

Structural and Dynamical Patterns on Online Social Networks: The Spanish May 15th Movement as a Case Study

Javier Borge-Holthoefer¹, Alejandro Rivero¹, Iñigo García³, Elisa Cauhé¹, Alfredo Ferrer¹, Darío Ferrer¹, David Francos¹, David Iñiguez^{1,4}, María Pilar Pérez¹, Gonzalo Ruiz¹, Francisco Sanz¹, Fermín Serrano¹, Cristina Viñas¹, Alfonso Tarancón^{1,2}, Yamir Moreno^{1,2,5*}

1 Instituto de Biocomputación y Física de Sistemas Complejos (BIFI), Universidad de Zaragoza, Zaragoza, Spain, **2** Departamento de Física Teórica, Universidad de Zaragoza, Zaragoza, Spain, **3** Cierzo Development S.L., Zaragoza, Spain, **4** Fundación ARAID, Diputación General de Aragón, Zaragoza, Spain, **5** Complex Networks and Systems Lagrange Lab, Institute for Scientific Interchange, Torino, Italy

Abstract

The number of people using online social networks in their everyday life is continuously growing at a pace never saw before. This new kind of communication has an enormous impact on opinions, cultural trends, information spreading and even in the commercial success of new products. More importantly, social online networks have revealed as a fundamental organizing mechanism in recent country-wide social movements. In this paper, we provide a quantitative analysis of the structural and dynamical patterns emerging from the activity of an online social network around the ongoing May 15th (15M) movement in Spain. Our network is made up by users that exchanged tweets in a time period of one month, which includes the birth and stabilization of the 15M movement. We characterize in depth the growth of such dynamical network and find that it is scale-free with communities at the mesoscale. We also find that its dynamics exhibits typical features of critical systems such as robustness and power-law distributions for several quantities. Remarkably, we report that the patterns characterizing the spreading dynamics are asymmetric, giving rise to a clear distinction between information sources and sinks. Our study represents a first step towards the use of data from online social media to comprehend modern societal dynamics.

Citation: Borge-Holthoefer J, Rivero A, García I, Cauhé E, Ferrer A, et al. (2011) Structural and Dynamical Patterns on Online Social Networks: The Spanish May 15th Movement as a Case Study. PLoS ONE 6(8): e23883. doi:10.1371/journal.pone.0023883

Editor: Matjaz Perc, University of Maribor, Slovenia

Received: July 8, 2011; **Accepted:** July 30, 2011; **Published:** August 19, 2011

Copyright: © 2011 Borge-Holthoefer et al. This is an open-access article distributed under the terms of the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

Funding: J.B.-H. is supported by Spanish Ministerio de Ciencia e Innovación (MICINN) through project FIS2008-01240. D.I. is supported by the Government of Aragón through a Fundación Aragón I+D contract. Y.M. was partially supported by the Future and Emerging Technologies Open project DYNANETS (grant no. 233847) funded by the European Commission, by Spanish Ministerio de Ciencia e Innovación (MICINN) through projects FIS2008-01240 and FIS2009-13364-C02-01 and by Comunidad de Aragón (Spain) through the project FMI22/10. The authors also acknowledge the Spanish Ministerio de Ciencia e Innovación (MICINN) for financial support through the project FIS-164-50. The funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript.

Competing Interests: Y.M. is a PLoS ONE Academic Editor. The authors declare that no other competing interests exist.

* E-mail: yamir.moreno@gmail.com

Introduction

Modern online socio-technological systems are producing a deep change in our traditional networking paradigms and in the way we communicate with each other. At the same time, online social media constitute nowadays efficient and fast means to group together many social agents around a common issue. In this way, new types of economic, financial and social phenomena are arising. An example of the latter is given by the so-called Arab revolts, which have materialized thanks to these new communication platforms. The protests have been mediated by the use of social networks such as Facebook, Twitter and YouTube, which have been critical for the birth and consolidation of campaigns involving strikes, demonstrations, marches and rallies.

On the other hand, online social networks not only modify in a radical way the dynamics of information and opinion spreading, but are also making our world even more global. More importantly, these platforms generate an enormous amount of time-stamped data, making it possible for the first time to study the

fast dynamics associated to different spreading processes at a system-wide scale. These novel and rich data niches allow testing different social dynamics and models that would otherwise be highly elusive with traditional data-gathering methods. Additionally, the availability of data enables the study of phenomena that take place on time scales ranging from a few minutes or hours to a year-long duration. An example of the former kind of fast dynamics is given by large sport or cultural events, whereas cooperative content production such as the case of Wikipedia typically occurs in months or even years, thus, in a much slower time scale.

