



The weighted Bitcoin Lightning Network

Jian-Hong Lin^{a,b,c}, Emiliano Marchese^d, Claudio J. Tessone^{a,b}, Tiziano Squartini^{d,e,*}

^a Blockchain and Distributed Ledger Technologies, Institute of Informatics, University of Zürich, Andreasstrasse 15, CH-8050 Zürich, Switzerland

^b UZH Blockchain Center, University of Zürich, Andreasstrasse 15, CH-8050 Zürich, Switzerland

^c ETH Zürich, Department of Management, Technology and Economics, Scheuchzerstrasse 7, CH-8092 Zürich, Switzerland

^d IMT School for Advanced Studies Lucca, Piazza San Francesco 19, I-55100 Lucca, Italy

^e Institute for Advanced Study, University of Amsterdam, Oude Turfmarkt 145, NL-1012 GC Amsterdam, The Netherlands

ARTICLE INFO

Keywords:

Blockchain-based systems
Bitcoin Lightning Network
Null models
Centralization
Core-periphery

ABSTRACT

The Bitcoin Lightning Network (BLN) was launched in 2018 to scale up the number of transactions between Bitcoin owners. Although several contributions concerning the analysis of the BLN binary structure have recently appeared in the literature, the properties of its weighted counterpart are still largely unknown. The present contribution aims at filling this gap, by considering the Bitcoin Lightning Network over a period of 18 months, ranging from 12th January 2018 to 17th July 2019, and focusing on its weighted, undirected, daily snapshot representation - each weight representing the total capacity of the channels the two involved nodes have established on a given temporal snapshot. As the study of the BLN weighted structural properties reveals, it is becoming increasingly ‘centralized’ at different levels, just as its binary counterpart: (1) the Nakamoto coefficient shows that the percentage of nodes whose degrees/strengths ‘enclose’ the 51% of the total number of links/total weight is rapidly decreasing; (2) the Gini coefficient confirms that several weighted centrality measures are becoming increasingly unevenly distributed; (3) the weighted BLN topology is becoming increasingly compatible with a core-periphery structure, with the largest nodes ‘by strength’ constituting the core of such a network, whose size keeps shrinking as the BLN evolves. Further inspection of the resilience of the weighted BLN shows that removing such hubs leads to the network fragmentation into many components, an evidence indicating potential security threats — as the ones represented by the so called ‘split attacks’.

1. Introduction

The Bitcoin Lightning Network (BLN) [1] represents an attempt to overcome one of the main limitations of the Bitcoin technological design, i.e. *scalability*: at the moment, only a limited amount of transactions per second, whose number is proportional to the size of blocks and their release frequency, can be processed by Bitcoin, a major shortcoming preventing the adoption of this payment system at a global scale — especially when considering that classic payment mechanisms are able to achieve tens of thousands of transactions per second. Increasing the size of the blocks has been proposed as a solution; implementing this choice, however, would require (1) a larger validation time, (2) a larger storage capability and (3) larger bandwidth costs, hence favoring a more centralized validation process: in fact, fewer entities would become able to validate the new blocks, thus making the system as a whole more prone to faults and attacks.

Developers have tried to break the trade-off between block size and centralization by proposing to process transactions off-chain, i.e. by

means of a ‘Layer 2’ protocol that can operate on top of blockchain-based cryptocurrencies such as Bitcoin: nowadays, such a protocol is known with the name of *Bitcoin Lightning Network* (BLN) and works by creating payment channels across which any two users can exchange money without having the data related to their transactions burdening the entire blockchain [2,3] (see Fig. 1).

The BLN has recently raised a lot of interest: analyses concerning the BLN functioning have found it to allow for a large number of transactions to be processed although rewarding fees are not high enough to promote a large participation in the mining process [4]; Beres et al. [5] analyzed transaction fees and privacy provisions, finding that the participation of most nodes is, in fact, ‘economically irrational’.

On the other hand, analyses concerning the BLN network structure have found it to be characterized by a scale-free topology [6]; besides, Lin et al. [7] and Martinazzi et al. [8] considered the evolution of the former and found it to have become increasingly centralized at different levels. Both features — common to several cryptocurrency

* Corresponding author at: IMT School for Advanced Studies Lucca, Piazza San Francesco 19, I-55100 Lucca, Italy.

E-mail address: tiziano.squartini@imtlucca.it (T. Squartini).

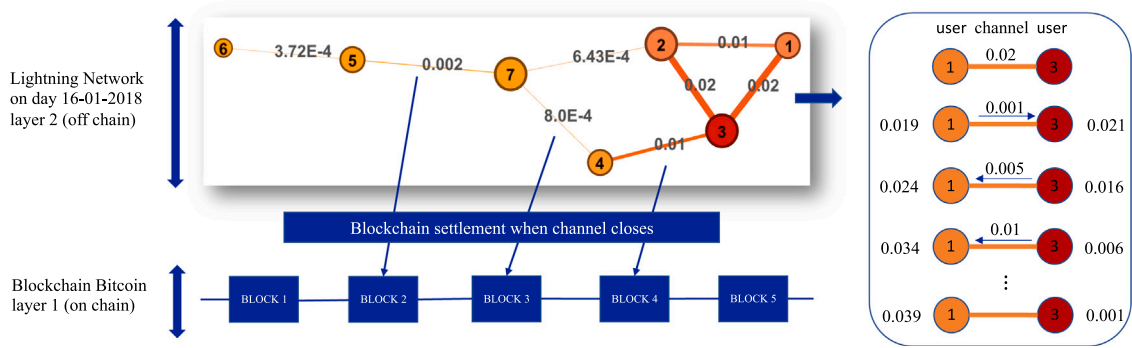


Fig. 1. Pictorial representation of the BLN functioning. The network structure refers to day 16-01-2018, i.e. when the network was constituted by seven nodes (representing users) and eight links (representing open channels). The size of each node is proportional to its degree (i.e. the bigger the node, the larger its degree) while the color of each node is proportional to its strength (i.e. the darker the node, the larger its strength). The width and color of each link are proportional to its weight, i.e. the capacity of the channel it depicts.. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

transaction networks such as Bitcoin, Bitcoin Cash, Dash, Dogecoin, Ethereum, Feathercoin, Litecoin, Monacoin and Z-Cash [9–12] - may have undesirable consequences such as causing a considerable fraction of payments to be easily de-anonymizable [5] and making it prone to channel exhaustion or attacks aimed at isolating nodes (thus, compromising their reachability, the payment success ratio, etc.) [13,14]; a similar conclusion is reached in [15], where the authors analyze the robustness of the BLN against three different types of attacks (locking channels, disconnecting pairs of nodes and isolating hubs) and find it to be disruptable at a relatively low cost; still, Conoscenti et al. [16] have suggested the BLN to be resilient against the removal of nodes that do not have a significant influence on the probability of success of a payment. More on the modeling side, Bartolucci et al. [17] proposed a fitness-based network model for the emergence of the BLN - a proposal that is reminiscent of the ‘fittest-gets-richer’ principle, suggested by Javarone et al. [18], as the mechanism driving the evolution of the global structure of the Bitcoin network and the Bitcoin Cash network.

