

Beyond Groups: Uncovering Dynamic Communities on the WhatsApp Network of Information Dissemination

Original

Beyond Groups: Uncovering Dynamic Communities on the WhatsApp Network of Information Dissemination / Nobre, G. P.; GOMES FERREIRA, CARLOS HENRIQUE; Almeida, J. M.. - STAMPA. - 12467:(2020), pp. 252-266. (Intervento presentato al convegno 12th International Conference on Social Informatics, SocInfo 2020 tenutosi a Pisa (Italy) nel October 6–9, 2020) [10.1007/978-3-030-60975-7_19].

Availability:

This version is available at: 11583/2932664 since: 2022-02-02T15:04:08Z

Publisher:

Springer Science and Business Media Deutschland GmbH

Published

DOI:10.1007/978-3-030-60975-7_19

Terms of use:

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright

Springer postprint/Author's Accepted Manuscript

This version of the article has been accepted for publication, after peer review (when applicable) and is subject to Springer Nature's AM terms of use, but is not the Version of Record and does not reflect post-acceptance improvements, or any corrections. The Version of Record is available online at: http://dx.doi.org/10.1007/978-3-030-60975-7_19

(Article begins on next page)

Beyond Groups: Uncovering Dynamic Communities on the WhatsApp Network of Information Dissemination

Gabriel P. Nobre¹, Carlos H. G. Ferreira^{1,2}, and Jussara M. Almeida¹

¹ Universidade Federal de Minas Gerais, Brazil

² Universidade Federal de Ouro Preto, Brazil

Abstract. In this paper, we investigate the network of information dissemination that emerges from group communication in the increasingly popular WhatsApp platform. We aim to reveal properties of the underlying structure that facilitates content spread in the system, despite limitations the application imposes in group membership. Our analyses reveal a number of strongly connected user communities that cross the boundaries of groups, suggesting that such boundaries offer little constraint to information spread. We also show that, despite frequent changes in community membership, there are consistent co-sharing activities among some users which, even while holding broad content diversity, lead to high coverage of the network in terms of groups and individual users.

Keywords: WhatsApp · Community Detection · Information Spread

1 Introduction

Social media applications are known as tools to connect people, enhance human interactions, ultimately contributing to information diffusion on the Internet, a widely explored subject of study [22, 25, 33]. WhatsApp is one such application with great popularity in many countries such as India, Brazil and Germany [41]. Indeed, the application, which has recently surpassed the mark of 2 billion monthly users worldwide [40], has been shown to play an important role as a vehicle for information dissemination and social mobilization [34].

WhatsApp allows for one-to-one and group conversations, both end-to-end encrypted. WhatsApp groups are structured as private chat rooms, generally under a certain topic and limited to only 256 users who can participate in multiple discussions at the same time. However, a group manager may choose to share an invitation link to join the group on websites and social networks, which effectively makes the group publicly accessible, as anyone with access to the link can join in and become a member. We note that groups are dynamic spaces of conversations, with users joining and leaving over time, at their will. Also, there is no limit in the number of groups a user can participate in at any time.

A number of recent studies have analyzed group conversations in WhatsApp [8, 13, 28, 33, 34], often discussing content and temporal properties of information dissemination within groups. A few have hinted at the potential for

information virality by taking a bird’s eye view of the network that emerges from information exchange in different groups [8, 34]. However, an investigation of the properties of this network, which may reveal an underlying structure that facilitates information dissemination at large, is still lacking.

In this paper, we delve deeper into the network of information dissemination in WhatsApp. Specifically, we use a dataset consisting of messages posted in over 150 publicly accessible WhatsApp groups in Brazil during a 6 week period, encompassing the 2018 general elections [34]. We build a sequence of media-centric networks so as to capture the relationships established among users who shared the same piece of content in one or more groups during pre-defined time windows. We then analyze properties of these networks and how they evolve over time. In particular, we are interested in investigating to which extent users in different groups, intentionally or not, build “communities” of content spread, which consistently help speeding up information dissemination within the system.

Specifically, we tackle the following two research questions (RQs):

- **RQ1:** To which extent cross-group user communities emerge from analyzing the WhatsApp media-centric networks?
- **RQ2:** What are key properties of these communities and how they evolve over time?

One challenge when analyzing the media-centric networks is that, by definition, they model interactions established among multiple users, that is, many-to-many interactions, as opposed to traditionally studied binary relationships. As recently argued [4, 10], such many-to-many interactions may lead to the formation of networks with multiple spurious and random edges. These spurious edges may ultimately hide the real underlying structure (often called the network backbone) representing the phenomenon under study, in our case, information spread across the monitored groups. Thus, as a first step to our study, we applied a technique [37] to extract the backbone of each media-centric network. We then used the Louvain algorithm [6] to extract communities in the backbone of each network, and analyzed their properties and how they evolve over time.