In this paper, we study the structural and dynamical patterns of the network made up by twitter users who have been involved in a social phenomenon that is currently taking place in Spain: the so-called May 15th movement (henceforth referred to as 15M). This movement-in-the-making had been brewing for a while in the social media, but took off on May 15th when the first demonstrators camped in a central square in Madrid, Spain. From that moment on, the protests and camps spread throughout the country. As many of the adherents are online social media users,

the growth and stabilization of the movement was closely reflected in time-stamped data of twitter messages, which we have gathered and analyzed. This will allow us to elucidate the mechanisms driving the emergence of this kind of social phenomenon, and whether it shares dynamical and structural features with other natural, social and technological processes [1,2]. Additionally, on more general scientific grounds, a social phenomenon like the 15M movement is an excellent opportunity to understand network formation processes and online spreading dynamics. The ultimate aim is to further advance our understanding of this kind of dynamics and eventually be able to make predictions based on real time data monitoring.

In what follows, we present the results of our analysis. On the one hand, we statistically characterize the structural patterns of the network of users who sent or received tweets containing keywords related to the 15M movement. We find that this network displays the typical features of other networks in Nature such as scale-free degree distributions, a community structure at the mesoscale and high structural robustness [1,2]. On the other hand, we have also analyzed the dynamical patterns characterizing the spreading of information over the 15M network. Our results show that the 15M diffusion dynamics is highly asymmetric. Admittedly, a relative large fraction of the system is actively trafficking, but a great part of the overall traffic is delivered to a few users that do not pass them anymore, thus constituting a sort of information sinks. We round off our analysis by comparing our results with those reported in the literature for other kinds of online dynamical processes.

Methods

Ethics Statement

The data are publicly available as they come from a public online social site (Twitter). However, any private information has been analyzed anonymously.

Data Analysis

The data used in this study is a set of messages (tweets) that were publicly exchanged through www.twitter.com. The whole time-stamped data collected comprises a period of one month (between April 25th, 2011 at 00:03:26 and May 26th, 2011 at 23:59:55) and it was archived by a local start-up company, *Cierzo Development Ltd* using the SMMART Platform. This platform is evolving into a new concept called “Open Social CRM,” which combines concepts in monitoring tools, CRM tools, social tools and a philosophy of open innovation. The company restricts its collection to messages in Spanish language that come preferentially from users within or related to Spain. The internals of data collection are private to the company, but basically 23 hours of data are homogeneously collected each day, always leaving the same timeframe (16:00 to 17:00 CET time) to readjust the database due to the introduction of new Spanish nodes, purging of the non-Spanish related ones, etc.

To filter out the whole sample and choose only those messages related to the 15M movement, we selected 70 keywords (*hashtags*), which were systematically used by the adherents to the demonstrations and camps. Next, the extracted sample was examined for missing hashtags — of the top ten only one of them was not in the selected set, this being likely related to its bilingual nature *#acampadabcn*. The filtered data set appears to be representative enough of the total traffic related to the 15M movement produced during the period analyzed. As a matter of fact, a comparison with other databases, such as *topsy.com*, which aims to collect the whole set of twitter messages, shows that for the

same period there were about 390.000 messages with the word “*acampadasol*” and 190.000 for the hashtag “*#nolesvotes*”. Our sample is made up of 189.000 and 66.000 messages and hashtags, respectively, i.e., roughly above a third of the total number of messages.

Once this process is finished, the final sample consists of 581.749 tweets, out of which 46.557 were identified as *retweets* of unknown origin, and therefore were discarded. On its turn, these tweets were produced by 85.851 different users. To complete the data set, we located the references to other users inside each message. These references are marked in the system by an arobase, “@username”. A user receives a notification, usually via email, each time a mention happens, and the messages having mentions are also copied to a special tab in the user interface. The total number of messages having at least a reference was 151.222. In some cases, the tweet is addressed to more than one user, so that the total pairs (source, target) extracted from these messages was actually higher: 206.592. This is the number of directed arrows in our network. We stress again that our network is a dynamical instance of a larger underlying network (i.e., that made up of followers and followings in twitter).

Finally, although not directly related to the study presented here but important for complementary studies, data for all the involved users were scrapped directly from *twitter.com* using a cloud of 128 different nodes of a subnet. The scrap was successful for 84.229 users, for whom we also obtained their official list of followers. Moreover, about a half of them can be associated to a location (city), which is later translated to geographical coordinates via a standard geo-localization service from *Yahoo*. It is worth remarking that the extraction of followers gave a list in the order of 3 millions users, which roughly coincides with the order of the audience estimated by *Twitter* in Spain.

Results and Discussion

The availability of time-stamped data allows us to closely track the birth and development of the network made up of users who exchanged tweets related to the 15M movement during the period analyzed. In this network, every node represents a user while a link between two nodes is established whenever they exchange a message. We have made a movie (see <http://15m.bifi.es/index.php> for a high resolution version with downloading options) that reproduces the temporal evolution of the networks and the dynamics of messages exchange during the period analyzed. The reader will note the highly dynamic character of the network as well as the dimension of the social phenomenon being analyzed.

