Emergent Overlays for Adaptive MANET Broadcast

Raziel Carvajal-Gómez⋆, Yérom-David Bromberg◇, Yehia Elkhatib•, Laurent Réveillère▽ and Etienne Rivière⋆

⋆ ICTEAM, UCLouvain, Belgium – raziel.carvajal@uclouvain.be, etienne.riviere@uclouvain.be
◇ Univ Rennes, Inria, CNRS, IRISA, France – david.bromberg@irisa.fr
• School of Computing and Communications, Lancaster University, United Kingdom – y.elkhatib@lancaster.ac.uk
▽ Univ. Bordeaux, CNRS, Bordeaux INP, LaBRI, UMR 5800, France – laurent.reveillere@u-bordeaux.fr

Abstract—Mobile Ad-Hoc Networks (MANETs) allow distributed applications where no fixed network infrastructure is available. MANETs use wireless communication subject to faults and uncertainty, and must support efficient *broadcast*. Controlled flooding is suitable for highly-dynamic networks, while overlay-based broadcast is suitable for dense and more static ones. Density and mobility vary significantly over a MANET deployment area. We present the design and implementation of *emergent overlays* for efficient and reliable broadcast in heterogeneous MANETs. This adaptation technique allows nodes to automatically switch from controlled flooding to the use of an overlay. Interoperability protocols support the integration of both protocols in a single heterogeneous system. Coordinated adaptation policies allow regions of nodes to autonomously and collectively emerge and dissolve overlays. Our simulation of the full network stack of 600 mobile nodes shows that emergent overlays reduce energy consumption, and improve reliability and coverage compared to single protocols and to two previously-proposed adaptation techniques.

Index Terms—MANET, Broadcast, Adaptation, Reliability

I. INTRODUCTION

Mobile Ad-hoc Networks (MANETs) are formed of mobile nodes relying on direct wireless communication subject to faults and uncertainty. MANETs allow networked applications in situations where no infrastructure is available, or where the available infrastructure does not suffice [1], [2]. A first target example is a music festival in a rural area, where setting up a permanent communication infrastructure is not economically viable. MANETs let a large number of attendees communicate with each other despite the saturation of the local cellular network. Another example is an emergency scenario where teams of rescuers coordinate over an area where communication infrastructures and power distribution networks are no longer available [3], [4].

MANETs typically involve a large number of nodes with arbitrary mobility. For human-attached nodes, models allow understanding general mobility patterns [5], [6]. These models concur with two major observations. First, *mobility* is heterogeneous, with nodes alternating between periods of high and low speed. Second, there are points of attraction around which users tend to aggregate [7], leading to zones of heterogeneous *density*. Heterogeneity in space was confirmed by an experiment at the Paléo music festival in Switzerland [8]. The authors performed an estimation of the population size and of its spatial distribution by collecting unique MAC addresses from a significant portion of attendees' mobile phones that were discoverable via bluetooth. The authors report differences in node density of up to one order of magnitude (10x) between different $15m \times 15m$ areas. A similar study during a street festival in Belgium [9] reports a high heterogeneity of density, and confirms that users tend to alternate between lowly-mobile zones (points of interests) with a higher number of users, and zones of lower density (3x in the collected sample).

Heterogeneity in space and mobility directly impacts data dissemination in MANETs, and particularly *broadcast* operations. Broadcast allows delivering messages from any source to all nodes in the network, e.g. to forward live announcements to all users of a music festival mobile application. Efficient routing also depends on broadcast [1]. Offering efficient and reliable broadcast in a MANET requires an appropriate forwarding strategy. The main hindrance to reliability is the risk of collisions, when two nodes in wireless range attempt to send at the same time and cancel each others' transmissions. Collisions are not always possible to detect, and are costly to address [10]. A MAC protocol [11] using too many resubmissions may result in lower performance and increase energy costs (the broadcast storm problem [12]).

Broadcast protocols for MANETs aim at reliably delivering messages to all nodes despite collisions. They can be classified in two families [13]: protocols based on *controlled flooding*, and those relying on the creation and maintenance of an *overlay*. With controlled flooding, any node receiving a message decides independently whether or not (and when) to retransmit it using local broadcast. Overlay-based broadcast on the other hand decouples the decision process of which node will, or will not, forward incoming messages, from the actual dissemination. A subset of *relay* nodes are selected *a priori*, through the exchange of control messages, to form a persistent dissemination backbone, or *overlay*.

Evaluation studies of broadcast algorithms [14], [15] show that their performance, reliability, and energy consumption are very sensitive to deployment conditions. Controlled flooding [16], [17] is efficient in sparsely-populated areas and when nodes are highly mobile. Overlay-based broadcast [18]–[20] performs better when nodes are relatively stable. It enables important saving in energy and retransmissions at densely-populated areas. There is no existing protocol that can offer reliable and efficient broadcast in any conditions and support heterogeneous deployments. Protocols of the two families are not designed to be interoperable: deploying different protocols in different regions leads to loss of reliability.
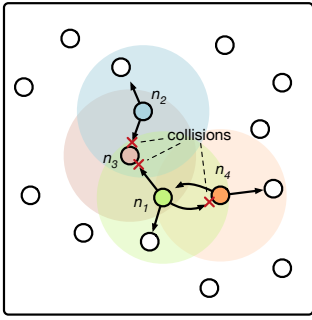
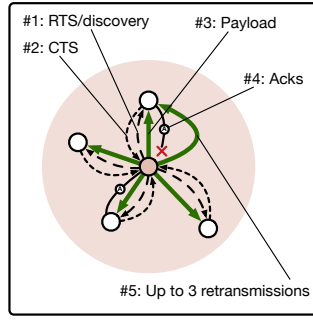Fig. 1. Communication model and illustration of network collisions.

Fig. 2. BMW MAC protocol for local broadcast: neighborhood discovery and frame retransmissions.

Contributions. We target efficient broadcast in *heterogeneous* MANETs. Our motivation is that controlled flooding is efficient in the general case and using an overlay is beneficial in zones where density increases while mobility decreases, typical of human behavior at points of interest. We contribute the novel notion of ***emergent overlays***, allowing a subset of the nodes in a MANET to autonomously decide to switch from the use of controlled flooding to the bootstrap and use of an overlay. The decision to emerge (or later dissolve) an overlay is based on local indicators of density and mobility, with the autonomous coordination of all nodes within a region. Interoperability protocols enable reliable broadcast between zones using heterogeneous protocols.

- We present our system model and our building blocks (§II): the Broadcast Medium Window (BMW), a MAC protocol for local broadcast [21]; *Counter-based* controlled flooding [22]; and the *MPR* overlay [20];
- We define the requirements for emergent overlays (§III);
- We address the interoperability challenge with a protocol allowing to guarantee continuity in the broadcast across heterogeneous zones using different protocols (§IV);
- We present how nodes are able to collect *observables* about their environment to drive the emergence of overlays, and how this emergence is coordinated autonomously across nodes in an homogeneous region of the MANET (§V);
- We evaluate emergent overlays using a full-stack simulation in OMNeT++/INET [23] (§VI). Our results show that emergent overlays allow reaching very good dissemination guarantees across heterogeneous zones, while reducing the load in the network and the amount of collisions. We show that our approach compares favorably to the use of single protocols and to two other adaptive broadcast approaches from the literature [17], [22].

We finally discuss related work (§VII) and conclude (§VIII).