For a summary of the existing literature on cryptocurrency transaction networks see Wu et al. [19].

2. Data

Most of the aforementioned contributions have just focused on the analysis of the BLN binary structure, leaving its weighted counterpart largely unexplored [7]. With the present paper we aim at filling this gap, by studying the weighted properties of the BLN daily snapshot representation, at the micro-, meso- and macro-scale, across a period of 18 months, i.e. from 12th January 2018 to 17th July 2019.

In order to properly understand the structural details of the aforementioned representation, let us briefly sum up how the BLN works. Payments in the BLN are *source-routed* and *onion-routed* - features that are there to enhance the users privacy: while ‘source-routed’ means that the source node is the one responsible for finding a route connecting itself to the final one, ‘onion-routed’ means that the intermediate nodes only know the ones lying ‘one step-before’ and ‘one step-after’ them, in the payment route.

In order to pre-compute the entire payment route, the sender must have a reasonably up-to-date view of the network topology. Nodes in the BLN regularly broadcast information about the channels they participate in: such a mechanism, called *gossip*, allows other nodes to keep their view of the network topology up-to-date.

The BLN topology can be visualized by means of the so-called *routing table*. For this paper, we took a snapshot of the routing table every 15 min, between January 12th 2018, at blockheight 503.816, to July 17th 2019, at blockheight 585.844 [20]: these snapshots were, then, aggregated into *timespans*, each timespan representing a constant state of a channel from its start to its end; for the present analysis, we considered the *daily snapshot* representation of the BLN, including all

channels that were found to be active during that day. Each weight defining the BLN representation considered in the present work, thus, indicates the total capacity of the channels the two involved nodes have established on a given temporal snapshot — the *channel capacity* coinciding with the total amount of bitcoins needed to keep it open as well as the maximum amount of tokens allowed to be exchanged across it (see Fig. 1).

Importantly, here we do not rest upon estimates of the number of daily blocks – obtainable by considering that the time between the appearance of two subsequent blocks, in the blockchain, is Poisson distributed with an expected value of 10 min – but on the exact time our channels have been opened: since every channel consists of an unspent transaction output on the blockchain, we can determine the size of a channel and its opening and closing time within minutes.

3. Methods

Notation. On a generic, daily snapshot t , the BLN can be described as a weighted, undirected network with total number of nodes $N^{(t)}$ and represented by an $N^{(t)} \times N^{(t)}$ symmetric matrix $\mathbf{W}^{(t)}$ whose generic entry $w_{ij}^{(t)}$ indicates the total capacity of the channels nodes i and j have established during the snapshot t [21,22]. Consistently, the generic entry of the BLN binary adjacency matrix $\mathbf{A}^{(t)}$ reads $a_{ij}^{(t)} = 1$ if $w_{ij}^{(t)} > 0$ and $a_{ij}^{(t)} = 0$ otherwise: the presence of a link between any two nodes i and j , i.e. $a_{ij}^{(t)} = 1$, indicates that one or more payment channels have been opened, between the same nodes, during the snapshot t . As a last remark, we will focus on the largest connected component (LCC) of the BLN, throughout its entire history — the percentage of nodes belonging to it being steadily above 90%.

For the sake of illustration, we will plot our results for four snapshots, i.e. the ones whose LCC is characterized by a number of nodes amounting at 100, 1.000, 3.000, 5.000 and corresponding to the days 24-01-2018, 30-03-2018, 19-12-2018 and 01-03-2019, respectively (see Fig. 2).

Degree and strength distributions. The total number of channels (i.e. *links*) that have been opened during the snapshot t is provided by $L^{(t)} = \sum_{i=1}^{N^{(t)}} \sum_{j=i+1}^{N^{(t)}} a_{ij}^{(t)}$; on the other hand, the total number of channels node i participates in coincides with its *degree*, i.e. $k_i^{(t)} = \sum_{j(\neq i)=1}^{N^{(t)}} a_{ij}^{(t)}$. The weighted counterparts of the notions above coincide with the total weight of the network, i.e. $W^{(t)} = \sum_{i=1}^{N^{(t)}} \sum_{j=i+1}^{N^{(t)}} w_{ij}^{(t)}$, and with the total amount of money exchanged by node i , i.e. $s_i^{(t)} = \sum_{j(\neq i)=1}^{N^{(t)}} w_{ij}^{(t)}$, a quantity often referred to as *node strength* or *node capacity*.

While inspecting the functional form of the degree and strength distributions may reveal the presence of hubs, i.e. ‘large’, single nodes, when dealing with cryptocurrencies it is of interest making a step further and inspecting the presence of ‘large subgraphs’ of nodes. The

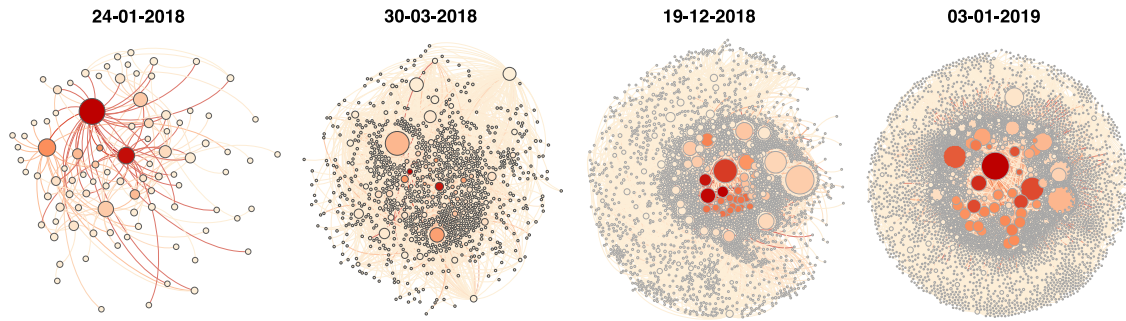


Fig. 2. Pictorial representation of the four snapshots of the BLN whose LCC is characterized by a number of nodes amounting at 100, 1,000, 3,000, 5,000 and corresponding to the days 24-01-2018, 30-03-2018, 19-12-2018 and 01-03-2019, respectively. The size of each node is proportional to its degree (i.e. the bigger the node, the larger its degree) while the color of each node is proportional to its strength (i.e. the darker the node, the larger its strength). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