The network constructed as described above is weighted and directed, i.e., a link from i to j means that i sent at least a message to j and the weight of the link $i \rightarrow j$ stands for the actual number of such messages. Therefore, the adjacency matrix of the network is not symmetric (j does not necessarily send a message to i). Moreover, it is worth noticing that we always work with accumulated data, such that the network at a certain time t includes every message (link) produced at any time $t' \leq t$, i.e., once a link is established, it connects the two end nodes forever. Once the network is built, we are able to characterize it from a topological point of view, within the framework of complex network theory [1,2]. In doing so, we discuss local and global descriptors, as well as the network structure at the mesoscale level. Additionally, we also analyze the dynamics of information spreading over the 15M network and compare our results with those already reported for other dynamical processes that are mediated by the Web 2.0.

Network Growth and Structure

The first point of interest concerns the structural growth pattern. We wonder whether a collective mobilization of thousands of agents demands a slow, progressive increase in size; or rather social networking platforms enable an abrupt emergence. In Figure 1 (top) we present three snapshots of the system for different days, relative to day D (May 15th). Colors stand for the “age” of the node: early active users are coded in yellow, where those that adhere the network in successive days are coded in green, red, etc. Black is left for the latest adopters (people whose activity began at $D+10$). Besides, the size of the nodes has been made proportional to their activity, taking into account both incoming and outgoing tweets (however, for the sake of clarity, such proportion has been truncated at $k_{in}+k_{out}=200$ in the networks displayed in the figure). Even this simple representation of the evolution of the 15M network is already indicative of the growth in the number of agents once the movement took off and time goes on.

The results depicted in the bottom panel of Figure 1 further illustrate the way in which the network evolves by gaining adherents. The figure represents the proportion of active nodes at time t (with a resolution of 12 hours) in the giant component relative to the total number of users in the network at the end of the growth process. As we can see from the figure, the formation of the network and its later increase in size does not proceed in a gradual proportional process but in a sequence of bursts concentrated in just a few days (from day D to day $D+7$). Obviously this process is driven by the events surrounding the movement: as mentioned, at day D the protesters decided to camp at *Puerta del Sol* square in Madrid, which in turn elicited huge attention from the media and made the difference as far as the spread of the movement to a country wide scale concerns. Besides, from our data, it appears that the number of active users saturates after $D+7$. It is interesting to note that in May 21st ($D+6$), the

day preceding local and regional elections, more than the 80% of the network was already formed.

Beyond structural growth, a second key aspect of the network under study concerns the distributions of strengths. The strength s of a given node i is defined, as usual, by the sum of the weights of the links that are incoming and outgoing to node i . In our case, it is also important for the discussion that will follow, to further divide this magnitude into two contributions. One the one hand, we have the strength derived from the weights of links incident to the node, s_{in} . This magnitude accounts for the total traffic (number of tweets) that a given node receives from its neighbors. Conversely, s_{out} represents the sum of the traffic generated at a node, i.e., the number of tweets this user sends out. Additionally, let $P(s_{in})$ and $P(s_{out})$ be the cumulative distributions of both strengths, which we can be measured at different instants t of the network development.

Figure 2 shows the cumulative distributions of the previous quantities for several times. As can be seen, even before the occurrence of the events that triggered public protests on day D , both $P(s_{in})$ and $P(s_{out})$ follow power-laws $P(s) \sim s^{-\gamma}$, but with different exponents ($\gamma_{in}=1.1$ and $\gamma_{out}=2.3$, respectively, as measured at $D+10$). Similar plots for the degree of the nodes exhibit the same behavior. It is well-known that the statistical properties of these variables in other technological, social and natural systems are also heterogeneously distributed. Therefore, the fat-tailed distributions that characterize the topology of the 15M network are not unique, but are rather ubiquitous in Nature. Nonetheless, the fact that the 15M network is scale-free, has deep consequences regarding a number of relevant issues including its origin, complexity, robustness and, from a dynamical point of view, the way in which information flows over the system. As the network obtained comes from the activity of the nodes, the heavy-tailed distribution of both nodes’ degrees and strengths suggests a dynamics lacking any typical or characteristic scale.