Our study revealed a number of strongly connected user communities that extrapolate the boundaries of groups, suggesting that such boundaries offer little constraint (if any) to information spread. Indeed, it is often the case that the same community covers a large number of different groups, and users with higher centrality in these communities are the ones with most impact on the community’s group coverage and on the content uniqueness spread through the network. We also observed that around 30% of the users persist in the network backbone over time, whereas those with highest activity tend to remain even in the same community. These results suggest that, though WhatsApp groups are limited to small sizes, the platform effectively has an underlying network that greatly facilitates information spread, revealing consistent co-sharing activities among the users which, even while holding broad content diversity, lead to high user and group coverage. While corroborating arguments in [34], our study offers a much deeper analysis, revealing properties that help understand the information dissemination phenomenon in a very popular communication platform.

2 Related Work

There is a rich literature on online information spread and its related phenomena involving the modeling of online user interactions [2, 16], the extraction of user communities [6, 17], the identification of important users [15] and the analysis of the information diffusion process [8, 25, 34]. In the context of community extraction, there is a plethora of techniques [12, 32] exploiting from statistical inference [1, 19] to modularity optimization [44] and dynamic clustering [12, 35]. Others have offered analyses of communities’ structural properties [17, 24] and temporal behavior [2]. In contrast, the authors of [16] described a system that allows the detection of influential users in dynamic social networks. Other authors have examined multi-model networks by proposing a framework for detecting community evolution over time [36, 39].

WhatsApp has emerged as a global tool for communication [40] and has driven many recent researches. In [8], the authors analyzed information spread within groups from the perspective of user attention on different topics, while in [7] the authors focused on partisan activities, aiming to distinguish left-wing and right-wing groups. The work in [33] describes a system for collecting shared messages in publicly accessible WhatsApp groups. The collected content was later used in analyses of content properties and temporal dynamics of the dissemination of images [34], textual messages [33] and audio content [28].

Despite WhatsApp group communication being restricted to small groups (up to 256 users), there has been evidence of its use for information dissemination at large [18, 27]. The only prior work that hinted at possible reasons for that was [34], which presented the network structure of the monitored groups, briefly analyzing some of its basic structural properties. Also, the authors of [13] showed that recent limitations on messaging forwarding [42] are not effective to block viral content spread across the network (despite contributing to delay it).

Our goal here is to delve deeper into information spread in WhatsApp groups, unveiling the underlying media sharing network that emerges from communication in those groups and investigating the extent to which the formation of user communities that cross group boundaries occur, favoring information spread. We thus complement prior analyses of WhatsApp, focusing on the network structure and how its properties relate to information dissemination in the platform.

3 Methodology

As stated, this work aims to investigate the formation of user communities that may favor information dissemination in WhatsApp. To that end, we adopted the three-step methodology depicted in Figure 1. First, we collected a dataset containing messages shared in a number of WhatsApp groups during a six-week period. The dataset was then broken into six non-overlapping snapshots, one for each week. The data was then processed to filter less relevant data and to identify (near-)duplicated content (Step 1). Next, we built a *media-centric network* for each snapshot. For each network, we proceeded to retrieve its *backbone* which,




Fig. 1: Overview of our Methodology

as argued later, is a necessary step to remove random and spurious edges (Step 2). Finally, we ran a community detection algorithm [6] to extract communities from each backbone and characterized their structural properties and temporal dynamics (Step 3). These steps are further discussed in the following sections.

3.1 Data Collection and Post-Processing

We used a dataset gathered by the WhatsApp Monitor [33], consisting of messages shared in political-oriented publicly accessible WhatsApp groups in Brazil. As described in [33], these groups were detected by searching for invitation links on public websites and online social networks. Our dataset covers six weeks around the 2018 general elections in Brazil (1st and 2nd rounds in October 7th and 28th, respectively), ranging from September 17th until October 28th in 2018. The dataset was broken into six non-overlapping one-week snapshots, and we restricted our analyses to 155 groups for which data was available in all snapshots.

Table 1 provides an overview of our dataset, showing the numbers of users, groups and messages (text, images, audios and videos) shared per week in the selected groups. It also shows the average number of users active per group, average number of messages shared per group, and average number of messages shared by a user in a group, for each week. Weeks including election dates (3 and 6) are highlighted in bold. As seen in the table, the numbers of users and messages vary over the snapshots. However, there tends to be an increase in activity around the dates of the two rounds of the election.