II. KEY CONCEPTS

We start by detailing our system and communication models. We follow-up with the description of the three base protocols we combine in our approach.

A. System and communication models

A set N of nodes $n_i, n_j, \ldots$ are mobile in a 2-dimensional space. We do not consider the availability of geolocation information (GPS) on the devices. Nodes communicate using wireless transceivers with a limited and fixed transmission range Tx. Nodes cannot send and receive simultaneously: transceivers are either in sending or listening mode. When nodes emit network frames, only nodes in their coverage area can receive them. We refer to this as a *local broadcast*. Figure 1 presents the coverage areas for nodes $n_{1\ldots4}$. For example, node n_3 is in range of both n_1 and n_2.

The main obstacle to reliable local broadcast is the occurrence of *collisions*, which may happen when two nodes have overlapping coverage areas. If nodes n_1 and n_2 send frames at the same time then it is impossible for n_3 to successfully receive both of the frames, or even any of them. The same applies to a node such as n_4 that is in sending mode and cannot receive frames from other nodes in range. Detecting collisions as they happen is not possible [11]. Reliable local broadcast must therefore rely on a medium access control protocol.

B. MAC protocol for local broadcast

A MAC (*Medium Access Control*) protocol arbitrates the use of the wireless medium to minimize collisions and their impact on communication reliability. It detects missed receptions and implements frame retransmissions when necessary.

We use broadcast medium window (BMW), a MAC protocol for reliable local broadcast [21]. The protocol is illustrated by Figure 2. BMW relies on *virtual carrier sensing*, as used for point-to-point communication in the classical IEEE 802.11 coordination protocol [11]. A node initiates a local broadcast by sending a RTS (Request To Send) frame, with the goal of *discovering* local neighbors. These neighbors reply with CTS (Clear To Send) messages containing their identity. Physical carrier sensing [11] reduces the risk of collisions between CTS replies: nodes wait for an initial random duration and probe the wireless medium before sending; If it is currently used, they back off and retry after an (exponentially increasing) duration. The payload frame is then transmitted, and neighbors acknowledge its reception, again using physical carrier sensing to avoid collisions. Up to two additional retransmissions happen in case of missed or error-bearing acknowledgments.

C. (Global) Broadcast protocols

The goal of *global* broadcast (hereafter referred simply as broadcast) is to reliably send a message to all nodes in the network, despite mobility and collisions. The source can be any node in the network, and the message may require multiple frames. Nodes *may* forward a message using local broadcast upon its first reception. We denote such nodes as *relay* nodes. Subsequent receptions are *duplicates* that are simply ignored. The target is to reduce the number of relay nodes while ensuring a high (ideally full) *reach*, i.e. the average proportion of nodes that receive each broadcast.[1] Broadcast protocols use

---

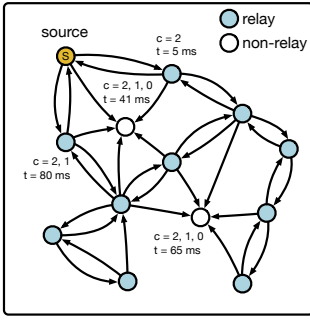[1] Incomplete broadcasts may be completed using periodic anti-entropy [24].

Fig. 3. Broadcast using counter-based controlled flooding, example with $c = 2$, and details of counter and timer initial values for 4 sample nodes.
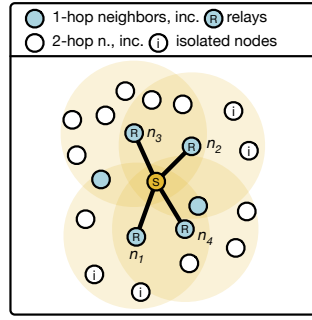

Fig. 4. MPR selection of overlay nodes: n_1 and n_2 are selected during the first phase as they connect to isolated nodes, n_3 and n_4 in the second phase as they allow to cover the most two-hop neighbors.

Algorithm 1 Selection of relays at node n_i with MPR

```
1:  input
2:      OneH : set                                    ▷ one-hop neighbors
3:      TwoH[ ] : map (n_{j≠i} ∈ OneH) → {n_k, … }    ▷ two-hop neighbors

4:  procedure chooseRelays()
5:      uncovered = ⋃_{n_j ∈ OneH} TwoH[n_j]          ▷ all two-hop neighbors
6:      relays ← ∅
    // Phase 1: select 1-hop neighbors connecting isolated 2-hop neighbors
7:      for every n_j in OneH do
8:          if ∃ n_iso ∈ TwoH[n_j] s.t. ∀n_{k≠j}, n_iso ∉ TwoH[n_k] then
9:              relays ← relays ∪ {n_j}
10:             uncovered ← uncovered - TwoH[n_j]
    // Phase 2: select high-degree 1-hop neighbors until coverage
11:     while uncovered ≠ ∅ do
12:         select n_j ∈ OneH maximizing |TwoH[n_j] ∩ uncovered|
13:         relays ← relays ∪ {n_j}
14:         uncovered ← uncovered - TwoH[n_j]
15:     return relays
```

---

either controlled flooding or an overlay [13]. We present next the two protocols, one of each type, that we later combine to enable emergent overlays.

1) Counter-based controlled flooding: Uncontrolled ("pure") flooding is the simplest broadcast protocol where every node acts as a relay for every message. It generates a large number of duplicates, and many concurrent retransmissions in dense areas. The risk of collision may be reduced by waiting before retransmitting, either for a random amount of time [12] or, when GPS coordinates are available, for a time that depends on a local density estimation [25] but these mitigations do not reduce the number of duplicates.

Controlled Flooding (CF) enables nodes to decide *not to* retransmit, when they estimate that their neighborhood has been sufficiently covered. We use the classical counter-based approach [12] that is efficient in a variety of practical situations, as comparative studies have shown [14], [15]. Upon the first reception of a message, a node sets up a timer t and a counter c. Every reception of a duplicate decreases c by one. When t expires, the node only acts as a relay if $c > 0$.

Controlled flooding is intuitively good for regions with high mobility. The decision of whether or not to be a relay is made by each node independently, based on its observation of the local neighborhood at dissemination time. Optimal values for t and c depend on nodes density. For sparse areas, a low value of c can lead to incomplete coverage. In dense areas, a high value of c will lead to many duplicates. We choose to favor coverage in sparse areas and use a random $t \in [10^{-6}, 10^{-4}]$ seconds and $c = 3$. An illustration of counter-based controlled flooding with $c = 2$ is given by Figure 3.

2) Multipoint Relay (MPR): Overlay-based approaches pre-compute which nodes should act as relays. This requires exchange of information between nodes outside of actual message broadcasts. The goal is to form a *connected dominating set*: the combined coverage areas of all relays must contain all nodes in the network. The overlay should be connected but at the same time contain as few relays as possible. It can be *explicit* when nodes are aware of their status of relay

nodes [18], [19]. It may also be implicit, when nodes are informed they should act as a relay from the source or previous relays in the broadcast process. We adopt a protocol of the latter type, MPR [20].