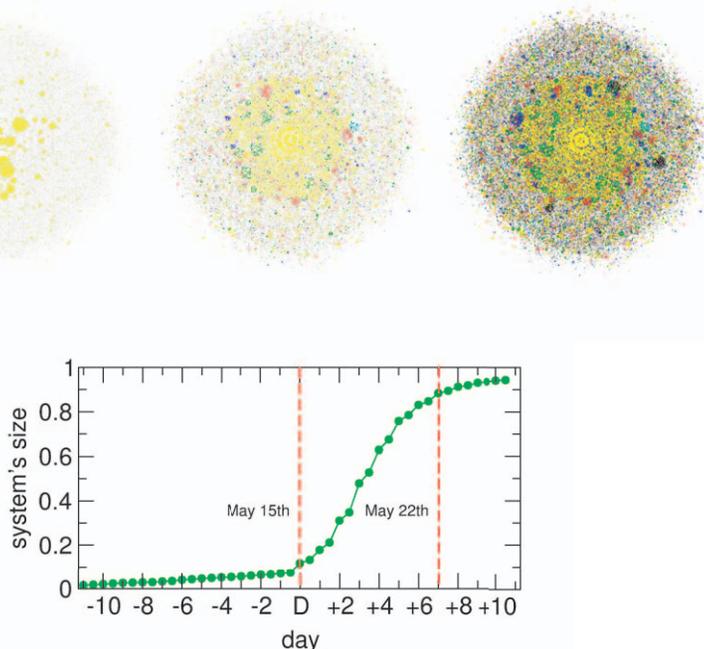


Figure 1. Evolution of the 15M Network. (Top) Cumulative snapshots of the system (at day D , $D+5$ and $D+10$). Different colors stand the age of the nodes (yellow for former adherents and black for latecomers). (Bottom) Evolution of the system’s giant component, relative to the final size of the network. Note that growth does not proceed progressively, but it rather explodes from day D due to socially relevant events on that date. doi:10.1371/journal.pone.0023883.g001

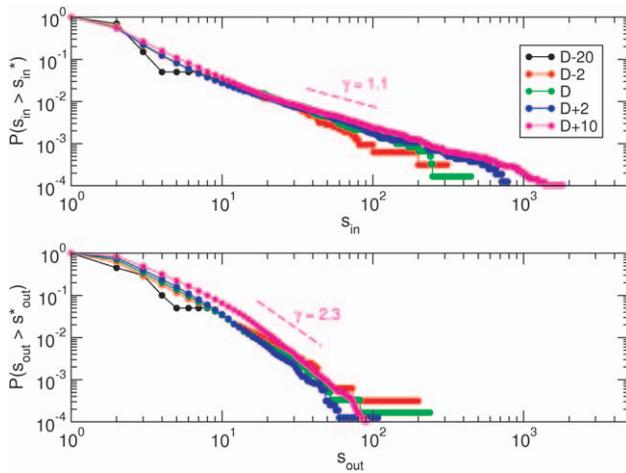


Figure 2. Cumulative in- and out-strength distributions. Strength distributions for both received (top) and sent (bottom) messages display a power-law behavior as early as at $D-2$. The fat-tailed distributions indicate that the 15M network is scale-free, with the implications this fact bears. Note that the exponents that define the power-laws differ significantly between sent and received messages. doi:10.1371/journal.pone.0023883.g002

On the other hand, the dynamical asymmetry between incoming and outgoing degrees or strengths is not surprising either. Indeed, individual behavior, which ultimately determines the resulting (out) dynamics, is an intended social action, but the emergent properties of the collective behavior of agents are unintended [3]. Essentially, subjects decide when and to whom a given message is sent. Therefore, the aggregate behavior of all agents and their popularity (i.e., how many incoming links a node has) result from individual choices. This is what is reflected in the in and out distributions. As a matter of fact, the exponent of the power law characterizing the degree probability distribution $p(k)$ lies in the interval [2,3], as usually found in most real-world networks. Interestingly, spreading dynamics such as rumor and disease propagation processes are most efficient for scale-free networks whose exponent is precisely in this range [4–6]. Finally, the strength distribution for the tweets sent, $p(s_{out})$, also resembles a power law function with an exponent larger than 3, although in this case the distribution exhibits an exponential cut-off. This might be due to the fact that sending messages has an associated cost in terms of bandwidth availability, the cognitive capacity to

produce different messages and ultimately an unavoidable physical limitation to type them [7–9].

Another aspect of capital interest regards how the overall traffic is generated and delivered. One of the main consequences of the functional form of the strength distributions is presented in Figure 3. The emergence of hubs, namely, the signaling feature of scale-free networks, leads to a predictable oligopoly in the way information is spread. In Figure 3, we observe that the number of tweets sent grows with the number of active users of the network. The curves corresponding to different days (i.e., instances of the network) nearly collapse into a single one. This means that as users join the network, the traffic generated scales accordingly. Moreover, the figure indicates that, for instance, roughly the 10% of active subjects generate the 52% of the total traffic. This is another indication of the dynamical robustness of the network to random failures but at the same time of its fragility to attacks directed towards that 10% of users. More remarkably, the results depicted in the figure are in sharp contrast with the activity patterns corresponding to received tweets. In this latter case, as time goes on, the number of in-strength hubs decreases. As shown in the figure, by $D+10$, less than 1% of users receive more than 50% of the information. As we will show later on, these nodes correspond to authorities or mass media, which the adherents identify as main receptors (government) of or potential spreaders (mass media) for their messages. However, what a priori seems to be a good choice, turns out to be harmful for the process of information spreading. As a matter of fact, we have checked that these hubs, which we call *information sinks* do receive a lot of messages but rarely act as spreaders within the network. As a consequence, almost all messages that arrive to those nodes are not redelivered and hence lost. In this sense, our results show that while the delivering of information is shared by a relative large number of users that keep the “social temperature” of the movement, most of this information is simply directed towards a few highly connected targets that might not pass the voice any longer (i.e., they are not active spreaders). Nonetheless, the information exchanged is public and users can therefore access it. This would however imply an individual action (to check a given user’s timeline) that is not captured in our twitter data.

