B. Schneider, "Tolerating malicious gossip," *Distrib. Comput.*, vol. 16, no. 1, pp. 49–68, Feb. 2003.
- [43] S. K. Tatarave, S. Tripathy, and S. Peri, "S-Gossip: Security enhanced gossip protocol for unstructured P2P networks," in *Proc. Int. Conf. Distrib. Comput. Internet Technol.*, 2015, pp. 288–298.
- [44] C. Esposito, A. Castiglione, F. Palmieri, and M. Ficco, "Improving the gossiping effectiveness with distributed strategic learning," *Future Gener. Comput. Syst.*, vol. 71, pp. 221–233, Jun. 2017.
- [45] J. M. Kizza, *Guide to Computer Network Security*. London, U.K.: Springer-Verlag, 2017.
- [46] *Transmission Control Protocol specifications*, document RFC 793, IETF, Sep. 1981. [Online]. Available: <https://tools.ietf.org/html/rfc793>
- [47] C. Diot, W. Dabbous, and J. Crowcroft, "Multipoint communication: A survey of protocols, functions, and mechanisms," *IEEE J. Sel. Areas Commun.*, vol. 15, no. 3, pp. 277–290, Apr. 1997.
- [48] S. M. Hedetniemi, S. T. Hedetniemi, and A. L. Liestman, "A survey of gossiping and broadcasting in communication networks," *Networks*, vol. 18, no. 4, pp. 319–349, 1988.
- [49] G. P. Jesi, E. Mollona, S. K. Nair, and M. van Steen, "Prestige-based peer sampling service: Interdisciplinary approach to secure gossip," in *Proc. ACM Symp. Appl. Comput. (SAC)*, 2009, pp. 1209–1213.
- [50] M. Jelasity, R. Guerraoui, A. M. Kermerrec, and M. Van Steen, "The peer sampling service: Experimental evaluation of unstructured gossip-based implementations," in *Proc. ACM/FIP/USENIX Int. Conf. Distrib. Syst. Platforms Open Distrib. Process.*, 2004, pp. 79–98.
- [51] H. D. Johansen, R. V. Renesse, Y. Vigfusson, and D. Johansen, "Fireflies: A secure and scalable membership and gossip service," *ACM Trans. Comput. Syst.*, vol. 33, no. 2, pp. 1–32, Jun. 2015.
- [52] M. Gerla and L. Kleinrock, "Vehicular networks and the future of the mobile Internet," *Comput. Netw.*, vol. 55, no. 2, pp. 457–469, Feb. 2011.
- [53] K. P. Kihlstrom and R. S. Elliott, "Performance of an intrusion-tolerant gossip protocol," in *Proc. 21st IASTED Int. Conf. Parallel Distrib. Comput. Syst.*, 2009, pp. 63–68.
- [54] D. Malkhi, Y. Mansour, and M. K. Reiter, "Diffusion without false rumors: On propagating updates in a byzantine environment," *Theor. Comput. Sci.*, vol. 299, nos. 1–3, pp. 289–306, Apr. 2003.
- [55] G. P. Jesi, D. Hales, and M. van Steen, "Identifying malicious peers before it's too late: A decentralized secure peer sampling service," in *Proc. 1st Int. Conf. Self-Adapt. Self-Organizing Syst. (SASO)*, Jul. 2007, pp. 237–246.
- [56] H. C. Li, A. Clement, E. L. Wong, J. Napper, I. Roy, L. Alvisi, and M. Dahlin, "BAR gossip," in *Proc. 7th Symp. Operating Syst. Design Implement.*, 2006, pp. 191–204.
- [57] O. F. W. Onifade, O. B. Longe, and A. S. Akanmu, "Congestion controlled anonymous gossip: A scalable method for providing probabilistic guarantees to multicast reliability in mobile ad-hoc networks," *Int. J. Comput. ICT Res.*, vol. 1, no. 2, pp. 42–49, 2007.
- [58] R. Chandra, V. Ramasubramanian, and K. Birman, "Anonymous gossip: Improving multicast reliability in mobile ad-hoc networks," in *Proc. 21st Int. Conf. Distrib. Comput. Syst.*, 2001, pp. 275–283.
- [59] C. Esposito, S. Russo, R. Beraldi, M. Platania, and R. Baldoni, "Achieving reliable and timely event dissemination over WAN," in *Proc. 13th Int. Conf. Distrib. Comput. Netw.*, in Lecture Notes in Computer Science, vol. 7129. Berlin, Germany: Springer, 2012, pp. 265–280.
- [60] F. Mattern, "Efficient algorithms for distributed snapshots and global virtual time approximation," *J. Parallel Distrib. Comput.*, vol. 18, no. 4, pp. 423–434, Aug. 1993.
- [61] K. M. Chandy and L. Lamport, "Distributed snapshots: Determining global states of distributed systems," *ACM Trans. Comput. Syst.*, vol. 3, no. 1, pp. 63–75, Feb. 1985.