meaning of this sentence can be made more precise upon considering the metric designed by Srinivasan et al. [23] to measure the number of addresses required (to collude) for gathering over the 51% of the overall mining power and named *Nakamoto index*: a high Nakamoto coefficient indicates that many miners, or mining pools, need to combine their power to reach the 51% threshold needed to take over the blockchain. Here, we adapt it to quantify a ‘topological’ kind of majority, by defining

$$N_k = \min\{i \in [1 \dots N] : \sum_i^i f_i \geq 0.51\} \quad (1)$$

where $f_i = k_i/2L$ and

$$N_s = \min\{i \in [1 \dots N] : \sum_i^i s_i \geq 0.51\} \quad (2)$$

where $f_i = s_i/2W$: the first variant of the Nakamoto index can be calculated by starting from the (node with) largest degree and add them up until the condition above is satisfied; analogously, for the second variant.

Assortativity and hierarchy. In order to gain insight into the higher-order structure of the BLN, we can consider the quantities known as *average nearest neighbors degree*, defined as

$$\text{ANND}_i = \frac{\sum_{j(\neq i)=1}^N a_{ij}k_j}{k_i}, \quad \forall i \quad (3)$$

and *average nearest neighbors strength*, defined as

$$\text{ANNS}_i = \frac{\sum_{j(\neq i)=1}^N a_{ij}s_j}{k_i}, \quad \forall i; \quad (4)$$

while plotting ANND_i versus k_i reveals the (either positive or negative) assortative character of a network, i.e. the presence of (either positive or negative) correlations between degrees, plotting ANNS_i versus k_i reveals the presence of (either positive or negative) correlations between degrees and strengths. On the other hand, the ‘cohesiveness’ of the neighborhood of each node can be inspected by calculating the *binary clustering coefficient*

$$\text{BCC}_i = \frac{\sum_{j(\neq i)=1}^N \sum_{k(\neq i,j)=1}^N a_{ij}a_{jk}a_{ki}}{k_i(k_i - 1)}, \quad \forall i \quad (5)$$

defined as the percentage of triangles established by any two neighbors of each node and the node itself. Its weighted counterpart reads

$$\text{WCC}_i = \frac{\sum_{j(\neq i)=1}^N \sum_{k(\neq i,j)=1}^N w_{ij}w_{jk}w_{ki}}{k_i(k_i - 1)}, \quad \forall i \quad (6)$$

and is intended to assign a ‘weight’ to each triangle counted by the BCC, by weighing the connections shaping it. Plotting BCC_i versus k_i reveals the (possibly) hierarchical character of a network, i.e. its organization in sub-modules; plotting WCC_i versus k_i , instead, provides a hint about

the magnitude of the nodes inter-connections as a function of the nodes connectivity.

Disparity. The *disparity index* is defined as

$$Y_i = \frac{\sum_{j(\neq i)=1}^N \left[\frac{w_{ij}}{s_i} \right]^2}{\sum_{j(\neq i)=1}^N w_{ij}^2} = \frac{\sum_{j(\neq i)=1}^N w_{ij}^2}{s_i^2} = \frac{\sum_{j(\neq i)=1}^N w_{ij}^2}{\left[\sum_{j(\neq i)=1}^N w_{ij} \right]^2}, \quad \forall i \quad (7)$$

and quantifies the (un)evenness of the distribution of the weights ‘constituting’ the i th strength over the k_i links characterizing the connectivity of node i . More specifically, the disparity index of node i reads $Y_i = 1/k_i$ in case weights are equally distributed among the connections established by it, i.e. $w_{ij} = a_{ij}s_i/k_i, \forall j$, any larger value signaling an excess concentration of weight in one or more links.

Centrality. Any index measuring the centrality of a node aims at quantifying its importance in the network, according to some specific topological criterion [24–27]. While the efforts of researchers have mainly focused on the definition of binary centrality measures, relatively little work has been done on their weighted counterparts. In what follows, we will consider possible extensions of the centrality measures employed in [7], i.e. the *degree*, *closeness*, *betweenness* and *eigenvector centrality*:

- the degree centrality [26,27] of node i coincides with its degree, normalized by the maximum attainable value, i.e. $\text{DC}_i = k_i/(N - 1)$: the strength centrality of node i generalizes it by simply replacing the total number of ‘node-specific’ connections with the total ‘node-specific’ weight. In what follows we will consider the (simpler) definition

$$\text{WDC}_i = s_i, \quad \forall i \quad (8)$$

from which it follows that the most central node, according to the strength variant, is the one characterized by the largest percentage of weight ‘embodied’ by (the totality of) its connections;

- the closeness centrality [26,27] of node i is defined as $\text{CC}_i = (N - 1)/\sum_{j(\neq i)=1}^N d_{ij}$ where d_{ij} is the topological distance between nodes i and j , i.e. the length of any shortest path connecting them. The definition of weighted closeness centrality of node i is based on the redefinition of shortest path length which, in turn, rests upon the redefinition of *weighted distance* between any two nodes, i.e. $d_{ij}^{(w)}$. Possible variants of the latter one read $d_{ij}^{(w)} = \min\{w_{ih} + \dots + w_{hj}\}$ and $d_{ij}^{(w)} = \min\left\{\frac{1}{w_{ih}} + \dots + \frac{1}{w_{hj}}\right\}$ where h indexes the intermediary vertices lying on the path between i and j , $w_{ih} \dots w_{hj}$ are the weights of the corresponding edges and the extremum is taken over all paths between i and j . Naturally, the meaning changes along with the chosen definition: while the first one describes any two nodes as ‘closer’, the smaller the weights of

the intermediate connections, the opposite is true when the second one is considered. Hereby, we opt for the following definition of weighted closeness centrality

$$WCC_i = \frac{N-1}{\sum_{j(\neq i)=1}^N d_{ij}^{(w)}}, \quad \forall i \quad (9)$$

with $d_{ij}^{(w)} = \min \left\{ \frac{1}{w_{ih}} + \dots + \frac{1}{w_{hj}} \right\}$. This choice also implies that once the path connecting nodes i and j has been individuated, the WCC is nothing else than the harmonic mean of the weights constituting it;