MPR lets each node n_i decide which of its one-hop (direct) neighbors should be instructed to act as relays during a broadcast in which n_i acts as a relay. MPR requires the discovery of all one-hop and two-hop neighbors using periodic control messages. Nodes periodically announce themselves using best-effort local broadcast, i.e. bypassing the BMW MAC protocol. These `hello` messages include the current neighborhood, collected by the MAC protocol in recent exchanges. Information about a neighbor is reset after a few missed announcements, to account for nodes leaving the coverage area. The ideal frequency of announcements depends on the network dynamics. Too frequent announcements bear the risk of extra collisions and energy costs, while infrequent ones may lead to overlays that are no longer connected or dominating.

The selection of relays is detailed by Algorithm 1 and illustrated by Figure 4. The goal is to ensure that all two-hop neighbors are *covered* by at least one relay. The selection is in two phases. The first phase (lines 7–10) selects relays that are *necessary*, i.e. one-hop neighbors that are the only connection to "isolated" two-hop neighbors. The second phase (lines 11–14) selects relays that are *sufficient* to complete the coverage. The heuristic is to add as relay the one-hop neighbors that cover the largest number of two-hop neighbors until all nodes are covered. As the set of uncovered nodes is the union of all two-hop neighbors (line 5), the `while` loop on lines 11–14 always terminate. MPR is a local heuristic that may not guarantee the optimal selection of relays, but it has the advantage of requiring little computation and no GPS coordinates. It also does not require specific synchronization between nodes to agree on an explicit overlay (e.g. to inform nodes who must act as relays), reducing its maintenance costs.

Overlay-based broadcast is expected to be useful in scenarios where nodes density is high, greatly reducing the number
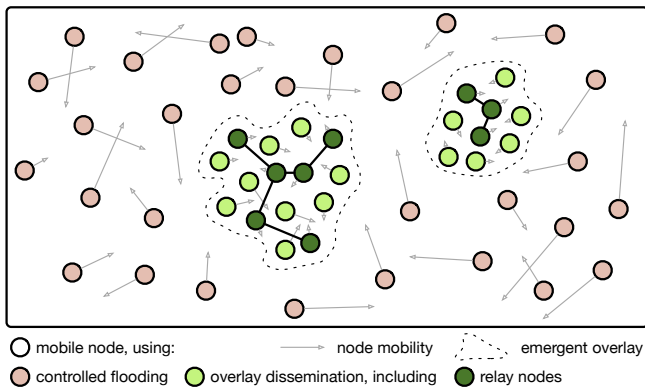
Fig. 5. Illustration of the general principle of emergent overlays.

of relays without requiring a difficult-to-set counter or time value. It is however necessary that the nodes positions be sufficiently stable, or the frequency of control messages will need to be increased to a point where they will create collisions and potentially bear higher costs than for dissemination itself.

III. TOWARDS EMERGENT OVERLAYS

Controlled flooding is efficient when node spatial density is small to moderate. It is also robust in dynamic environments, as the relaying decision is based on the *current* conditions during a dissemination. However, it falls short when density increases, leading to broadcast storms [12]. Overlay-based broadcast is, in contrast, adapted to dense and more stable environments, where a small number of well-selected relays are pre-selected for broadcast, minimizing the risk of collisions.

Human mobility models [5], [6], [26] indicates that a hybrid approach could benefit from the resiliency and robustness of controlled flooding in general, while leveraging overlays where and when they are beneficial, that is, around the *points of interests* (POI) identified at large gatherings [8], [9].

The core contribution of this paper is the notion of *emergent overlays* as illustrated by Figure 5. Nodes in a MANET with emergent overlays use controlled flooding by default, configured for a good performance/reliability tradeoff in sparse to moderately dense areas. Nodes in denser and more stable areas *autonomously* decide to switch from controlled flooding to the setup and use of an overlay. Realizing emergent overlays requires mechanisms for interoperability, observation, and coordinated adaptation.

Interoperability. Broadcast messages may be emitted by any node of the network, from a zone using controlled flooding or from one using an overlay. It is necessary that the broadcast reaches all nodes regardless of the protocols they use. This is a problem of protocol interoperability. Overlay-based protocols maintain efficient dissemination structure for nodes *within their boundaries*. They do not select relays for guaranteeing connectivity to nodes that do not participate in the protocol. In §IV, we develop interoperability mechanisms that properly connect and bound regions using protocols of different types, regardless of the size and shape of these regions.

Observation. Nodes must be able to evaluate local density and mobility. Actively probing this information would defeat the purpose of reducing costs. Instead, we rely on passive observation of information collected from the regular operation of the BMW MAC protocol. We describe the nature and accuracy of the resulting *observables* in the first part of §V.

Adaptation. An emergent overlay is controlled by an adaptation *policy* to decide when to *emerge* and when to *dissolve* an overlay. Adaptation must be decentralized, and we only allow information piggybacked over existing messages to enable coordinated decisions. The goal of the adaptation is to form connected regions where nodes collectively switch to the use of an overlay. We discuss adaptation in the second part of §V.

IV. INTEROPERABILITY

Although nodes can use different protocols, broadcast is required to reach all nodes. Emergent overlays construct dissemination structures that cover their boundaries, i.e. nodes running MPR only receive `hello` messages from other nodes participating to the same protocol, and are not aware of nodes using CF outside of the emerged overlay boundary. Similarly, nodes running CF and located at the horizon of an emerged overlay may not relay a message if their counter goes down to 0 before the timer expires. This may hinder propagation of this message inside the emerged overlay.

We present in Algorithm 2 the interoperability mechanisms that allow messages to flow from regions using CF to regions using MPR, and *vice-versa*. The original calls of CF and MPR for receiving a message m are named "handle(m)" and encapsulated in novel interoperability-enabling functions named "receive(m)". Lines 1–8 are common, while lines 9–18 and 19–28 apply to nodes running CF or MPR, respectively.

The core of the interoperability protocol is the autonomous detection of nodes that are at the *border* of a zone using one protocol. When a node detects that it should act as a border node, it sets a boolean b for a (extendable) duration of two MPR `hello` message local broadcast periods (lines 4–8). We present how interoperability is enabled between CF and MPR nodes in each direction.

From CF to MPR. Nodes running CF detect that they are border nodes by reacting to `hello` messages sent by nearby MPR nodes. A CF border node will systematically act as a relay when receiving a new message from another node (running either protocol). This is enabled by setting an infinite value for the CF counter, which will therefore be positive when the CF timer expires (line 18).

Selecting all CF nodes in range of a MPR zone as border nodes, and thus as relays, is not desirable. This would lead to massive amounts of redundant relaying and collisions in the frontier regions of an emergent overlay although only a few exchanges across this frontier are actually sufficient to enable the overlay-based broadcast to take over. To reduce this number, CF nodes set a random timer t_{n_j} when they receive a `hello` message from a MPR node n_j (lines 9–10). If they receive a `border` message from another CF node before the

Algorithm 2 Interoperability mechanisms on node n_i

```
1: constants
2:    Δ: time          ▷ period of sending of hello messages for MPR
3: variables
4:    b: boolean, init false    ▷ should the node act as a border node?
5:    t: time, init ⊥              ▷ only meaningful if b is true
6: procedure iAmBorder()
7:    if b ∧ (currentTime - t) > 2 × Δ then b ← false
8:    return b
```

Code for node running CF

```
9: procedure MAC.receive(MPR hello from n_j)
10:    t_{n_j} ← new timer with duration ∈ [10 : 100] ms
11: procedure CF.receive(border(n_j) from n_k)
12:    cancel t_{n_j} (if it exists)
13: when timer t_{n_j} expires
14:    b ← true ; t ← currentTime     ▷ now a border node for 2 × Δ s
15:    MAC.localBroadcast(border(n_j))
16: procedure CF.receive(message m) for the first time
17:    CF.handle(m)                    ▷ regular handling of the message
18:    if iAmBorder() then counter_m ← ∞    ▷ forces being a relay
```