- [62] M. Osborne and A. Rubinstein, *A Course in Game Theory*. Cambridge, MA, USA: MIT Press, 1994.
- [63] B. Yang and M. Johansson, "Distributed optimization and games: A tutorial overview," in *Networked Control Systems* (Lecture Notes in Control and Information Sciences), vol. 406. London, U.K.: Springer, 2010, pp. 109–148.
- [64] A. Rowstron, A.-M. Kermarrec, M. Castro, and P. Druschel, "SCRIBE: The design of a large-scale event notification infrastructure," in *Proc. Int. Workshop Netw. Group Commun.*, 2001, pp. 30–43.
- [65] H. Lu, "On the existence of pure-strategy Nash equilibrium," *Econ. Lett.*, vol. 94, no. 3, pp. 459–462, Mar. 2007.
- [66] H. Tembine, *Distributed Strategic Learning for Wireless Engineers*. Boca Raton, FL, USA: CRC Press, 2012.
- [67] A. Varga, "OMNeT++," in *Modeling and Tools for Network Simulation*. Berlin, Germany: Springer-Verlag, 2010, pp. 35–59.
- [68] P. Druschel and A. Rowstron, "PASTRY: Scalable, decentralized object location, and routing for large-scale peer-to-peer systems," in *Middleware* (Lecture Notes in Computer Science), vol. 2218. Berlin, Germany: Springer-Verlag, 2001, pp. 329–350.
- [69] E. O. Elliott, "Estimates of error rates for codes on burst-noise channels," *Bell Syst. Tech. J.*, vol. 42, no. 5, pp. 1977–1997, Sep. 1963.
- [70] G. Hasslinger and O. Hohlfeld, "The Gilbert-Elliott model for packet loss in real time services on the Internet," in *Proc. 14th GI/ITG Conf. Measuring, Modeling Eval. Comput. Commun. Syst.*, Mar./Apr. 2008, pp. 1–15.
- [71] S. Saha, S. Nandi, P. S. Paul, V. K. Shah, A. Roy, and S. K. Das, "Designing delay constrained hybrid ad hoc network infrastructure for post-disaster communication," *Ad Hoc Netw.*, vol. 25, pp. 406–429, Feb. 2015.
- [72] M. I. Channa and K. M. Ahmed, "A reliable routing scheme for post-disaster ad hoc communication networks," *J. Commun.*, vol. 6, no. 7, pp. 549–557, Oct. 2011.
- [73] L. Li, D. Alderson, W. Willinger, and J. Doyle, "A first-principles approach to understanding the Internet's router-level topology," in *Proc. Conf. Appl., Technol., Archit., Protocols Comput. Commun. (SIGCOMM)*, 2004, pp. 3–14.
- [74] Atmel. *Atmel AT86RF230 Datasheet*. Accessed: May 2020. [Online]. Available: <http://www.atmel.com/Images/doc5131.pdf>
- [75] R. Mzid, M. Boujelben, H. Youssef, and M. Abid, "Adapting TLS handshake protocol for heterogenous IP-based WSN using identity based cryptography," in *Proc. Int. Conf. Wireless Ubiquitous Syst.*, Oct. 2010, pp. 1–8.
- [76] D. Lewis, *Convention: A Philosophical Study*. Cambridge, MA, USA: Harvard Univ. Press, 1969.
- [77] C. Esposito and C. Choi, "Signaling game based strategy for secure positioning in wireless sensor networks," *Pervasive Mobile Comput.*, vol. 40, pp. 611–627, Sep. 2017.
- [78] C. Esposito, A. Castiglione, and F. Palmieri, "Information theoretic-based detection and removal of slander and/or false-praise attacks for robust trust management with Dempster-Shafer combination of linguistic fuzzy terms," *Concurrency Comput., Pract. Exp.*, vol. 30, no. 3, p. e4302, Feb. 2018.
- [79] Y. Chen, S. Shakkottai, and J. G. Andrews, "On the role of mobility for multmessage gossip," *IEEE Trans. Inf. Theory*, vol. 59, no. 6, pp. 3953–3970, Jun. 2013.
- [80] E. Hopkins and M. Posch, "Attainability of boundary points under reinforcement learning," *Games Econ. Behav.*, vol. 53, no. 1, pp. 110–125, Oct. 2005.
- [81] J. R. Marden, H. P. Young, and L. Y. Pao, "Achieving Pareto optimality through distributed learning," *SIAM J. Control Optim.*, vol. 52, no. 5, pp. 2753–2770, Jan. 2014.
- [82] A. Vasin, "The Folk theorems in the framework of evolution and cooperation," in *Advances in Dynamic Games*. Boston, MA, USA: Birkhäuser, 2006, pp. 197–207.
- [83] J. D. Cohen, S. M. McClure, and A. J. Yu, "Should i stay or should i go? How the human brain manages the trade-off between exploitation and exploration," *Phil. Trans. Roy. Soc. B, Biol. Sci.*, vol. 362, no. 1481, pp. 933–942, May 2007.