- the betweenness centrality [26,28–30] of node i is given by $BC_i = \sum_{s(\neq i)=1}^N \sum_{t(\neq i,s)=1}^N \sigma_{st}(i) / \sigma_{st}$ where σ_{st} is the total number of shortest paths between node s and t and $\sigma_{st}(i)$ is the number of shortest paths between nodes s and t that pass through node i . The weighted counterpart of it can be defined as

$$WBC_i = \sum_{s(\neq i)=1}^N \sum_{t(\neq i,s)=1}^N \frac{\sigma_{st}^{(w)}(i)}{\sigma_{st}^{(w)}}, \quad \forall i \quad (10)$$

where, now, $\sigma_{st}^{(w)}$ is the total number of weighted shortest paths between nodes s and t and $\sigma_{st}^{(w)}(i)$ is the number of weighted shortest paths between nodes s and t that pass through node i ;

- the eigenvector centrality [26,30,31] of node i is defined as the i th element of the eigenvector corresponding to the largest eigenvalue of the binary adjacency matrix — whose existence is guaranteed in case the Perron–Frobenius theorem holds true. According to the definition above, a node with large eigenvector centrality is connected to other ‘well connected’ nodes. Such a definition can be extended by considering the WEC_i , defined as the i th element of the eigenvector corresponding to the largest eigenvalue of the weighted adjacency matrix.

The Gini coefficient has been introduced to quantify ‘inequality’ in wealth distribution [32,33] and ranges between 0 and 1, a larger Gini coefficient indicating a larger ‘(un)evenness’ of the income distribution:

$$G_c = \frac{\sum_{i=1}^N \sum_{j=1}^N |c_i - c_j|}{2N \sum_{i=1}^N c_i}; \quad (11)$$

hereby, we apply it to the several definitions of centrality provided above. As a general comment, we would like to stress that a non-normalized centrality measure cannot be employed to compare nodes, across different configurations, in a fully consistent way. However, if our only interest is that of quantifying the (un)evenness of the distribution of our centrality measures, the absence of a normalization term does not make any difference: in fact, the Gini coefficient is not affected by it.

Small-world -ness. The study of the BLN centralization can be approached from a slightly different perspective by asking if the BLN is (increasingly) becoming a small-world system [34–36]. The usual way of proceeding to answer such a question prescribes to check if

$$\bar{d} = \frac{\sum_{i=1}^N \sum_{j(\neq i)=1}^N d_{ij}}{N(N-1)} \sim \ln N \quad (12)$$

i.e. if the average path length grows logarithmically with the number of nodes and if the average clustering coefficient $BCC_i = \sum_{j=1}^N BCC_{ij} / N$ is larger than the one predicted by an *Undirected Random Graph Model* (URGM) tuned to reproduce the empirical density of links. Recently, however, it has been argued that the same question can be answered by considering the quantity named *global efficiency*, defined as

$$E_g = \frac{\sum_{i=1}^N \sum_{j(\neq i)=1}^N d_{ij}^{-1}}{N(N-1)}, \quad (13)$$

understood as an indicator of the ‘traffic capacity’ of a network and, quite remarkably, not affected by the analytical problems suffered by

the average path length [35] - potentially diverging due to the presence of couples of nodes belonging to disconnected components. Latora et al. [35] have also defined a *local efficiency* as

$$E_l = \frac{1}{N} \sum_{i=1}^N E(\mathbf{G}_i), \quad (14)$$

a quantity that can be evaluated by, first, calculating the efficiency of the subgraph induced by the nearest neighbors of each node, upon removing it and, then, averaging such numbers. Latora et al. [35] have argued that while E_g plays a role analogous to the inverse of the average path length, E_l plays a role analogous to the average clustering coefficient: hence, small-world networks should have both a large E_g and a large E_l , i.e. should be very efficient in allowing nodes to communicate in both a global and a local fashion.

Core-periphery detection. As it has emerged quite clearly from the binary analysis of the BLN, just inspecting the evolution of centrality measures can return a too simplistic picture of the network under consideration. For this reason, we have checked for the presence of mesoscopic ‘centralized’ structures such as the *core-periphery* one [37–41], composed by a densely-connected subgraph of nodes surrounded by a periphery of loosely-connected vertices. In order to do so, we have implemented the approach recently proposed in [42] and prescribing to minimize the score function

$$\mathcal{W}_{//} = \sum_{w_* \geq w_*^*} \sum_{w_o \geq w_o^*} \frac{\binom{V_c}{w_*} \binom{V_o}{w_o} \binom{V-V_c-V_o}{W-w_*-w_o}}{\binom{V}{W}} \quad (15)$$

known as *bimodular surprise*¹; here, $V = N(N-1)/2$ is the total number of node pairs, $W = \sum_{i=1}^N \sum_{j=i+1}^N w_{ij}$ is the total weight of the network, V_c is the number of node pairs in the core portion of the network, V_o is the number of node pairs in the periphery portion of the network, w_*^* is the observed number of core links and w_o^* is the observed number of periphery links. From a technical point of view, $\mathcal{W}_{//}$ is the p -value of a multivariate negative hypergeometric distribution and the multiset notation, according to which $\binom{V_c}{w_*} = \binom{V_c+w_o-1}{w_*}$ allows $\mathcal{W}_{//}$ to be compactly rewritten in a way that nicely mirrors that of its binary counterpart [42,43].