Code for node running MPR

```
19: procedure MAC.receive(border(n_j) from n_k)
20:    b ← true ; t ← currentTime     ▷ now a border node for 2 × Δs
21: procedure MPR.receive(message m) for the first time
22:    if m was sent using MPR then
23:        if n_i ∉ m.relays ∧ iAmBorder() then
24:            CF.handle(m)             ▷ if leaf and border node, use CF
25:        else
26:            if n_i ∈ m.relays then b ← false    ▷ already a relay
27:            MPR.handle(m)            ▷ regular handling of the message
28:    else MPR.broadcast(m)           ▷ from CF, initiate MPR broadcast
```

timer expires, they cancel this timer (lines 11-12). Otherwise, they tag themselves to act as a border node for the next $2 \times \Delta$ seconds, and send their own `border` frame to announce their existence to n_j (lines 13–15).

This mechanism allows a sufficient number of CF nodes in range of MPR nodes to be aware of the need to act as interoperability relays, while avoiding broadcast storms on the border regions.

From MPR to CF. Nodes running MPR must be aware of their border position. This happens through the reception of `border` messages from in-range CF nodes (lines 15 & 19–20). Border MPR nodes may be relays in the overlay or *leaf* nodes. Upon reception of a message using MPR (lines 22-27), the node first checks if it is a MPR leaf node but also a border node (line 23). This check must be performed because the overlay is implicit. A MPR leaf node that is at the border may have to forward the message in order to reach CF nodes across the overlay boundary. In contrast to CF nodes, MPR relay nodes do not pro-actively reduce their numbers to avoid broadcast storms. Instead, they re-use the controlled flooding message handling (using timers and counters) which may relay the message between the two zones (line 24). A MPR relay does not act as a border node as it will already forward the message during the regular MPR message handling procedure (line 26). Finally, when a MPR node receives an unseen message m from a CF node, it acts as a source and forwards m to its relay one-hop neighbors (line 28).

V. Autonomous adaptation

We now present how overlays emerge (and *dissolve*) autonomously in regions of a MANET where all nodes initially use CF. We start by presenting how nodes leverage the MAC protocol to collect observables about their environment. An adaptation *policy* triggers the creation of overlays based on thresholds over observables. The decision is collectively enforced for the necessary lifetime of the overlay, by piggybacking control information over existing message broadcasts.

A. Observables: density and stability

The emergence of an overlay is driven by the observed density and mobility. Those could be measured with periodic probing messages, as used by Tseng et al. [22] and Viswanath and Obraczka [17] to adapt the parameters of controlled flooding. The use of extra messages go against our objective of reducing costs and risk of collisions. Instead, we observe that the necessary information can be collected passively from the BMW MAC protocol. CTS requests are sent for every local broadcast, and nodes in range reply with RTS frames. This enables the collection of one-hop neighbors identities. We collect the set of repliers for each RTS/CTS exchange, and keep a limited-time history of these neighborhoods (last 10 seconds in our implementation).

Density for a node n_i is measured as its degree, or number of observed neighbors. The measurements in the time window for n_i form a sequence $d_{-x}, d_{-w}, d_{-v} \ldots, d_0$, where $x \leq 10$ is the time of the latest measurement. Using only d_0 may lead to sudden fluctuations that do not reflect medium-term increases in density around the node. We compute instead the density average over periods of time, e.g. $\bar{d}_{-1}, \bar{d}_{-2}, \bar{d}_{-5}, \bar{d}_{-10}$. As measurements are not taken with a fixed frequency, we weight each measurement using its validity window (e.g., $d_{-w} - d_{-x}$).

Mobility is measured indirectly through the *stability* of n_i's neighborhood, measured as the average time unique neighbors are observed. Again, we compute stability over varying periods of time, resulting in $\bar{s}_{-1}, \bar{s}_{-2}, \bar{s}_{-5}, \bar{s}_{-10}$.

B. Protocol switching thresholds

Nodes decide to emerge an overlay based on both $\bar{s}_{-\delta}$ and $\bar{d}_{-\delta}$ for a time window δ. In order to avoid oscillations where nodes would switch back and forth between using CF and MPR, we use a hysteresis, with two pairs of thresholds: $(s_{\min}, d_{\text{emerge}})$ and $(s_{\min}, d_{\text{keep}})$, where $d_{\text{keep}} < d_{\text{emerge}}$. The criterium $s_{\min}$ is a minimal stability required in all cases. The two density criteria d_{emerge} and d_{keep} indicate when to respectively emerge and maintain an overlay. The values of these criteria depend on the nature of the MANET, the wireless range of its nodes, and deployment factor. We define appropriate values for our test scenario in our evaluation.

C. Protocol switching policies

A simple policy would be to let nodes decide independently when to switch from CF to MPR, when the locally observed value for $(\bar{s}_{-\delta}, \bar{d}_{-\delta})$ exceeds both thresholds $(s_{\min}, d_{\text{emerge}})$. As long as these values remain over thresholds $(s_{\min}, d_{\text{keep}})$,

Algorithm 3 Coordinated adaptation on node n_i

1: **constants**
2:   Δ: time   ▷ period of sending of `hello` messages for MPR
3:   $s_{\min}, d_{\text{emerge}}, d_{\text{keep}}$ ▷ policy thresholds
4: **variables**
5:   $\bar{s}_{-\delta}, \bar{d}_{-\delta}$ ▷ observables maintained by MAC protocol
6:   $P \in \{$ CF, MPR $\}$ ▷ currently running protocol
7:   t_{MPR} ▷ when running MPR, dissolution timer

8: **procedure** emerge(): **return** $(\bar{s}_{-\delta} \geq s_{\min} \wedge \bar{d}_{-\delta} \geq d_{\text{emerge}})$

9: **procedure** keep(): **return** $(\bar{s}_{-\delta} \geq s_{\min} \wedge \bar{d}_{-\delta} \geq d_{\text{keep}})$

10: **upon** processing of message m for the first time
11:   **if** emerge() **then**   ▷ act as a proposer
12:     m.em $\leftarrow$ true, useMPR()
13:   **else if** m.em **then**   ▷ react to proposal:
14:     **if** keep() **then** useMPR()   ▷ start or renew MPR use, or
15:     **else** m.em $\leftarrow$ false   ▷ do not propagate proposal further
16: **procedure** useMPR()
17:   **if** $P =$ CF **then**
18:     MPR.init()   ▷ initialize `hello` messages
19:     wait($2 \times \Delta$), then $P \leftarrow$ MPR   ▷ stabilize then switch
20:   $t_{\text{MPR}} \leftarrow$ currentTime() $+10 \times \Delta$ ▷ set dissolution timer
21: **when** timer t_{MPR} expires or if no MPR hello received in 2Δ
22:   MPR.terminate()
23:   $P \leftarrow$ CF   ▷ dissolve overlay and return to using CF

Fig. 6. Initial position of 600 nodes in the simulation field.

a node would continue using MPR, but switch back to CF otherwise. While this *single-node* policy will lead homogeneously dense and stable regions to eventually switch to MPR, it has the risk of creating small MPR regions surrounded by CF zones, when nodes density is distributed between d_{keep} and d_{emerge}. A *coordinated* mechanism is given in Algorithm 3, that addresses these shortcomings. First, it allows entire *regions* of the MANET to switch simultaneously to use MPR, thereby avoiding single isolated nodes using an overlay. Second, it allows the created overlays to gracefully *dissolve* when they are no longer necessary, without unnecessarily keeping isolated residual regions.