The collected data was then processed to extract and store the following information associated with each message: timestamp, anonymized user identifiers³, group identifier and the media(s) (text, image, audio or video) shared through the message. Next, we filtered out text messages shorter than 180 characters, as suggested in [33], so as to retain only those that most probably carry relevant information. Then, as a final processing, we ran a number of heuristics to identify (near)-duplicate content, a necessary step to build the media-centric networks (next section). The specific heuristic depends on the type of media. For textual content, we used the Jaccard Similarity Coefficient [20] to perform paired comparisons, using a threshold of 70% of similarity to detect near duplicates (as performed in [33]). For images, we followed [34]: we generated the Perceptual Hash (pHash) [43] of each file and grouped together those with same hash. The

³ Indeed, our data contains only cellular phone numbers. Thus, we are not able to identify the same user with multiple phone numbers.

Table 1: Overview of our Dataset (155 WhatsApp groups, 09/17 - 10/28/2018)

Weeks	1	2	3	4	5	6
# Unique users	4,994	4,774	5,115	4,815	4,439	4,914
Average # users/group	34.68	33.21	35.60	33.27	31.14	34.09
Average # messages/group	575	598	590	536	490	599
Avg # msgs/user (in a group)	16.59	18.01	16.58	16.13	15.75	17.59
# Text messages	89,136	92,650	91,438	83,118	75,982	92,840
# Image messages	13,018	13,208	13,274	13,471	11,922	17,113
# Audio messages	1,614	1,644	2,000	1,842	1,621	2,059
# Video messages	10,168	9,515	9,142	9,508	9,193	12,344

pHash algorithm works by detecting color variations on the image resulting on a hash value. By comparing hash values we are able to detect resized and modified images that are indeed the same content. Finally, for audios and videos, we used the name associated to each media file by WhatsApp during the data transfer.

We note that our near-duplicate identification process is limited by the approximation techniques used. As future work, we intend to explore more sophisticated and possibly accurate techniques, such as word embeddings [29] (for text messages), product quantization [21] (for images) and techniques based on audio and video content analysis [23, 38], which may enhance the generated network.

3.2 Network Model and Backbone Extraction

The network model used in this work creates an abstraction for user interactions in WhatsApp groups as a vehicle for information dissemination. Given that goal, the model focuses on the *content* shared, by connecting users who shared the “same message” at least once, regardless of whether they shared it on the same group or on different groups. By “same message” we mean messages that were identified as carrying near duplicate content, as described in Section 3.1. We refer to such networks as *media-centric* networks.

Specifically, given our dataset, we created a set of graphs $\mathcal{G} = \{G^1, G^2, \dots, G^{\Delta T}\}$, in which each G^w models user interactions during week w (i.e., $\Delta T = 6$ in our case). Each graph $G^w(V, E)$ is structured as follows. Each vertex in V refers to a user who posted a message during week w in one of the groups. An undirected edge $e = (v_i, v_j)$ exists in E if users corresponding to vertices v_i and v_j shared at least one message in common during week w . The weight of e corresponds to the number of messages both users shared in common during w .

As defined, our network model captures *many-to-many* interactions, i.e., interactions that occur among multiple (possibly more than two) users at once - in our case, co-sharing the same media content. This kind of interaction occurs in a range of other environments, such as the networks that emerge from relationships based on co-authorship, emails sent to a group of people and congressmen voting sessions [5, 11, 30], and raises different modeling challenges if compared to traditionally studied *one-to-one* interactions (e.g., friendship links) [4, 14].

In particular, modeling sequences of many-to-many interactions into a network may lead to the emergence of a (potentially large) number of spurious edges, reflecting random or sporadic user activities. Such spurious edges may pollute the network, obfuscating the real underlying structure that better represents the phenomenon under study. We illustrate this problem by taking a fictitious example in our context. Suppose two different scenarios: (1) one particular viral content is massively disseminated through the WhatsApp network as many users shared it in different groups and (2) a smaller set of users repeatedly spread the same content, reaching different audience which ultimately leads to a large spread. By simply looking at the topology of the network that emerges from these two scenarios, one may consider both groups of users in (1) and (2) as communities. However, we are here interested in identifying strong and consistent co-sharing behavior, as in (2), as opposed to sporadic connections, as possibly in (1). As more many-to-many interactions are added to the network, more edges of different natures are added, resulting in a richer but quite noisier topology. We want to remove this noise to be able to focus on what is fundamental to the large scale dissemination of content in the network.

In other words, we want to identify pairs of users (i.e., edges) who disproportionately shared messages in common, filtering out edges resulted from randomness and sporadic co-sharing, thus revealing the underlying *network backbone* [9]. By definition, the backbone contains only the *salient* edges more fundamentally related to the phenomenon under study (information dissemination, in our case).