Community Structure

The modular structure is pervasive in many natural, social and technological networks. Generally speaking, modules are islands of highly connected nodes separated by a relatively small number of links. This meso-level skeleton is likely to be relevant to understanding dynamical processes in networked systems. Agents

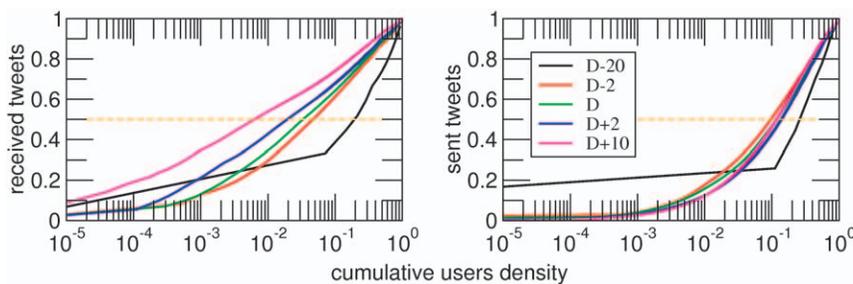


Figure 3. Information flow. The figure represents the density of tweets received (left) and sent (right) as a function of the cumulative fraction of active users. For each day, data are normalized by the number of active users at that date. As a reference, the horizontal line corresponds to 50% of emitted/received tweets. Note that, on $D+10$, less than 1% of the nodes receive half of the messages. On the contrary, the pattern of tweets sent hardly evolves from the beginning of the movement: 10% of the active nodes produce 50% of the messages. This asymmetry is coherent with the differences observed for the strength distributions. doi:10.1371/journal.pone.0023883.g003

information. Our results show that modules are highly hierarchical and that nodes that are central to their communities, i.e., *local hubs*, are mostly hubs at the global scale as well.

The mesoscale structure allows to get deeper insights into social aspects of our case study. First, tags identifying the 30 largest communities are highly heterogeneous. 6 of these modules correspond to important mass media (newspapers and television), which points to the otherwise intuitive fact that users rely on these agents to amplify their opinion. The same can be said of 3 modules corresponding to famous journalists. More interestingly, 7 modules correspond to on-line activists and/or veteran bloggers. These agents are unknown to most people, but they are present in the network from its birth and enjoy a solid reputation that facilitates their being considered a reference in the movement. Remarkably, 7 modules are formed by camps in 7 different cities. Madrid is of course the main one, as the movement began there (*acampadasol*, which comprehends over 3000 nodes). Other cities are Barcelona, Granada, Zaragoza, Valencia, Seville and Pamplona.

The fact that communities are geographically defined suggests some additional conclusions: (i) the mesoscale reflects the autonomy of each of the assemblies throughout the Spanish geography. Each of these modules hardly connects to any other, indicating a low communication between them; (ii) the exception to the previous point is Madrid: each minor camp holds a strong communication interchange with the community represented by *acampadasol*. Taking points (i) and (ii) together, it can be safely said that the movement is highly centralized, because in most cases a peripheral settlement is only influenced by Madrid and one or two minor ones. Finally, (iii) despite the potential of Web 2.0 communication platforms, data indicates that these media are mostly used to communicate with geographically close people. In other words, the network is *global*, but communication is mainly *local*. This is further verified in Table 1, where we have summarized the percentage of people whose geolocated information coincides with that of the module (city).

Popularity Evolution

The Web 2.0 has brought to network science the challenge to deal with highly dynamic, changing structures. Besides source and target nodes, and a (perhaps weighted) link between them, one must now consider a new ingredient: time. In this context, an interesting issue is related to the evolution of particular nodes:

understanding how an element (be it a Wikipedia entry or a novel trend in social networks) comes to existence (appears in the network) and develops. Of further interest is to elucidate how a subset of these network’s components ends up as a “popular” entity. This is a key aspect in network development, as popular agents eventually have an impact on other agents’ opinions, acting as a referent, be those opinions related to politics, culture or business.

To capture the dynamics of popularity, we follow the framework recently proposed in [3,20]. The natural quantity to measure popularity in a communication network is the number of messages that arrive at a node, which corresponds to that node’s in-strength s_{in} , and the rate at which s_{in} changes. Hence, a way to grasp how the activity of a node evolves is to consider its logarithmic derivative $[\Delta s/s]_t = (s_t - s_{t-1})/s_{t-1}$, i.e., the relative variation of strength in a time unit (we omit subindex “in” for clarity). Figure 5 displays the evolution of the latter variable for some arbitrarily chosen nodes among those that are information sinks. Beyond an initial surge typically observed in many nodes, the time series of the logarithmic derivative evidence a bursty behavior. Fluctuations depend on exogenous events, in a strong parallelism with the external circumstances that drive the whole network’s strong changes. It is noteworthy that these patterns closely resemble other, less conflictive, examples of popularity evolution in the Wikipedia or the Web [20].