Benchmarking the observations. Along the guidelines of the analysis carried out in [7], in what follows we benchmark our observations by employing the recently-proposed null model called *CREMA* - the acronym standing for *Conditional Reconstruction Model A* [44,45] - that allows binary and weighted constraints to be defined in a disentangled fashion. From a purely theoretical point of view, it is defined by the maximization of the *conditional Shannon entropy*

$$S(\mathcal{W}|\mathcal{A}) = - \sum_{\mathbf{A} \in \mathcal{A}} P(\mathbf{A}) \int_{\mathbb{W}_{\mathbf{A}}} Q(\mathbf{W}|\mathbf{A}) \ln Q(\mathbf{W}|\mathbf{A}) d\mathbf{W} \quad (16)$$

constrained to reproduce the strengths $\{s_i\}_{i=1}^N$; the (conditional) weighted distribution output by such an optimization procedure reads

$$Q(\mathbf{W}|\mathbf{A}) = \frac{e^{-H(\mathbf{W})}}{Z_{\mathbf{A}}} = \prod_{i=1}^N \prod_{j=i+1}^N q_{ij}(w_{ij}|a_{ij}) = \prod_{i=1}^N \prod_{j=i+1}^N (\beta_i + \beta_j)^{a_{ij}} e^{-(\beta_i + \beta_j)w_{ij}}; \quad (17)$$

notice the conditional character of the distribution above, embodied by the term a_{ij} at the exponent — as a simple consistency check, the probability that $w_{ij} = 0$ in case there is no link is $q(w_{ij} = 0|a_{ij} = 0) = 1$ as it should be. The vector of parameters $\{\beta_i\}_{i=1}^N$ defining

¹ The Python package for surprise optimization, called ‘SurpriseMeMore’, is freely downloadable at the following URL: <https://github.com/EmilianoMarchese/SurpriseMeMore>.

the distribution above can be estimated via a (generalized) *likelihood maximization* procedure [44] that leads to the system of N equations

$$s_i = \sum_{j(\neq i)=1}^N \langle w_{ij} \rangle = \sum_{j(\neq i)=1}^N \frac{p_{ij}}{\beta_i + \beta_j}, \quad \forall i; \quad (18)$$

the coefficients $\{p_{ij}\}_{i,j=1}^N$, instead, are treated as ‘prior information’ and, as such, left ‘untouched’ by the estimation procedure above. In a sense, we are free to combine the (conditional) weighted distribution above with the purely binary probability mass function ‘best’ encoding the available information about the network structure. In what follows, we have considered

- the one defining the *Undirected Binary Configuration Model* (UBCM) and following from the maximization of the traditional Shannon entropy $S = -\sum_{\mathbf{A}} P(\mathbf{A}) \ln P(\mathbf{A})$ constrained to reproduce the degrees $\{k_i\}_{i=1}^N$: the UBCM captures the idea that the probability for any two nodes to establish a connection (solely) depends on their degrees and can be fully determined by solving the N equations

$$k_i = \sum_{j(\neq i)=1}^N p_{ij}^{\text{UBCM}} = \sum_{j(\neq i)=1}^N \frac{x_i x_j}{1 + x_i x_j}, \quad \forall i; \quad (19)$$

- the deterministic recipe $p_{ij} \equiv a_{ij}, \forall i < j$, accounting for the case in which the prior knowledge concerns the entire network topological structure, now treated as given.

While, in the first case, the generic set of coefficients $\{p_{ij}\}_{i,j=1}^N$ is instantiated upon identifying $p_{ij} \equiv p_{ij}^{\text{UBCM}}, \forall i < j$, in the second one the, identification simply reads $p_{ij} \equiv a_{ij}, \forall i < j$; in both cases, the resolution of the related system of equations, carried out via the Python package called ‘NEMTROPY²’, leads us to numerically determine the corresponding vector of parameters $\{\beta_i\}_{i=1}^N$.

Benchmarking a set of observations ultimately boils down at verifying their ‘compatibility’ with the predictions output by a chosen null model, by testing their statistical significance against the null model itself. To this aim, one can proceed as follows: first, sampling the ensemble induced by the chosen null model, by generating a ‘sufficiently large’ number of configurations (in all our experiments, 100); second, calculating the value of any quantity of interest over each configuration; third, deriving the corresponding ensemble CDF. At this point, a p -value remains naturally defined; in what follows, we will employ it to carry out one-tailed tests. Whenever tests of this kind are considered, one may be interested in calculating either the (ensemble) probability $Q(X \geq X^*)$ of observing a value, for the quantity of interest X , that is *larger* than the empirical one, X^* , or the (ensemble) probability $Q(X \leq X^*)$ of observing a value, for the quantity of interest X , that is *smaller* than it; in both cases, if such a probability is found to be smaller than a given threshold, the quantity is deemed as statistically significant, hence incompatible with the description provided by the chosen null model — which (significantly) underestimates or overestimates it, respectively.

4. Results

Degree and strength distributions. Giving a look at the four snapshots depicted in Fig. 2 reveals the presence of a large heterogeneity, with nodes having a large degree/strength co-existing with nodes having a small degree/strength; moreover, while nodes with a large degree also have a large strength (i.e. larger nodes are also darker), small, dark nodes can be observed as well: in other words, an overall positive correlation between degrees and strengths co-exists with a large variability of the

² The acronym stands for ‘Network Entropy Maximization: a Toolbox Running On Python’ and the package is freely downloadable at the following URL: <https://pypi.org/project/NEMTropy/>.

strength values — especially for what concerns the nodes with a small connectivity (see Fig. 3).

As a first empirical analysis, we have inspected the functional form of the degree distribution for four distinct snapshots, i.e. the days 24-01-2018, 30-03-2018, 19-12-2018 and 01-03-2019; to this aim, we have plotted the cumulative density function (CDF), defined as $\text{CDF}(k) = \sum_{h \geq k} f(h)$ where $f(h)$ is the fraction of nodes whose degree is h . As shown in Fig. 4, the degree distribution becomes broader as the BLN evolves; moreover, running the code released by Clauset et al. [46] to fit the functional form $\text{PDF}(k) = (\alpha - 1) k_{\min}^{\alpha-1} k^{-\alpha}$ on the data returns the values $\alpha = 1.9, 2.0, 2.1, 2.2$ and $k_{\min} = 1, 3, 14, 26$ while the Kolmogorov–Smirnov test returns the p -values $p = 0.02, 0.03, 0.04, 0.5$. Hence, the null hypothesis that the degrees are distributed according to a power-law is never rejected, at the 1% significance level — while it is, for the first three snapshots, at the 5% significance level. Overall, the null hypothesis that the degrees are distributed according to a power law is not rejected for the 85% of the total number of snapshots, at the 1% significance level, and for the 71% of the total number of snapshots, at the 5% significance level.

As a second empirical analysis, we have calculated the evolution of the CDF of the weights, defined as $\text{CDF}(w) = \sum_{v \geq w} f(v)$. Analogously to what has been observed for the degrees, even the support of the weight distribution has broadened throughout the entire BLN history (see Fig. 5), although to a lesser extent. Fitting a log-normal distribution, whose functional form reads $\text{PDF}(w) = (w\sigma\sqrt{2\pi})^{-1} e^{-\frac{(\ln w - \mu)^2}{2\sigma^2}}$, on the data reveals that, at both the 1% and the 5% significance levels, the Kolmogorov–Smirnov test does not reject the hypothesis that weights are log-normally distributed when $N < 94$ (i.e. from the fourth day to the twelfth day). For our four snapshots, instead, the hypothesis is rejected — notice that day 24-01-2018 is the thirteenth.