There is a rich and vast literature on methods to extract the backbone from networks [9, 30, 37]. Here, we aimed to identify when an edge connecting two vertices reflects a strong connection between them when compared to the other spurious connections they both may have with other peers. We experimented with two state-of-the-art backbone extraction methods which are driven by that goal, namely Noise Corrected Method [9] and Disparity Filter Method [37], selecting the latter. Our choice follows the approach in [9]: we selected the method that, according to our experiments, delivers the best trade-off between the number of edges with lower weight removed and the structural connectivity of the resulted backbone. The latter was measured in terms of clustering coefficient and modularity metrics (discussed in the next section), preserving those edges with higher weights. The Disparity Filter Method was able to remove a larger number of spurious/sporadic edges while still maintaining modularity and clustering coefficient measures comparable to the complete network.

The Disparity Filter algorithm works as follows. Let's define the strength of vertex v_i as the sum of all edge weights attached to v_i . The algorithm considers an edge attached to v_i as salient if it represents a "large fraction" of v_i 's strength. Specifically, each edge attached to v_i is tested against the null hypothesis that the weights of all edges of v_i are uniformly distributed. Salient edges are those whose weights deviate significantly from this hypothesis. Notice that an edge is tested twice, once for each vertex it is incident to, and it is considered salient if it is statistically significant for both vertices when compared to a p -value. Edges that are not considered salient are removed from the graph. In our experiments

we adopted a p -value of 0.1. This value was selected in preliminary experiments, by running the algorithm with various options (ranging from 0.01 to 0.32) and choosing the one that led to the best tradeoff between statistical significance, number of remaining vertices in the backbone and backbone connectivity. This choice of p -value is consistent with prior studies on backbone extraction, which report that very small p -values lead to a large number of nodes removed, ultimately breaking the original network and turning the analysis unfeasible [37]. Thus the need for a choice that meets the aforementioned tradeoff.

3.3 Community Identification and Characterization

As a final step, we identify groupings of users who impact information spread in the network by, intentionally or not, sharing common content in a disproportionately high frequency. Studying these groupings, here referred to as *communities* (to avoid confusion with the original WhatsApp groups), reveals how they are structured, how such structure relates to the Whatsapp groups and how they evolve over time. Ultimately, we aim at bringing novel insights into how information virality may occur [34], despite the restrictions in group membership.

We identify user communities in the backbone extracted from each graph G^w ($w = 1 \dots 6$) using the Louvain algorithm [6]. It is a widely used community detection algorithm [32] that relies on a heuristic approach to iteratively build hierarchical partitions of the backbone. Specifically, it is based on a greedy optimization of the modularity, which is a metric of quality of these partitions. The modularity Q is defined [6] as $Q = \frac{1}{2M} \sum_{ij} \left[A_{ij} - \frac{k_i k_j}{2m} \right] \delta(c_i, c_j)$, where A_{ij} is the weight of edge (v_i, v_j) ; k_i (k_j) is the sum of the weights of the edges attached to v_i (v_j); m is the sum of all of the edge weights in the graph; c_i (c_j) is the community assigned to v_i (v_j); and $\delta(c_i, c_j) = 1$ if $c_i = c_j$ or 0 otherwise.

Note that the communities are built from pairs of users who share similar content more often than expected by chance. Such groupings could be driven by intentional behavior (i.e., by orchestration), by coincidence (as side effect of the general information diffusion process) or by a mix of both. We here do not distinguish between these effects, although the stronger the edge weights the greater the chance that some sort of coordination is in place. Characterizing the effects behind community formation is an interesting line of future research.

Our community identification process revealed a number of very small communities (e.g., three-four nodes), often organized as small trees. We chose to disregard small groupings (fewer than 5 vertices), focusing our analyses instead on the larger and more impactful communities. These communities were analyzed according to two dimensions: structural properties and temporal dynamics.

The analysis of structural properties aimed to quantify the quality of each grouping and what it represents to the diffusion of information during the period under analysis. For the former, we make use of the clustering coefficient metric computed for each community [31], which measures the degree to which vertices in the community tend to cluster together. For the latter, we compute the group coverage and the content coverage of each community. Group coverage is the

fraction of all monitored groups that were reached by posts from community members. Content coverage is the fraction of all contents shared during the period that were shared by the community members. Larger values of either metric reflect greater importance of the community to the information spread.

The analysis of temporal dynamics aimed to characterize the evolution of communities in the backbone and quantify changes in community membership over time. To that end, we used two metrics (as in [11]), always computed for snapshot w with respect to snapshot $w-1$. The first metric, called *persistence*, captures the continuous presence of the same users in the network backbone over time. It is computed as the fraction of users in the backbone at snapshot $w-1$ that remain in the backbone in snapshot w . This metric quantifies the presence of users who remain important for content dissemination over time.

Note that persistence does not distinguish between users who, despite remaining in the backbone over time, often change community from those who remain in the same community. The latter might reflect a potential coordinated effort to boost information spread. To capture the permanence of members in the same community, we adopted the Normalized Mutual Information (NMI) [26].