On the other hand, Figure 6 shows how bursts are distributed according to their magnitude for two different time intervals but the same time granularity (1 day). The observed pattern, which is the same regardless of the time intervals under consideration, clearly shows heavy-tailed distributions, again in close resemblance to results already reported for other web-mediated dynamics [20] and a variety of critical phenomena in physical, economic and social systems. As a matter of fact, a simple model can account for the observed bursts distribution. The so-called *rank model* is specially conceived for networks in which prestige, rather than degree-based preferential attachment, plays a central role to determine nodes’ connectivity [21]. The rank model depends on a prestige measure that is used to rank nodes. In this model, the probability that a new node that joins the network at $t+1$ connects to an older one j is given by

$$p(t+1 \rightarrow j) = \frac{R_j^{-\alpha}}{\sum_{i=1}^t R_i^{-\alpha}} \tag{1}$$

where R_j is the rank of node j and $\alpha > 0$ determines the exponent γ of the resulting power-law degree distribution $p(k)$, such that $\gamma = 1 + \frac{1}{\alpha}$. Figure 7 compares the bursts distributions resulting from the data and from the model. The results shown correspond to the case in which nodes are ranked according to their age and α has been set to 0.9091 (for this value we obtained the best fit to the actual degree distribution of the empirical network). Moreover, we note that in order to simulate the non steady growth of the network, we have considered for the synthetic case that 1 day has gone by when the size of the network being generated coincides with that of the empirically assembled network for that time.

As we can see from the figure, even under the simplest rank rule, the burst magnitude distribution is nicely reproduced. As for our specific social context, the previous results do not imply that the driving mechanism behind the evolution of the 15M network is that simple, but illustrate that bursty activity of this sort can be produced by generic mechanisms. Given that this self-organized activity is widespread in nature, there are no reasons to consider that the network has its origin in external actions. Nonetheless, it

Table 1. Geographic origin of nodes in region-based communities.

community tag	area	fraction of users from same area
@acampadasol	Madrid	54%
@acampadabcn	Barcelona	81%
@acampadavlc	Valencia	63%
@acampadazgz	Zaragoza	82%
@acampadagranada	Granada	53%
@acampadasevilla	Sevilla	83%
@15MPamplona	Pamplona	71%

Region-based modules are mostly formed by nodes whose geographical origin coincides with that of the most central node in the community. This statement is clear for almost all these modules, except in the case of Madrid and Granada. The case of Madrid is not surprising, given that @acampadasol is the reference of the whole movement, thus the community organized around this actant is a more heterogeneous one. Granada is a more intriguing exception.
doi:10.1371/journal.pone.0023883.t001

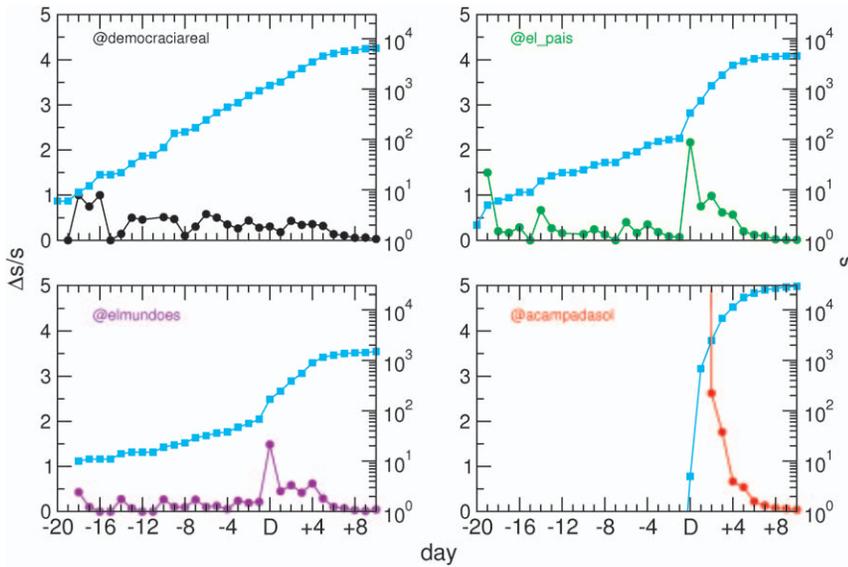


Figure 5. Popularity dynamics. The increase in the strength of the selected nodes (cyan curves) markedly changes after day D . The logarithmic derivative $\Delta s/s$ provides a finer interpretation in terms of the bursts in popularity. *democraciareal* (upper-left) had been an active agent long time before D , seemingly the movement was gathering strength (note some remarkable increments between $D-20$ and $D-16$). A newer user in the social network (*acampadasol*; lower-right) emerges and quickly takes over from D onwards as the reference of the whole movement. On the other hand, two right- and left-wing newspapers, *El Mundo* and *El País* respectively, undertake similar changes, which indicates that a large sets of nodes identified them as relevant actants in the context of protests. Note, however, that this does not imply that the same nodes are active senders too. As a matter of fact, they are not.
doi:10.1371/journal.pone.0023883.g005

must be pointed out that although the rank model grasps the main observed trends, both regarding structural characteristics and dynamical facts, yet other important aspects extracted from the data are not captured by the model. These limitations demand model refinements beyond the scope of the present study.