The policy leverages nodes in very dense regions to lead the emergence and maintenance of overlays. The entry point to the policy algorithm (line 10) is weaved with the processing of a message m by either of the protocols. Nodes whose local thresholds are above $(s_{\min}, d_{\text{emerge}})$ will *invite* other nodes contacted by the broadcast of m to switch or keep using MPR. The invitation is a single bit em (for *emerge!*) piggybacked in the header of m. We note that multiple nodes in a dense regions may set em, and possibly switch themselves to MPR (line 12). Other nodes receiving a message with the em bit set will check whether their local observables are above the $(s_{\min}, d_{\text{keep}})$ thresholds (line 14). If it is the case, a node will start or continue using MPR. The value of em is unmodified and the message is further broadcast. Nodes that do not pass any of the thresholds unset em, marking the boundary of the emerging overlay's region (line 15).

The switch from CF to MPR is not immediate. After starting to send `hello` frames (line 18), nodes wait for a period of $2 \times \Delta$ before switching to MPR (line 19). The lifetime of an MPR overlay is governed by the timer t_{MPR} at each node, which is always set to $10 \times \Delta$ when a node switches to, or continues using, MPR (line 20). If no further invitation to
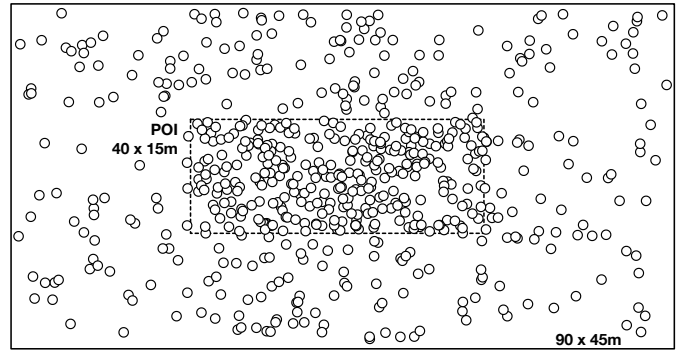
use MPR has been received when the timer expires, the node switches back to CF. This can happen either as nodes are in a moderately dense area where no node has consistently seen a density justifying sending an invitation, or if an increase in mobility jeopardizes the interest of using an overlay. An exception is made for nodes running MPR who do not receive a `hello` frame for more than a period of 2Δ. This accounts for mobile nodes who have moved out of a dense zone, and are no longer in range of any MPR node.

VI. EVALUATION

We evaluate the performance and reliability of broadcast using full-stack simulations in Omnet++ v5.4 [23] and its plug-in INET v4.1. We compare emergent overlays with baseline protocols and with two approaches to adaptive broadcast from the literature [17], [22].

Our evaluation aims at answering the following research questions: (i) Do emergent overlays allow improving the cost of broadcast in zones where CF is inefficient? (ii) Is autonomous adaptation effective in emerging overlays without centralized control? (iii) How do emergent overlays compare to previous approaches of adaptive broadcasting?

A. Experimental setup

Simulation settings. We simulate a full MANET network stack, from the wireless medium to application-level protocols. The physical and data-link layers follow IEEE 802.11 specifications, with the additional use of BMW for MAC-level local broadcast. Every node uses a single transceiver with a fixed transmission range of 10 m, representative of the range of bluetooth and similar radio devices. This transceiver uses the 2.4 GHz frequency band, with a maximum achievable data rate of 11 Mbps.

Use case and experiment model. We consider a set of 600 mobile nodes, the placement and mobility of which are inspired by real-life measurements of crowds at large events [8], [9]. Figure 6 portrays the starting node positions in the simulated setup. Nodes are deployed over a 90×45 m field. A 40×15 m area in the middle of this field represents a POI where a larger number of nodes tend to aggregate.

At the beginning of the simulation, 300 nodes are randomly placed inside the POI and 300 others outside. During the simulation, nodes are mobile. To simulate human mobility, we use the *truncated levy walks* [6] model, which combines frequent short walks with occasional rides to distant locations. The velocity of every walk is chosen randomly from an interval: [0:0.5] m/s at the POI and [1.5:2] m/s outside it. These values are representative of mobility in large gatherings. Nodes move freely within the entire communication area. We verify upon generation of a mobility trace that at any point in time, the graph formed by all nodes and direct neighborhood relations is connected (therefore, the potential for complete coverage is theoretically 100%). We discard the infrequent random seeds that do not ensure this connectedness.

Each experiment lasts for 10 minutes, with one new broadcast every two seconds, starting from a randomly-selected source node. Messages are always 280 Bytes in size, similar to the length of a typical message on social media. We report the aggregates of 5 experiment runs with different random seeds.

Comparison to previous approaches. We consider our baseline protocols (counter-based CF and MPR) used in isolation, and combined for emergent overlays. In addition, we compare with two state-of-the-art adaptive approaches. We implemented these two protocols in the same simulation environment, and both use the BMW MAC protocol for reliable local broadcast.

Our first comparison point is *adaptive counter-based controlled flooding*, or ACF [22]. It is an adaptive variant of counter-based CF (§II). It uses control messages to collect density information around each node. We rely instead on the observable from the MAC protocol. The setting of timers is unchanged, but the initial value of counters is a function of density. We use the adaptation function suggested by the authors [22, Figure 6]. Initial counter values vary from 2 to 5 for different densities (i.e., from 2 to 5 for density values 1 to 4, and then slowly down to a value of 2 for 12+ neighbors).

Our second comparison point is *S-H flooding* [17]. It uses two types of controlled flooding protocols, *scoped flooding* for low-mobility nodes and *hyper flooding* for high-mobility nodes. The decision of which protocol to use is made independently by each node based on its measured *velocity*, therefore requiring accelerometers. We use 1 m/s (a typical walking speed) as the switching threshold. S-H flooding needs the list of one-hop neighbors, which is piggybacked on all forwarded messages. We do not use explicit control messages but obtain these lists passively from the MAC protocol, as we do for ACF. A node using scoped flooding compares the list of neighbors of the forwarder to its own and only retransmits if less than 85% of the two sets overlap. Hyper flooding is a variant of timer-based CF, without counters, where a timed retransmission at n_i is only cancelled when the duplicate is received by a node present in n_i's currently-known neighbor set.