Let X and Y be the sets of communities identified in snapshots $w-1$ and w , respectively. Let also $P(x)$ be the probability of a randomly selected user being assigned to community x in X , and $P(x, y)$ be the probability of a randomly selected user being assigned to communities x in X and y in Y . Finally, let $H(x)$ be the Shannon entropy for X defined as $H(X) = -\sum_x P(x) \log P(x)$. The NMI X and Y is defined as $NMI(X, Y) = \frac{\sum_x \sum_y P(x, y) \log \frac{P(x, y)}{P(x)P(y)}}{\sqrt{H(X)H(Y)}}$. It can be thought as the information “overlap” between X and Y , or how much we learn about Y from X (and vice-versa). Its value ranges from 0 (all members changed their communities) to 1 (all members remained in the same communities).

In addition to characterizing communities, we also analyzed the importance of users to the information dissemination process. We did so by computing the impact on the (group and content) coverage metrics as vertices are removed from the backbone according to a ranking of importance. We experimented with different user/vertex rankings built based on metrics of activity level (number of messages shared and number of groups the user participates in) and metrics of centrality in the backbone. Our goal by comparing those rankings is to assess the extent to which the backbone and its communities are able to reveal important users to information dissemination in the system (as captured by the coverage metrics), compared to the activity metrics, priorly analyzed in [8, 28, 34].

We built four vertex rankings based on degree centrality and closeness (which relates to the average distance of the vertex to all other vertices) [3]. In each case, we use the standard metric, computed over the complete backbone, along with a tuned variation, referred to as community centrality. The latter, defined in [15], considers the embeddedness of the vertex in its community and the relations between this community and the others, and quantifies the vertex’s ability to disseminate information on its own community and on the overall network.

Specifically, the community centrality of a vertex v_i assigned to community c is computed by combining two components, a local one and a global one [15].

Table 2: Structural Properties of the Extracted Backbones

Metrics	Weeks					
	1	2	3	4	5	6
# Users	114	162	132	216	143	338
# Edges	346	767	500	1,676	617	3,499
# Connected components	1	1	1	1	2	2
Average clustering coefficient	0.49	0.61	0.58	0.63	0.57	0.63
Average degree	6.07	9.47	7.58	15.52	8.63	20.70
Average edge weight	6.95	9.10	8.05	12.66	7.36	8.91
# Communities	5	6	6	5	7	8
Modularity	0.70	0.61	0.55	0.54	0.59	0.57

The local component α_i^L quantifies the (degree or closeness) centrality of v_i in its own community c ; whereas the global component α_i^G quantifies the centrality of v_i in the backbone by considering only edges connecting vertices of different communities (i.e., inter-community edges). These two components are combined by a weighting factor μ_c that is the fraction of all inter-community edges that are incident to community c . In other words, the community (degree or closeness) centrality of vertex v_i , α_i , is defined as $\alpha_i = (1 - \mu_c) * \alpha_i^L + \mu_c * \alpha_i^G$.

4 Characterization Results

We now discuss the results of the characterization of the communities that emerged from the networks generated according to methodology described in section 3. Table 2 presents a summary of the structural properties of the backbones extracted from the networks in set \mathcal{G} . Once again, we show the snapshots containing 1st and 2nd rounds of the general elections (weeks 3 and 6) in bold.

We first note that the size of the backbone, in number of users and edges, varies greatly over the weeks. This variation is consistent, though in higher degree, with the variations in the amounts of participation and sharing activity in the monitored groups over the period, presented in Table 1. We also note that the backbones are formed by at most 2 connected components (often only one), with a reasonably strong average clustering coefficient (ranging from 0.49 to 0.63) and large average degree (6.07 to 20.7). These measures suggest well connected topologies and also hint at the formation of communities.

The table also shows the average edge weight in each backbone. These numbers should be analyzed in light of the average number of messages per user in a group, shown in Table 1. Note that the backbone edges represent a large fraction of all messages shared by the users, on average. Moreover, by combining average degree and average edge weight, we observe that each user in the backbone simultaneously shared multiple contents with many other users. All these results show the multiplicity of media co-sharing among users and highlight the need for investigating higher level user structures, notably communities. Indeed, as shown in the last two rows of Table 2, we identified from 5 to 8 communities in the backbones, with an overall quality in terms of structural connectivity




Fig. 2: Community Properties: (a) Size; (b) Average Clustering Coefficient; (c) Group Coverage and (d) Similarity of Group Coverage.

rather high (modularity between 0.54 and 0.70). Next, we delve deeper into these communities.