Conclusions

The social phenomenon here presented as a case study is a collective endeavor that is expressed at many levels, ranging from public demonstrations and camps to the presence of news in the

mass media. In this work we have analyzed data from time-stamped, online activity in a specific social networking site during the formation and stabilization of this social movement. From a scientific point of view, these data represent a challenge as most network studies typically deal with a static structure.

Undoubtedly, there are many facts that can be easily identified as the grounds of the 15M movement. Among them, the worldwide economic crisis and the impact it has had on society. Nonetheless, the particular events that triggered the growth of the whole movement remain unknown and are beyond the scope of

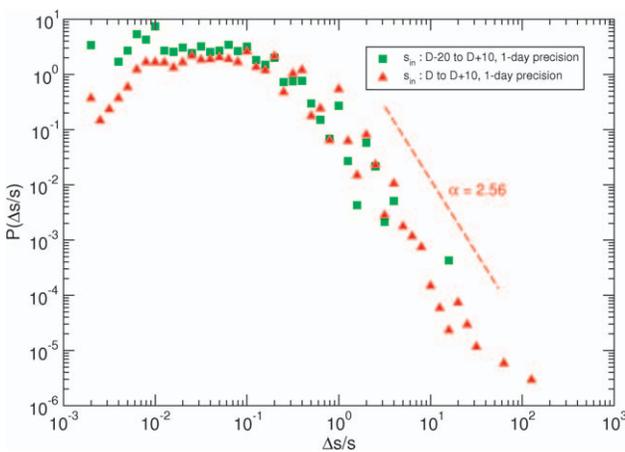


Figure 6. Burst size distribution. Popularity “efervescence” in the whole system (each node). Unsurprisingly, most nodes hardly undergo noticeable changes in their popularity. However, a small fraction of nodes does experience significant increases (heavy tail). This pattern is not exclusive of this communication network.
doi:10.1371/journal.pone.0023883.g006

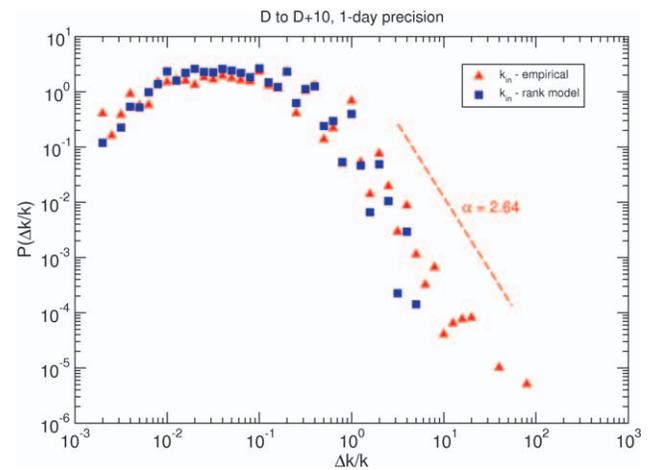


Figure 7. Empirical and simulated burst size distribution. The rank model, even at its simplest form, successfully reproduces not only empirical structural features, but also the prestige or popularity evolution over time. See the text for more details about the model implemented.
doi:10.1371/journal.pone.0023883.g007

this work. Addressing these questions would probably require an in-depth semantic analysis of the contents of interchanged messages. From its onset, our statistical characterization of the communication network built from tweets exchanged between adherents (and opponents) reveals a strong resemblance to well-known phenomena in natural and manmade systems, which are admittedly self-organized.

Additionally, the 15M movement also raises relevant questions with sociological consequences. We argue that information centralization (Figure 3), as well as patterns of popularity growth (Figures 5 and 6) are indicative of a tendency towards a hierarchical structure. Opinion leaders emerge spontaneously and minor actants devote much energy to communicate with them (be it to have their ideas echoed, or to influence such leaders). This proclivity is coherent with economy of attention [22], i.e., the system tends to avoid the overabundance of opinions to prevent scarcity of attention, but raises doubts about the possibility of converging to an egalitarian social system in which information flows and is received in an efficient manner. As far as our analysis concerns, we have shown that in a dynamics such that the one at work for our case, a relative large number of information sources exist, which is behind the robust functioning of the system. Conversely, communication sinks, where information is lost, are also generated.