CHRISTIAN ESPOSITO (Member, IEEE) received the Ph.D. degree in computer engineering and automation from the University of Napoli Federico II, in 2009. He is currently a tenured Assistant Professor with the University of Salerno, a non-tenured Assistant Professor with the University of Napoli Federico II, and a Research Fellow with the University of Salerno and the Institute for High Performance Computing and Networking, The National Research Council (ICAR-CNR). He has been involved in the organization of about 40 international conferences workshops. His research interests include reliable and secure communications, middleware, distributed systems, positioning systems, multi-objective optimization, and game theory. He has served as a Reviewer and the Guest Editor for several international journals and conferences (with about 200 completed reviews). He is also an Associate Editor of IEEE Access.



ZHONGLIANG ZHAO (Member, IEEE) received the Ph.D. degree from the University of Bern, in 2014. In 2014, he holds an appointment of a Senior Researcher with the University of Bern. He is currently an Associate Professor with the School of Electronic and Information Engineering, Beihang University, China. His main research interests include UAC ad-hoc networking, smart transportation networks, and edge intelligence for communication networks. He has been the TPC Co-Chair of IEEE WONS 2019, the Chair of IEEE INFOCOM 2020 EINSTEIN workshop, and a TPC Member of many conferences.



JACEK RAK (Senior Member, IEEE) received the M.Sc., Ph.D., and D.Sc. (Habilitation) degrees from the Gdańsk University of Technology, Gdańsk, Poland, in 2003, 2009, and 2016, respectively.

From 2016 to 2020, he was leading the COST CA15127 Action "Resilient Communication Services Protecting End-User Applications from Disaster-Based Failures" (RECODIS) involving over 170 members from 31 countries. He is currently an Associate Professor and the Head of the Department of Computer Communications, Gdańsk University of Technology. He has authored over 100 publications, including the book *Resilient Routing in Communication Networks* (Springer, 2015). His main research interests include the resilience of communication networks and networked systems.

Dr. Rak is a member of the Editorial Board of Optical Switching and Networking, Elsevier and the Founder of the International Workshop on Resilient Networks Design and Modeling (RNDM). He has also served as a TPC member of numerous conferences and journals. He has been the General Chair of ITS-T 2017 and MMM-ACNS 2017, the General Co-Chair of NETWORKS 2016, the TPC Chair of ONDM 2017, and the TPC Co-Chair of IFIP Networking 2019.

...