Figure 2 provides an overview of different community properties. Figure 2a shows the sizes (in number of nodes) of the communities identified in each snapshot (week), with each point representing a different community. We observe a great diversity of community sizes, normally constrained to fewer than 40 users, but it is also noticeable communities with more than 50 members. As an example, we identified a community with 26 members who shared political-driven content about presidential candidates and online campaigns on over 20 distinct WhatsApp groups in the week preceding the 1st round of election. This result suggests that the communities are dynamically built over time with variable number of users. We also correlate community size with average clustering coefficient, computed for community nodes, which is a measure of internal connectivity. Figure 2b shows these results for communities in all snapshots. We observe that most communities are strongly connected (even the larger ones) as the vast majority of them have average clustering coefficient above 0.50. Thus, in essence, the identified user communities are well structured, and, despite some size diversity over the snapshots, offer clear indications of consistent user co-sharing activity.

We now analyze the reach of these communities. Figure 2c shows a scatter plot with community size versus group coverage, for all communities in all snapshots. Recall that the latter is the number of groups the members of the communities participate in. There is a strong one-to-one relationship, in which the community size is strongly correlated with the number of groups it reaches. For larger communities, the sizes often are greater than the number of groups covered. These results suggest a broad reach in the ability to disseminate information, since communities often have members participating in multiple groups during the same time period. Moreover, the redundancy in some larger communities suggests some degree of robustness as well.

Taking a step further, we analyze the intersection of group coverage for different pairs of communities. Given two communities A and B , the intersection of B with A is the fraction of all groups covered by A that are also covered by B . Figure 2d shows these fractions for one snapshot – the third week ($w = 3$), which contains the 1st election round. Results for the other snapshots are similar, being thus omitted. In the figure, each cell in the upper triangle shows the intersection of community B (column) with community A (row), and each cell




Fig. 3: Impact of Removing Users in Decreasing Centrality Degree on Group (a) and Content Coverage (c) and in Decreasing Closeness (b and d, respectively).

in the lower triangle shows the intersection of A with B . In general, around half of the communities are built over distinct groups, due to the small similarity values (around 20%). On the other hand, some communities are built over similar groups, reaching up to 70% of coverage similarity. Such communities contribute to robust, distinct and parallel content dissemination on the WhatsApp network.

We also analyze communities' dissemination potential and their ability to endure perturbations, by assessing their robustness to member removals. We experiment with removing members based on attributes related to activity level, standard network centrality and node community centrality, and evaluate community robustness in terms of group coverage and content coverage. Results for one snapshot – the third week ($w = 3$) – are shown in Figure 3, where the x -axis represents the percentage of the top users, according to each attribute, that are removed ($n\%$ cut). Results for the other snapshots are quite similar.

Let's start by looking into group coverage. Figures 3a and 3b show that, by comparing the results for the attributes related to activity (number of messages shared and number of groups the user participates in) with those for standard (std.) network centrality metrics, we see that the latter are more relevant to detect the most important users for group coverage. This is true for both centrality metrics considered, i.e., centrality degree, which captures the node's ability to retrieve information from the network, and closeness, which relates to the node's efficiency to spread information in the network. Yet, the results for the node community centrality metrics (red lines in both figures) reveal that the communities are very important to the information dissemination through the network. This is due to the successful identification of the most important users and the positions they occupy both within its community and in relation to the whole network. Analogously, the same trend is observed for content coverage, as shown in Figures 3c and 3d. Thus, using the community structure for identifying the most important users to information dissemination is relatively more effective than simply checking the activity level or the standard centrality metrics.

Finally, we also analyze the temporal evolution of the backbone and its communities. First, we quantify the persistence of users in the backbone over consecutive weeks. As shown in Figure 4, a considerable amount of users (from 20% to 42%) remain in the backbone, repeatedly engaging on the weekly sharing activities. Yet, there is a large fraction of newcomers week after week. It could be due to new users who joined the groups (see Table 1) or simply reflect the replace-




Fig. 4: Persistence




Fig. 5: NMI




Fig. 6: Community Evolution

ment of non-persistent users by others as result of a natural diversity of user behavior over time. Regardless, the results indicate a highly dynamic backbone.

Focusing on the persistent users, we analyze how community members change over time based on NMI. As shown in Figure 5, we compute NMI considering all persistent users (blue line) and only the top 10 users who shared more content in the week (green line). Considering all persistent users, there is a high mobility of users across communities (low NMI). This is illustrated in Figure 6, which shows, from one week to the next, events associated with community membership such as splits, merges, births and deaths. In this figure, the wider the line, the greater the number of members migrating between communities, and the larger the diameter of the circle, the greater the number of members of that community. Once again, we see that persistent users in general often change community, engaging in different but strong co-sharing activities over time. However, if we zoom in the top 10 most active users (green line in Figure 5), we observe a stronger tendency to continue in the same community (higher NMI), suggesting that these most active users regularly share common content over time.