Metrics. We use the following metrics:

- Reliability of the broadcast is measured by its **reach** (in %), the average proportion of nodes that receive each message;

Protocol	Region	Reach (%)	LBC (%)	Collisions
CF$_{c=3}$	All	98.90	78.51	8890.00
	Out of POI	100.00	94.40	3070.10
	At POI	97.30	35.09	5819.80
CF$_{c=4}$	All	99.91	74.63	12403.36
	Out of POI	100.00	90.05	3962.85
	At POI	98.67	32.69	8440.51
MPR$_{\Delta=20s}$	All	99.80	94.70	599.80
	Out of POI	99.60	98.00	231.50
	At POI	100.00	81.50	368.20
MPR$_{\Delta=30s}$	All	99.30	95.00	621.60
	Out of POI	98.70	98.40	245.20
	At POI	99.90	79.50	376.30

TABLE I
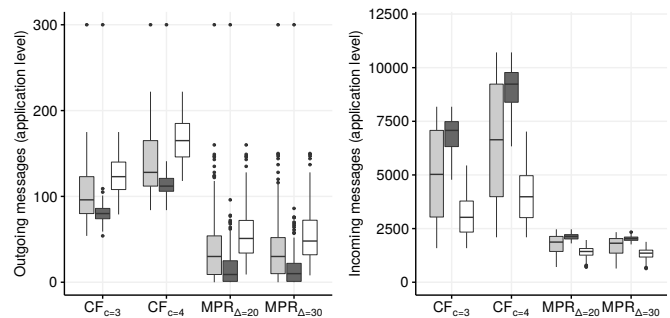PERFORMANCE AND RELIABILITY OF BASELINE PROTOCOLS.

Fig. 7. Performance of single-protocol deployments, for all of the communication area (light grey), at the POI (dark grey) and out of the POI (white).

- Reliability of local broadcast is measured at the MAC level by *local broadcast coverage* or **LBC** (in %), the average proportion of neighbors that successfully receive frames despite collisions. A low LBC indicates that nodes were not able to successfully disseminate frames to their neighbors despite using 3 MAC-level retransmissions. It is a clear sign of a strongly locally-contented wireless medium, which directly impacts on the reliability and cost of the broadcast;
- Overall contention on the wireless medium is measured as the number of observed **collisions**;
- Bandwidth cost for nodes is measured as the number of **sent and received messages**.

B. Performance of base protocols

We start by evaluating the performance and reliability of the base protocols. Table I reports reliability metrics for two variants of CF with initial counter values of 3 and 4, and two variants of MPR with period Δ of 20 and 30 seconds. Figure 7 details the network traffic at the application level (i.e. forwarded messages, without control traffic). We distinguish between outgoing messages (i.e. forwarded messages) and incoming messages (i.e. received messages). The difference between the former and the latter stands for duplicates. We present all metrics for the entire system, and the breakdown outside and inside the central point of interest.
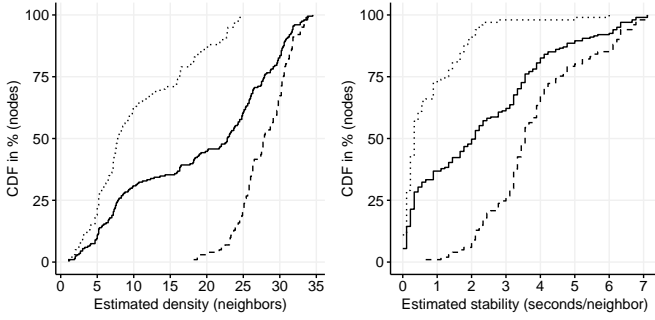
Fig. 8. Distributions of density $\bar{d}_5$ (left) and stability $\bar{s}_5$ (right) in the 3 regions: for all of the communication area (solid line), at the POI (dashed line) and out of the POI (dotted line).

We can observe that CF with an initial counter value of 3 achieves a good reach in the entire system, but only through a high number of duplicates. This results in low LBC and a high number of collisions. This is particularly true at the POI where the LBC plummets down to 35.09%, and where the number of collisions is significantly higher than outside of the POI for the same number of nodes. The use of a higher counter value (c = 4) improves reach but also results in a major inefficiency particularly at the POI.

MPR is able to attain total reach (100% with $\Delta = 20s$) for nodes at the POI, with a significantly better LBC and less collisions than CF. The reliability out of the POI is good but not as high as with CF. We note that MPR with control messages periods Δ lower than 20 seconds yields a significant number of collisions and, thus, suffers a loss in reliability.

The evaluation of individual protocols confirms that CF is inefficient in dense zones, where MPR is a better choice. We choose to use MPR with $\Delta = 30s$ for emergent overlays, and CF with c = 3 otherwise.

C. Quality of observables

We evaluate the quality of observables collected at the level of the MAC protocol, as defined in §V-A. Figure 8 presents the distribution of observed density ($\bar{d}_5$) and observed stability ($\bar{s}_5$) for an observation period of 5 seconds.[2] We present the distribution for both the entire system and regions at and outside the POI.

We can observe that the distribution of density outside and at the POI differ, with the former ranging from as low as 1 and up to 25 neighbors, while the latter ranges from 18 to 34. These estimations are in line with a baseline of exact densities that we compute offline from the mobility trace. We only observe a slight deviation to lower density estimates for nodes at the POI, that we explain as a result of collisions leading to missed CTS messages by the MAC protocol. Stability ranges from 0 to 6 and 1 to 7 seconds per neighbor on average, outside and at the POI, respectively, in coherence with our mobility model.

[2]We do not present distributions for other aggregation periods due to space limitations. Our experiments confirm that short periods (e.g. 1 second) yield similar distributions, but measurements for individual nodes are subject to higher noise. Longer periods (e.g. 10 seconds) do not bring significant benefits.
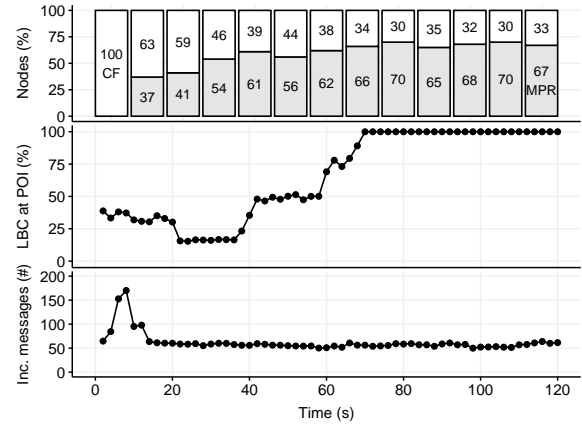

Fig. 9. Timeline of an overlay emergence, with the number of nodes using the two protocols (first plot), the average LBC for nodes inside the POI (second plot) and the average number of incoming messages (third plot).

D. Efficiency of emerging overlays

We evaluate the effectiveness of our adaptation protocol in emerging an overlay at the POI. The simulation starts with all nodes using CF, with new messages being sent from the beginning, and every 2 seconds.

We use the following parameters for our adaptation (§V): s_{min} is set to 2 seconds/node. d_{emerge} and d_{keep} are set to 20 and 10 observed neighbors respectively. These values were determined from the distribution of observables discussed in the previous subsection.

Figure 9 presents, over time and from top to bottom, the evolution of the number of nodes using either protocol, the LBC at the POI and the average number of incoming messages over the entire network. Reach is not shown but remains close to 100% for the entire duration of the experiment.

We observe a gradual increase in the number of nodes switching to MPR in the first plot. This number stabilizes at around two-thirds of the network for the rest of the simulation. The LBC at the POI is very low at first, as a result of the high number of collisions in the dense region using CF. This results in a higher number of retransmissions and therefore a higher cost to obtain a high reach. At time t=60s, we can observe a sharp increase in the LBC for nodes at the POI. This corresponds to the end of the transition period (line 19 of Algorithm 3), of duration $2 \times \Delta = 2 \times 30s$.