As a third empirical analysis, we have considered the evolution of the CDF of the strengths, defined as $\text{CDF}(s) = \sum_{t \geq s} f(s)$. The support of the distribution is enlarged of a few orders of magnitude during the BLN history (see Fig. 6). Analogously to the case of the weights, we have fitted a log-normal distribution, whose functional form reads $\text{PDF}(s) = (s\sigma\sqrt{2\pi})^{-1} e^{-\frac{(\ln s - \mu)^2}{2\sigma^2}}$, on the data: while the Kolmogorov–Smirnov test returns the p -values $p = 0.061, 0.006, 4.4 \cdot 10^{-7}$ and $1.6 \cdot 10^{-7}$, hence does not reject the hypothesis that strengths are log-normally distributed, at both significance levels, for the first snapshot considered here, it does so for the other three ones — an evidence seemingly indicating that, quite early in its history, the BLN has started deviating more and more from the picture provided by the distribution tested here. Overall, the null hypothesis that the strengths are distributed according to a log-normal is not rejected for the 16% of the total number of snapshots, at the 1% significance level, and for the 5% of the total number of snapshots, at the 5% significance level.

The picture provided by the three distributions above can be complemented by the information provided by the Nakamoto index (see Fig. 7). As its evolution clearly shows, the percentage of nodes ‘providing’ the 51% of the total number of links/the total weight progressively reduces, as the BLN size enlarges: in particular, the total weight seems to be distributed less evenly than the total number of connections — as fewer nodes are needed to embody the (same) required percentage. This seems to confirm the appearance of nodes constituting the aforementioned ‘topological’ majority.

Assortativity and hierarchy. Plotting the values of the average nearest neighbors degree versus the degrees reveals the disassortative character of the BLN, i.e. the presence of negative correlations between the degrees: in other words, nodes with a large degree are (preferentially) connected to nodes with a small degree and viceversa. To be noticed that the UBCM-induced CREM_A model successfully captures such a trend, indicating that the information encoded into the degree sequence, leading to

$$\langle \text{ANND}_i \rangle \simeq \frac{\sum_{j(\neq i)=1}^N \langle a_{ij} \rangle \langle k_j \rangle}{\langle k_i \rangle} = \frac{\sum_{j(\neq i)=1}^N p_{ij} k_j}{k_i}, \quad \forall i \quad (20)$$




Fig. 3. Scattering the strength sequence versus the degree sequence reveals the presence of positive correlations between the two sets of quantities: the Pearson coefficient describing them, on our usual four snapshots, amounts at $r = 0.84, 0.42, 0.66, 0.80$, respectively.




Fig. 4. Cumulative density function of the degrees, for our usual four snapshots. The support of the distribution has become broader as the BLN has evolved. Fitting a power-law PDF(k) = $(\alpha - 1)k_{\min}^{\alpha-1}k^{-\alpha}$ on the data, by running the code released by Clauset et al. [46], returns values amounting at $\alpha = 1.9, 2.0, 2.1, 2.2$ and $k_{\min} = 1, 3, 14, 26$ (see the dashed lines) while the Kolmogorov–Smirnov test returns the p-values $p = 0.02, 0.03, 0.04, 0.5$. Hence, the null hypothesis that the degrees are distributed according to a power-law is never rejected, at the 1% significance level — while it is, for the first three snapshots, at the 5% significance level. Overall, the null hypothesis that the degrees are distributed according to a power-law is not rejected for the 85% of the total number of snapshots, at the 1% significance level, and for the 71% of the total number of snapshots, at the 5% significance level.




Fig. 5. Cumulative density function of the weights, for our usual four snapshots. The support of the distribution has become slightly broader as the BLN has evolved. Fitting a log-normal distribution PDF(w) = $(w\sigma\sqrt{2\pi})^{-1}e^{-\frac{(\ln w - \mu)^2}{2\sigma^2}}$ on the data reveals that, at both the 1% and the 5% significance levels, the Kolmogorov–Smirnov test does not reject the hypothesis that weights are log-normally distributed when $N < 94$ (i.e. from the fourth day to the twelfth day); for our four snapshots, instead, the hypothesis is rejected — notice that day 24-01-2018 is the thirteenth. Each dashed line represents the best fit for the specific snapshot: the corresponding parameters read $\mu = -6.6, -8.1, -6.1, -5.5$ and $\sigma = 1.7, 2.1, 2.3, 2.0$.




Fig. 6. Cumulative density function of the strengths, for our usual four snapshots. The support of the distribution is enlarged of a few orders of magnitude during the BLN history. A log-normal distribution PDF(s) = $(s\sigma\sqrt{2\pi})^{-1}e^{-\frac{(\ln s - \mu)^2}{2\sigma^2}}$, fitted on the data, lets the Kolmogorov–Smirnov test returns the p-values $p = 0.061, 0.006, 4.4 \cdot 10^{-7}$ and $1.6 \cdot 10^{-7}$. Hence, the null hypothesis that strengths are log-normally distributed is not rejected for the first snapshot while it is for the last three ones — at both significance levels. Overall, the null hypothesis that the strengths are distributed according to a log-normal is not rejected for the 16% of the total number of snapshots, at the 1% significance level, and for the 5% of the total number of snapshots, at the 5% significance level. Each dashed line represents the best fit for the specific snapshot: the corresponding parameters read $\mu = -6.0, -6.7, -4.7, -4.0$ and $\sigma = 2.0, 2.3, 3.0, 2.8$.