On the other hand, our analysis of the community structure reveals some interesting facts. Geo-centered modules are abundant, but ideological or fame-related ones are also remarkable. It is important to keep in mind that time-stamped data is the dynamic (or functional) result of the activity on top of a more stable underlying structure, that of “following” and “followers” in the

social networking site. For this reason, the communication network we observe is ever-changing and we argue that modules within it have not a straightforward interpretation.

Finally, we have studied the patterns by which nodes become increasingly more visible. Results indicate that popularity growth in the context of political conflict does not display significant differences from other, less fashionable examples. Popularity is dominated by a fluctuating behavior, and popularity burst distributions lack of a characteristic scale. This fact connects the dynamics of popularity with other critical phenomena in many natural and artificial systems.

In summary, online social networks and the Web 2.0 provide new challenges to network theory. Events in the real world, ranging from economic phenomena to political protest, stand as the driving forces leading to the emergence of complex, time-evolving communication patterns. In this scenario, network theory stands as a suitable tool to unfold the structural and dynamical facets of such emergent systems.

Acknowledgments

We are indebted to Beatriz Antolí, Guillermo Losilla, Rubén Valles and Isabel Vidal for their help and assistance at several stages of this study.

Author Contributions

Conceived and designed the experiments: JB-H AT YM. Performed the experiments: JB-H AR IG EC AF D. Ferrer D. Francos DI MPP GR F. Sanz F. Serrano CV AT YM. Analyzed the data: JB-H AR AT YM. Contributed reagents/materials/analysis tools: JB-H YM. Wrote the paper: JB-H YM.

References

- Boccaletti S, Latora V, Moreno Y, Chavez M, Hwang D (2006) Complex networks: structure and dynamics. *Phys Rep* 424: 175–308.
- Dorogovtsev SN, Goltsev AV, Mendes JFF (2008) Critical phenomena in complex networks. *Rev Mod Phys* 80: 1275–1336.
- Rybski D, Buldyrev S, Havlin S, Liljeros F, Makse H (2009) Scaling laws of human interaction activity. *Proc Nat Acad Sci* 106: 12640.
- Pastor-Satorras R, Vespignani AA (2001) Epidemic Spreading in Scale-Free Networks. *Phys Rev Lett* 86: 3200–3203.
- Moreno Y, Pastor-Satorras R, Vespignani A (2002) Epidemic outbreaks in complex heterogeneous networks. *Eur Phys J B* 26: 521–529.
- Moreno Y, Nekovee M, Vespignani A (2004) Efficiency and reliability of epidemic data dissemination in complex networks. *Phys Rev E* 69: 055101(R).
- Dunbar RIM (1992) Neocortex size as a constraint on group size in primates. *Journal of Human Evolution* 22(6): 469–493.
- Dunbar RIM (1993) Coevolution of neocortical size, group size and language in humans. *Behavioral and Brain Sciences* 16(4): 681–735.
- Goncalves B, Perra N, Vespignani A (2011) Validation of Dunbar’s number in Twitter conversations. *PLoS ONE*, in press.
- Arenas A, Danon L, Diaz-Guilera A, Gleiser PM, Guimerà R (2004) Community analysis in social networks. *Eur Phys J B* 38: 373–380.
- Lozano S, Arenas A, Sanchez A (2008) Community connectivity and heterogeneity: clues and insights on cooperation on social networks. *Journal of Economic Interaction and Coordination* 3: 183–199.
- Blondel V, Guillaume J, Lambiotte R, Lefebvre E (2008) Fast unfolding of communities in large networks. *Journal of Statistical Mechanics: Theory and Experiment* 2008: P10008.
- Fortunato S (2010) Community detection in graphs. *Physics Reports* 486: 75–174.
- Adamic L, Glance N (2005) The political blogosphere and the 2004 us election: divided they blog. In: *Proceedings of the 3rd International Workshop on Link Discovery ACM*. pp 36–43.
- Conover M, Ratkiewicz J, Francisco M, Gonçalves B, Flammini A, et al. (2011) Political polarization on twitter. *Proc. 5th Intl. Conference on Weblogs and Social Media*. 89 p.
- Rosvall M, Bergstrom C (2008) Maps of random walks on complex networks reveal community structure. *Proceedings of the National Academy of Sciences* 105: 1118.
- Pons P, Latapy M (2005) Computing communities in large networks using random walks. *Lecture notes in computer science* 3733: 284.
- Borge-Holthoefer J, Arenas A (2010) Categorizing words through semantic memory navigation. *Eur Phys J B* 74: 265–270.
- Edler D, Rosvall M. The map generator software package. Available: URL <http://www.mapequation.org>. Accessed 2011 Jun 15.
- Ratkiewicz J, Fortunato S, Flammini A, Menczer F, Vespignani A (2010) Characterizing and modeling the dynamics of online popularity. *Physical Review Letters* 105: 158701.
- Fortunato S, Flammini A, Menczer F (2006) Scale-free network growth by ranking. *Physical Review Letters* 96: 218701.
- Simon H (1971) *Computers, Communication, and the Public Interest*, Baltimore: Johns Hopkins. pp 33–72.