This experiment shows the capacity of nodes to trigger the adaptation. Actually, we can observe that more nodes that those strictly at the POI itself switch to MPR. This is a result of measurements for nodes out of the POI but close to it, who are neighbors with the denser area and pass the density threshold d_{keep}.

E. Performance of emerged overlays

We finally evaluate the performance of the emerged overlay, and compare it with the result of the two other adaptive approaches, ACF and S-H flooding.

Table II presents the reliability metrics, and Figure 10 depicts the number of forwarded and incoming messages. As for previous experiments we present the metrics for the entire system and a breakdown over the two zones.

Protocol	Region	Reach (%)	LBC (%)	Collisions
Emerg. Ovl.	All	99.99	87.35	3552.83
$CF_{c=3}$	Out of POI	99.99	73.13	1871.77
$MPR_{\Delta=30s}$	At POI	100.00	93.40	1681.05
	All	98.90	82.25	5058.26
ACF	Out of POI	100.00	99.83	1515.69
	At POI	98.90	50.27	3542.56
S-H	All	92.67	62.07	17501.30
Flooding	Out of POI	90.99	100.00	6609.07
	At POI	96.34	25.23	10892.22

TABLE II
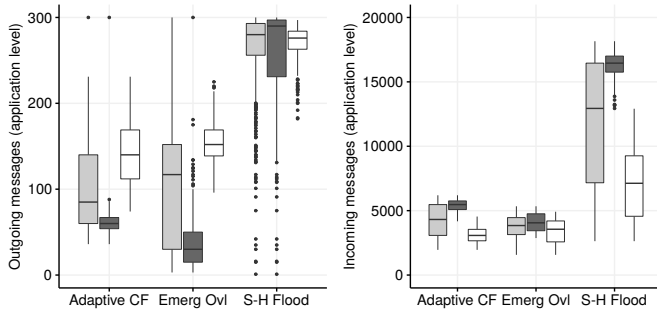PERFORMANCE AND RELIABILITY OF ADAPTIVE APPROACHES.

Fig. 10. Performance of adaptive approaches, for all of the communication area (light grey), at the POI (dark grey) and out of the POI (white).

We first observe that our approach reaches an excellent reach of 99.99% out of the POI, and 100% at the POI where the overlay is used. The LBC at the POI is over 90%, and over 70% outside of it. The number of collisions is below that of CF used alone, for nodes out of the POI. This is explained by the presence of border nodes who disseminate to a significant portion of these nodes, thereby increasing the coverage of the MPR overlay and reducing the number of necessary CF forwarding. This number of collisions for nodes at the POI is however higher than for the same nodes when the entire network uses only MPR. This is a result of the additional control messages and traffic generated by border nodes, that interfere with the overlay dissemination itself. The emergent overlay realizes a compromise between the performance of MPR at the POI and the reliability of CF in the rest of the system. Its cost is slightly higher than MPR alone due to the use of interoperability protocols and border nodes, but it remains significantly less than CF alone.

ACF offers good performance out of the POI with 100% reach and a high LBC. However, the use of controlled flooding at the POI, despite the adaptation of its parameters, yields an inefficient broadcast for nodes at the POI, with a LBC reaching only 50.27%, and with a number of collisions that is more than twice that of the emergent overlay. ACF indeed uses a similar number of incoming and outgoing messages outside of the POI but requires almost twice as much of the former, and results in about 20% more of the latter. The performance of S-H flooding, on the other hand, is deceptive. Perhaps surprisingly, the reach is better at the POI than outside of it, despite the protocol being a controlled flooding approach. LBC out of the POI is 100% but the reach is only 90.99%. The reach is higher

at the POI, but at the cost of a very high number of messages, and therefore duplicates and collisions. This results in a LBC that only reaches 25.23%.

Our evaluation and the comparison to other approaches show that emergent overlays allow combining the interest of controlled flooding in the sparse and mobile zone, out of the POI, while benefitting from the performance and reliability of the overlay-based broadcast at the POI. The overall reliability is high and the cost remains reasonable compared to the use of a single protocol and to other broadcast approaches adapting the parameters of controlled flooding.

VII. RELATED WORK

A survey by Ruiz and Bouvry [13] provides a comprehensive taxonomy of controlled flooding and overlay-based protocols, also covered in surveys by Reina *et al.* [27] and Santi [28]. Our previous work [14] and the work of Williams and Camp [15] present comparative evaluations of broadcast protocols from different classes. Hess *et al.* [5] survey human mobility models for mobile networking. Levy walk are a *general* model [6]. Other proposals allow modeling mobility in specific scenarios, such as disaster areas [26].

The *broadcast storm* problem associated with flooding is identified by Ni et al. [12], who also propose *controlled flooding* using counters and timers as we use in this paper. The collection of topology information [29] or GPS locations [25] allows the optimizing of controlled flooding. However, they are costlier and more sensitive to high mobility.

Overlay-based broadcast is based on the construction of connected dominating sets (CDS) [19], with the goal of minimizing the number of necessary relay nodes. MPR [20] or TSS [18] create CDS by guaranteeing the coverage of two-hop neighbors by at least one relay node. Ahn and Lee [30] improve MPR to guarantee that each 2-hop neighbor is covered by m such relays. Other approaches [31], [32] consider variable-range local broadcast to optimize the CDS.

Determining the optimal configuration for a MANET broadcast algorithm is difficult. Kokuti and Vilmos [33] propose adaptation techniques to dynamically decide on parameters in controlled flooding. Tseng, Ni and Shih [22] similarly propose the ACF adaptation mechanism that we used in our evaluation (§VI). Controlled flooding adaptation can use measures of node velocity, as for S-H flooding [17] detailed in §VI, or measures of observed contention and collisions [22]. Liarokapis and Shahrabi [34] propose techniques to switch from simple to controlled flooding, and adapt switching threshold to locally measured density. Artimy [35] proposes density estimation and dynamic range adaptation for vehicular networks, distinguishing between free-flow and congested phases.

The term *emergent overlay* has been previously used in another context, wired unstructured P2P network [36], to denote the *on demand* construction of distributed hash tables [37] following a centralized decision. Bellavista *et al.* [2] combine a static sensor network using an overlay and a MANET with controlled flooding for data collection for an IoT scenario, but the use of either protocol is determined in advance. We are

not aware of any previous approach combining in a single MANET controlled flooding and overlays.

VIII. Conclusion

We presented the design and implementation of *emergent overlays* for efficient broadcast in heterogeneous MANETs. Emergent overlays allow interoperability between regions of the MANET using controlled flooding and regions using a MPR overlay. Our decentralized adaptation mechanism allows such regions to autonomously coordinate on the emergence of overlays, based on density and mobility. Our evaluation shows that emergent overlays offer good delivery guarantees with a lower cost than state-of-the-art adaptation techniques for controlled flooding. They reduce the number of messages and collisions and therefore the cost of broadcast.

This works open interesting perspectives that we consider for our future work. A first one is the use of emergent overlays for *routing*, e.g. based on the discovering at runtime of commonly-used paths between sources and destinations. A second one is adaptation strategies to automatically derive thresholds used for emerging overlays, based on observations of broadcast reliability and performance.

References

[1] J. Hoebeke, I. Moerman, B. Dhoedt, and P. Demeester, "An overview of mobile ad hoc networks: Applications and challenges," *Communications Network*, vol. 3, no. 3, 2004.

[2] P. Bellavista, G. Cardone, A. Corradi, and L. Foschini, "Convergence of MANET and WSN in IoT urban scenarios," *IEEE Sensors Journal*, vol. 13, no. 10, 2013.