Fig. 7. Evolution of the Nakamoto index for the degrees and the strengths, plotted versus the total number of nodes: as the size of the system enlarges, the percentage of nodes ‘providing’ the 51% of the total number of links/the total weight progressively reduces, an evidence pointing out that nodes embodying a ‘topological’ kind of majority indeed appear. Moreover, the total weight seems to be distributed less evenly than the total number of connections.

with $p_{ij} \equiv p_{ij}^{\text{UBCM}}$, $\forall i < j$ and where the symbol \simeq indicates that we have approximated the expected value of a ratio as the ratio of the expected values, is enough to account for the correlations between the degrees as well. An analogous decreasing trend characterizes the values of the average nearest neighbors strength when plotted versus the degrees, i.e. nodes with a large degree are (preferentially) connected to nodes with a small strength and viceversa; as for its binary counterpart, the UBCM-induced CReM_A model successfully reproduces the empirical ANNS values, indicating that the information encoded into the degree and the strength sequences, leading to

$$\langle \text{ANNS}_i \rangle \simeq \frac{\sum_{j(\neq i)=1}^N \langle a_{ij} \rangle \langle s_j \rangle}{\langle k_i \rangle} = \frac{\sum_{j(\neq i)=1}^N p_{ij} s_j}{k_i}, \quad \forall i \quad (21)$$

with $p_{ij} \equiv p_{ij}^{\text{UBCM}}$, $\forall i < j$ successfully accounts for the correlations between the degrees and the strengths as well (see Fig. 8). On the other hand, plotting the values of the clustering coefficient versus the degrees reveals the hierarchical character of the BLN: nodes with a larger degree tend to participate into a smaller number of connected triples than nodes with a smaller degree and viceversa; the UBCM-induced CReM_A model, leading to

$$\begin{aligned} \langle \text{BCC}_i \rangle &\simeq \frac{\sum_{j(\neq i)=1}^N \sum_{k(\neq i,j)=1}^N \langle a_{ij} \rangle \langle a_{jk} \rangle \langle a_{ki} \rangle}{\sum_{j(\neq i)=1}^N \sum_{k(\neq i,j)=1}^N \langle a_{ij} \rangle \langle a_{ik} \rangle} \\ &= \frac{\sum_{j(\neq i)=1}^N \sum_{k(\neq i,j)=1}^N p_{ij} p_{jk} p_{ki}}{\sum_{j(\neq i)=1}^N \sum_{k(\neq i,j)=1}^N p_{ij} p_{ik}}, \quad \forall i \end{aligned} \quad (22)$$

with $p_{ij} \equiv p_{ij}^{\text{UBCM}}$, $\forall i < j$ is able to capture such a trend as well. The same decreasing trend, instead, does not characterize the values of the weighted clustering coefficient when plotted versus the degrees which, instead, appears as rather flat — interestingly, this is no longer true when the weighted clustering coefficient values are plotted versus the strengths: in this case, a clear rising trend is visible, signaling that nodes with a larger strength tend to participate into ‘heavier’ connected triples of nodes. Again, the UBCM-induced CReM_A model, predicting

$$\begin{aligned} \langle \text{WCC}_i \rangle &\simeq \frac{\sum_{j(\neq i)=1}^N \sum_{k(\neq i,j)=1}^N \langle w_{ij} \rangle \langle w_{jk} \rangle \langle w_{ki} \rangle}{\sum_{j(\neq i)=1}^N \sum_{k(\neq i,j)=1}^N \langle a_{ij} \rangle \langle a_{ik} \rangle} \\ &= \frac{\sum_{j(\neq i)=1}^N \sum_{k(\neq i,j)=1}^N \langle w_{ij} \rangle \langle w_{jk} \rangle \langle w_{ki} \rangle}{\sum_{j(\neq i)=1}^N \sum_{k(\neq i,j)=1}^N p_{ij} p_{ik}}, \quad \forall i \end{aligned} \quad (23)$$

with $p_{ij} \equiv p_{ij}^{\text{UBCM}}$, $\forall i < j$ successfully reproduces the empirical WCC values, indicating that the information encoded into the degree and the strength sequences successfully accounts for the behavior of third-order properties as well (see Fig. 9).

Disparity. As anticipated in the paragraph introducing such a quantity, the disparity index of node i reads $Y_i = 1/k_i$ in case weights are equally distributed among the neighbors of node i . Fig. 10 shows the scatter plot of Y_i as a function of k_i (since it is plotted in a log–log scale, the function $y = -x$ becomes the trend signaling that weights are uniformly distributed among the neighbors of each node): generally speaking, many values lie above the $y = -x$ line, an evidence indicating that some kind of ‘excess concentration’ of weight (in one or more links) is indeed present - a tendency which is particularly evident for nodes with smaller degree.

Let us now compare the empirical disparity values with the predictions of the null models defined within our CReM_A framework. To this aim, let us explicitly calculate the expected value of disparity, that reads

$$\begin{aligned} \langle Y_i \rangle &\simeq \frac{\sum_{j(\neq i)=1}^N \langle w_{ij}^2 \rangle}{\langle s_i^2 \rangle} = \frac{\sum_{j(\neq i)=1}^N (\text{Var}[w_{ij}] + \langle w_{ij} \rangle^2)}{\text{Var}[s_i] + \langle s_i \rangle^2} \\ &= \frac{\text{Var}[s_i] + \sum_{j(\neq i)=1}^N \langle w_{ij} \rangle^2}{\text{Var}[s_i] + s_i^2}, \quad \forall i \end{aligned} \quad (24)$$

where

$$\langle w_{ij} \rangle = \frac{p_{ij}}{\beta_i + \beta_j}, \quad \forall i < j \quad (25)$$

and

$$\text{Var}[s_i] = \sum_{j(\neq i)=1}^N \text{Var}[w_{ij}] = \sum_{j(\neq i)=1}^N \frac{p_{ij}}{(\beta_i + \beta_j)^2}, \quad \forall i \quad (26)$$

(naturally, for the present analysis we have considered both the case $p_{ij} \equiv p_{ij}^{\text{UBCM}}$, $\forall i < j$ and the case $p_{ij} \equiv a_{ij}$, $\forall i < j$). As Fig. 10 shows, disparity is, generally speaking, in agreement with our benchmark models. However, the calculation of the percentage of nodes for which $Q(Y_i \geq Y_i^*) < 0.05$, for our usual four snapshots, reveals it to be 0%, 3.0%, 9.1%, 11% for the UBCM-induced CReM_A model and 0%, 5.6%, 15%, 17% for the deterministic CReM_A model: in other words, the percentage of nodes whose empirical disparity is significantly larger than predicted by one of our two null models is rising throughout the entire BLN history. This evidence suggests that, as the BLN evolves, its vertices treat their neighbors less and less equally: indeed, they seem to place weights in a way that increasingly ‘favors’ some of the links surrounding them - a result that remains true even when a null model constraining the entire topology of the BLN is employed.³

Centrality. Let us now comment the results concerning the weighted centrality measures considered in the present work. As a general observation, the weighted cases are characterized by trends which are overall similar to the trends characterizing the binary cases. As already observed for the purely binary BLN structure, the evolution of the Gini index for most centrality measures points out the latter ones to grow (strongly) unevenly distributed throughout the entire BLN history. While the rise of the Gini coefficient for the weighted degree, betweenness and eigenvector centrality measures suggests the appearance of nodes with ‘heavy’ connections - further confirmed by the strength distribution, which is a fat-tailed one - likely crossed by many paths and well connected between themselves, the flat trend characterizing the evolution of the closeness centrality confirms what has been already observed in the purely binary case, i.e. that the aforementioned ‘hubs’ ensure the vast majority of nodes to be reachable (hence, to be close to each other) quite easily.