5 Conclusions and Future Work

This work analyzed the underlying network structure of information dissemination on WhatsApp publicly accessible groups in Brazil. By monitoring over 150 groups, our study revealed the formation of strongly connected user communities that cross the boundaries of traditional groups, fostering content spread at large. We found that these communities co-exist in the same groups constantly sharing broad content. By analyzing backbone and community centrality metrics, we were able to uncover users who are very important to the information dissemination, and those users would not be found if we looked only at their activity levels (numbers of messages and groups), as done in prior work. Moreover, by analyzing the temporal dynamics of the communities, we found that, despite constant changes in community membership, there is a number of users who persist in the network backbone over time, some of whom even remain tightly connected in the same community, suggesting coordinated efforts to boost information spread.

Future work includes analyzing the relations between content and community properties, zooming into spread of particular types of content (e.g., misinformation), and assessing user engagement and coordinated efforts in the communities.

Acknowledgements

This work was partially supported by grants from FAPEMIG, CNPQ and CAPES.

References

1. Abbe, E.: Community detection and stochastic block models: recent developments. *The Journal of Machine Learning Research* pp. 6446–6531 (2017)
2. Agrawal, K., Garg, S., Patel, P.: Spatio-temporal outlier detection technique. *International Journal of Computer Science & Communication* pp. 330–337 (2015)
3. Barabási, A.L., et al.: *Network science*. Cambridge University Press (2016)
4. Benson, A.R., Abebe, R., Schaub, M.T., Jadbabaie, A., Kleinberg, J.: Simplicial closure and higher-order link prediction. *Proceedings of the National Academy of Sciences* pp. 11221–11230 (2018)
5. Benson, A.R., Kumar, R., Tomkins, A.: Sequences of sets. In: *Proc. of the 24th ACM International Conference on Knowledge Discovery & Data Mining* (2018)
6. Blondel, V.D., Guillaume, J.L., Lambiotte, R., Lefebvre, E.: Fast unfolding of communities in large networks. *Journal of Statistical Mechanics: Theory and Experiment* (2008)
7. Bursztyn, V.S., Birnbaum, L.: Thousands of small, constant rallies: A large-scale analysis of partisan whatsapp groups. In: *Proceedings of the IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining* (2019)
8. Caetano, J.A., Magno, G., Gonçalves, M.A., Almeida, J.M., Marques-Neto, H.T., Almeida, V.A.F.: Characterizing attention cascades in whatsapp groups. In: Boldi, P., Welles, B.F., Kinder-Kurlanda, K., Wilson, C., Peters, I., Jr., W.M. (eds.) *Proceedings of the 10th ACM Conference on Web Science*. pp. 27–36 (2019)
9. Coscia, M., Neffke, F.M.: Network backboning with noisy data. In: *2017 IEEE 33rd International Conference on Data Engineering (ICDE)* (2018)
10. Ferreira, C.H.G., de Sousa Matos, B., Almeida, J.M.: Analyzing dynamic ideological communities in congressional voting networks. In: *International Conference on Social Informatics*. pp. 257–273. Springer (2018)
11. Ferreira, C.H., Murai, F., de Souza Matos, B., de Almeida, J.M.: Modeling dynamic ideological behavior in political networks. *The Journal of Web Science* (2019)
12. Fortunato, S., Hric, D.: Community detection in networks: A user guide. *Physics Reports* pp. 1 – 44 (2016)
13. de Freitas Melo, P., Vieira, C.C., Garimella, K., de Melo, P.O.S.V., Benevenuto, F.: Can whatsapp counter misinformation by limiting message forwarding? In: *International Conference on Complex Networks and Their Applications* (2019)
14. Fu, X., Yu, S., Benson, A.R.: Modelling and analysis of tagging networks in Stack Exchange communities. *Journal of Complex Networks* (2019)
15. Ghalmane, Z., Cherifi, C., Cherifi, H., El Hassouni, M.: Centrality in complex networks with overlapping community structure. *Scientific Reports* (2019)
16. Gilbert, F., Simonetto, P., Zaidi, F., Jourdan, F., Bourqui, R.: Communities and hierarchical structures in dynamic social networks: analysis and visualization. *Social Network Analysis and Mining* pp. 83–95 (2011)
17. Girvan, M., Newman, M.E.: Community structure in social and biological networks. *Proceedings of the national academy of sciences* pp. 7821–7826 (2002)
18. Goel, V.: How whatsapp leads mobs to murder in india (July 2018), <https://www.nytimes.com/interactive/2018/07/18/technology/whatsapp-india-killings.html>, "[Online; Posted on 18-Jul-2018]"
19. Hoffmann, T., Peel, L., Lambiotte, R., Jones, N.S.: Community detection in networks without observing edges. *Science Advances* (2020)
20. Jaccard, P.: Etude de la distribution florale dans une portion des alpes et du jura. *Bulletin de la Societe Vaudoise des Sciences Naturelles* pp. 547–579 (1901)