[3] P. Lieser, F. Alvarez, P. Gardner-Stephen, M. Hollick, and D. Boehnstedt, "Architecture for responsive emergency communications networks," in *IEEE Global Humanitarian Technology Conference*, 2017.

[4] H. Verma and N. Chauhan, "MANET based emergency communication system for natural disasters," in *IEEE Intl. Conference on Computing, Communication & Automation*, 2015.

[5] A. Hess, K. A. Hummel, W. N. Gansterer, and G. Haring, "Data-driven human mobility modeling: a survey and engineering guidance for mobile networking," *ACM Comput. Surv.*, vol. 48, no. 3, 2016.

[6] I. Rhee, M. Shin, S. Hong, K. Lee, S. J. Kim, and S. Chong, "On the levy-walk nature of human mobility," *IEEE/ACM Trans. on Networking*, vol. 19, no. 3, 2011.

[7] I. Stepanov, P. J. Marron, and K. Rothermel, "Mobility modeling of outdoor scenarios for MANETs," in *Symp. on Simulation*, 2005.

[8] F. M. Naini, O. Dousse, P. Thiran, and M. Vetterli, "Opportunistic sampling for joint population size and density estimation," *IEEE Trans. on Mobile Computing*, vol. 14, no. 12, 2015.

[9] M. Versichele, T. Neutens, M. Delafontaine, and N. Van de Weghe, "The use of bluetooth for analysing spatiotemporal dynamics of human movement at mass events: A case study of the Ghent festivities," *Applied Geography*, vol. 32, no. 2, 2012.

[10] A. Tsertou and D. I. Laurenson, "Revisiting the hidden terminal problem in a CSMA/CA wireless network," *IEEE Trans. on Mobile Computing*, vol. 7, no. 7, 2008.

[11] "Wireless lan medium access control (mac) and physical layer (phy) specifications," IEEE 802.11 Working Group, Tech. Rep., 1997.

[12] S.-Y. Ni, Y.-C. Tseng, Y.-S. Chen, and J.-P. Sheu, "The broadcast storm problem in a mobile ad hoc network," in *ACM/IEEE Conference on Mobile computing and networking*, ser. MobiCom, 1999.

[13] P. Ruiz and P. Bouvry, "Survey on broadcast algorithms for mobile ad hoc networks," *ACM Comput. Surv.*, vol. 48, no. 1, 2015.

[14] R. Carvajal Gómez, I. Gonzalez-Herrera, Y.-D. Bromberg, L. Réveillère, and E. Rivière, "Density and mobility-driven evaluation of broadcast algorithms for MANETs," in *IEEE Intl. Conference on Distributed Computing Systems*, ser. ICDCS, 2017.

[15] B. Williams and T. Camp, "Comparison of broadcasting techniques for mobile ad hoc networks," in *Symp. on Mobile Ad Hoc Networking & Computing*, ser. MobiHoc, 2002.

[16] K. Obraczka, G. Tsudik, and K. Viswanath, "Pushing the limits of multicast in ad hoc networks," in *IEEE Intl. Conference on Distributed Computing Systems*, ser. ICDCS, 2001.

[17] K. Viswanath and K. Obraczka, "An adaptive approach to group communications in multi hop ad hoc networks," in *IEEE Intl. Symp. on Computers and Communications*, ser. ISCC, 2002.

[18] M. Nikolov and Z. J. Haas, "Towards optimal broadcast in wireless networks," *IEEE Trans. on Mobile Computing*, vol. 14, no. 7, 2015.

[19] J. Wu, F. Dai, M. Gao, and I. Stojmenovic, "On calculating power-aware connected dominating sets for efficient routing in ad hoc wireless networks," *Journal of Communications and Networks*, vol. 4, no. 1, 2002.

[20] A. Qayyum, L. Viennot, and A. Laouiti, "Multipoint relaying for flooding broadcast messages in mobile wireless networks," in *IEEE annual Hawaii Intl. Conf. on System Sciences*, 2002.

[21] K. Tang and M. Gerla, "MAC reliable broadcast in ad hoc networks," in *Communications for Network-Centric Operations*, ser. MILCOM, 2001.

[22] Y.-C. Tseng, S.-Y. Ni, and E.-Y. Shih, "Adaptive approaches to relieving broadcast storms in a wireless multi-hop mobile ad hoc network," *IEEE Trans. on Computers*, no. 5, 2003.

[23] A. Varga and R. Hornig, "An overview of the OMNeT++ simulation environment," in *Conference on Simulation Tools and Techniques for Communications, Networks and Systems*, ser. SIMUtools, 2008.

[24] E. Leme, N. Ivaki, N. Laranjeiro, and R. Moraes, "Analyzing gossip protocols for reliable MANET applications," in *IEEE Intl. Conference on Edge Computing*, 2017.

[25] F. J. Ovalle-Martínez, A. Nayak, I. Stojmenović, J. Carle, and D. Simplot-Ryl, "Area based beaconless reliable broadcasting in sensor networks," in *Conf. on Algo. Aspects of Wireless Sensor Net.*, 2006.

[26] N. Aschenbruck, E. Gerhards-Padilla, and P. Martini, "Modeling mobility in disaster area scenarios," *Perf. Evaluation*, vol. 66, no. 12, 2009.

[27] D. Reina, S. Toral, P. Johnson, and F. Barrero, "A survey on probabilistic broadcast schemes for wireless ad hoc networks," *Ad Hoc Networks*, vol. 25, 2015.

[28] P. Santi, "Topology control in wireless ad hoc and sensor networks," *ACM Comput. Surv.*, vol. 37, no. 2, 2005.

[29] J. Cartigny and D. Simplot, "Border node retransmission based probabilistic broadcast protocols in ad-hoc networks," *Telecommun. Syst.*, vol. 22, no. 1-4, 2003.

[30] J. H. Ahn and T.-J. Lee, "Multipoint relay selection for robust broadcast in ad hoc networks," *Ad Hoc Networks*, vol. 17, 2014.

[31] W. Woon and K. L. Yeung, "Variable power broadcasting based on local information for source-dependent broadcasting protocols," in *Wireless Communications and Networking Conference*, ser. WCNC, 2011.

[32] S. Park and S.-M. Yoo, "An efficient reliable one-hop broadcast in mobile ad hoc networks," *Ad Hoc Networks*, vol. 11, no. 1, 2013.

[33] A. Kokuti and V. Simon, "Adaptive multihop broadcast protocols for ad hoc networks," in *Symp. on Communication Systems, Networks & Digital Signal Processing*, ser. CSNDSP, 2012.

[34] D. Liarokapis and A. Shahrabi, "A probability-based adaptive scheme for broadcasting in MANETs," in *Conference on Mobile Technology, Application & Systems*, 2009.

[35] M. Artimy, "Local density estimation and dynamic transmission-range assignment in vehicular ad hoc networks," *IEEE Trans. on Intelligent Transportation Systems*, vol. 8, no. 3, 2007.

[36] W. Galuba and K. Aberer, "Generic emergent overlays in arbitrary peer identifier spaces," in *Workshop on Self-Organizing Systems*, 2007.

[37] M. Jelasity, A. Montresor, and O. Babaoglu, "T-man: Gossip-based fast overlay topology construction," *Computer Networks*, vol. 53:13, 2009.