21. Jégou, H., Douze, M., Schmid, C.: Product quantization for nearest neighbor search. *IEEE Transactions on Pattern Analysis and Machine Intelligence* pp. 117–128 (2011)
22. Kietzmann, J.H., Hermkens, K., McCarthy, I.P., Silvestre, B.S.: Social media? get serious! understanding the functional building blocks of social media. *Business Horizons* pp. 241 – 251 (2011)
23. Kordopatis-Zilos, G., Papadopoulos, S., Patras, I., Kompatsiaris, I.: Fivr: Fine-grained incident video retrieval. *IEEE Transactions on Multimedia* pp. 2638–2652 (2019)
24. Kumar, R., Novak, J., Tomkins, A.: Structure and evolution of online social networks. In: *Proceedings of the 20th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (2006)
25. Kumar, S., Shah, N.: False information on web and social media: A survey. *Social Media Analytics: Advances and Applications* (2018)
26. Lancichinetti, A., Fortunato, S.: Detecting the overlapping and hierarchical community structure in complex networks. *New Journal of Physics* (2009)
27. Magenta, M., Gagnani, J., Souza, F.: How whatsapp is being abused in brazil’s elections (October 2018), <https://www.bbc.com/news/technology-45956557>, "[Accessed on 24-May-2020]"
28. Maros, A., Almeida, J., Benevenuto, F., Vasconcelos, M.: Analyzing the use of audio messages in whatsapp groups. In: *Proceedings of The Web Conference* (2020)
29. Mikolov, T., Sutskever, I., Chen, K., Corrado, G.S., Dean, J.: Distributed representations of words and phrases and their compositionality. In: *Advances in Neural Information Processing Systems* 26, pp. 3111–3119. Curran Associates, Inc. (2013)
30. Newman, M.: Network structure from rich but noisy data. *Nature Physics* pp. 542–545 (2018)
31. Onnela, J.P., Saramäki, J., Kertész, J., Kaski, K.: Intensity and coherence of motifs in weighted complex networks. *Physical Review E* (2005)
32. Papadopoulos, S., Kompatsiaris, I., Vakali, A., Spyridonos, P.: Community detection in social media. *Data Mining and Knowledge Discovery* pp. 515–554 (2012)
33. Resende, G., Melo, P., C. S. Reis, J., Vasconcelos, M., Almeida, J.M., Benevenuto, F.: Analyzing textual (mis)information shared in whatsapp groups. In: *Proceedings of the 10th ACM Conference on Web Science. WebSci '19*, Association for Computing Machinery (2019)
34. Resende, G., Melo, P.F., Sousa, H., Messias, J., Vasconcelos, M., Almeida, J.M., Benevenuto, F.: (mis)information dissemination in whatsapp: Gathering, analyzing and countermeasures. In: *The World Wide Web Conference, WWW 2019*, San Francisco, CA, USA, May 13–17, 2019. pp. 818–828. ACM (2019)
35. Rossetti, G., Cazabet, R.: Community discovery in dynamic networks: a survey. *ACM Computing Surveys (CSUR)* pp. 1–37 (2018)
36. Rossetti, G., Pappalardo, L., Pedreschi, D., Giannotti, F.: Tiles: an online algorithm for community discovery in dynamic social networks. *Machine Learning* pp. 1213–1241 (2017)
37. Serrano, M.A., Boguna, M., Vespignani, A.: Extracting the multiscale backbone of complex weighted networks. *Proceedings of the National Academy of Sciences* p. 6483–6488 (2009)
38. Song, J., Yang, Y., Huang, Z., Shen, H.T., Hong, R.: Multiple feature hashing for real-time large scale near-duplicate video retrieval. In: *Proceedings of the 19th ACM International Conference on Multimedia*. p. 423–432 (2011)

39. Tang, L., Liu, H., Zhang, J., Nazeri, Z.: Community evolution in dynamic multi-mode networks. In: Proceedings of International Conference on Knowledge Discovery and Data Mining. pp. 677–685 (2008)
40. WhatsApp: Two Billion Users on WhatsApp. <https://blog.whatsapp.com/10000666/Two-Billion-Users-Connecting-the-World-Privately>, accessed 19 May 2020
41. WhatsApp: About WhatsApp. <https://www.whatsapp.com/about/> (2020), accessed 19 May 2020
42. WhatsApp: Keeping WhatsApp Personal and Private. <https://blog.whatsapp.com/Keeping-WhatsApp-Personal-and-Private> (2020), accessed 19 May 2020
43. Zauner, C., Steinebach, M., Hermann, E.: Rihamark: perceptual image hash benchmarking. In: Media watermarking, security, and forensics III. vol. 7880 (2011)
44. Zhang, X.S., Wang, R.S., Wang, Y., Wang, J., Qiu, Y., Wang, L., Chen, L.: Modularity optimization in community detection of complex networks. EPL (Europhysics Letters) (2009)