Let us now compare the empirical trends of our four centrality measures with the ones predicted by our two null models. Fig. 11 reveals

³ To be noticed that our null models also underestimate disparity values: however, the corresponding percentages (amounting at 3.2%, 4.3%, 5.8%, 5.2% for the UBCM-induced CReM_A model and at 12.8%, 10.8%, 13.0%, 13.6% for the deterministic CReM_A model, for our usual four snapshots), are typically lower and not increasing.

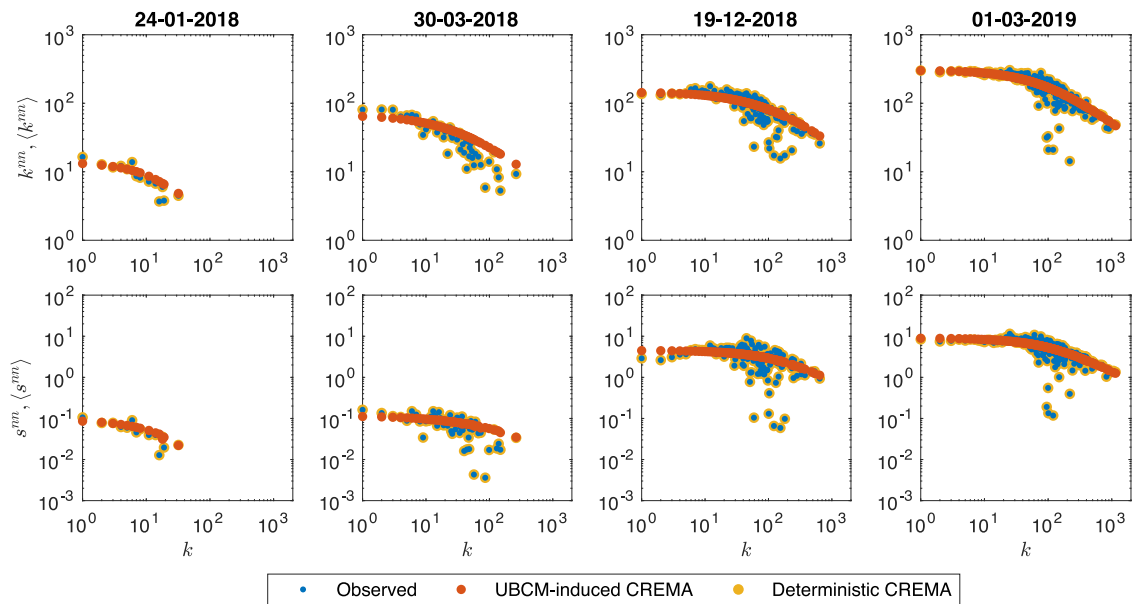


Fig. 8. $\langle ANND_i \rangle$, $\langle ANND_i \rangle$ values scattered versus k_i (upper panels) and $\langle ANNS_i \rangle$, $\langle ANNS_i \rangle$ values scattered versus k_i (bottom panels) for our usual four snapshots (all trends are averaged over the classes of nodes with the same degree). Both trends clearly signal a disassortative behavior, i.e. nodes with a large degree are (preferentially) connected to nodes with a small degree/small strength and viceversa. While the UBCM-induced CREMA model successfully captures such a disassortative trend, the deterministic CREMA model reproduces both the ANND and the ANNS values exactly.




Fig. 9. $\langle BCC_i \rangle$, $\langle BCC_i \rangle$ values scattered versus k_i (upper panels) and $\langle WCC_i \rangle$, $\langle WCC_i \rangle$ values scattered versus k_i (bottom panels) for our usual four snapshots (all trends are averaged over the classes of nodes with the same degree). While the trend of the BCC clearly signals a hierarchical behavior, i.e. the tendency of nodes with a larger degree to participate into a smaller number of connected triples than nodes with a smaller degree and viceversa, this does not seem to be the case for the WCC values when plotted versus the degrees. While the UBCM-induced CREMA model successfully captures both trends, the deterministic CREMA model reproduces only the BCC values exactly.

that the UBCM-induced CREMA model tends to overestimate the values of the Gini index for the weighted degree and closeness centrality, i.e. the empirical weighted degree and closeness centrality measures are always significantly lower than their predicted counterparts. For what concerns the weighted betweenness centrality, instead, the percentage of snapshots for which $Q(G_{WBC} \geq G_{WBC}^*) < 0.05$ amounts at $\approx 87\%$, i.e. the UBCM-induced CREMA model significantly underestimates the weighted betweenness centrality for $\approx 87\%$ of the total number of snapshots. Analogously, the same null model tends to underestimate the values of the Gini index for the weighted eigenvector centrality roughly one third of the times: in fact, the percentage of snapshots for which $Q(G_{WEC} \geq G_{WEC}^*) < 0.05$ amounts at $\approx 33\%$. Interestingly, the empirical

WBC and WEC values are compatible with the predictions output by the UBCM-induced CREMA model, on the ‘remaining’ snapshots.

The deterministic CREMA model, instead, performs slightly better in reproducing the centrality patterns characterizing the BLN: in fact, while it still overestimates the Gini index for the weighted degree centrality, the percentage of snapshots for which $Q(G_{WCC} \geq G_{WCC}^*) < 0.05$ amounts at $\approx 96\%$, i.e. the deterministic CREMA model significantly underestimates the weighted closeness centrality for $\approx 96\%$ of the total number of snapshots. For what concerns the weighted betweenness centrality, the percentage of snapshots for which $Q(G_{WBC} \geq G_{WBC}^*) < 0.05$ amounts at $\approx 50\%$, i.e. the deterministic CREMA model significantly underestimates the weighted betweenness centrality roughly half of the

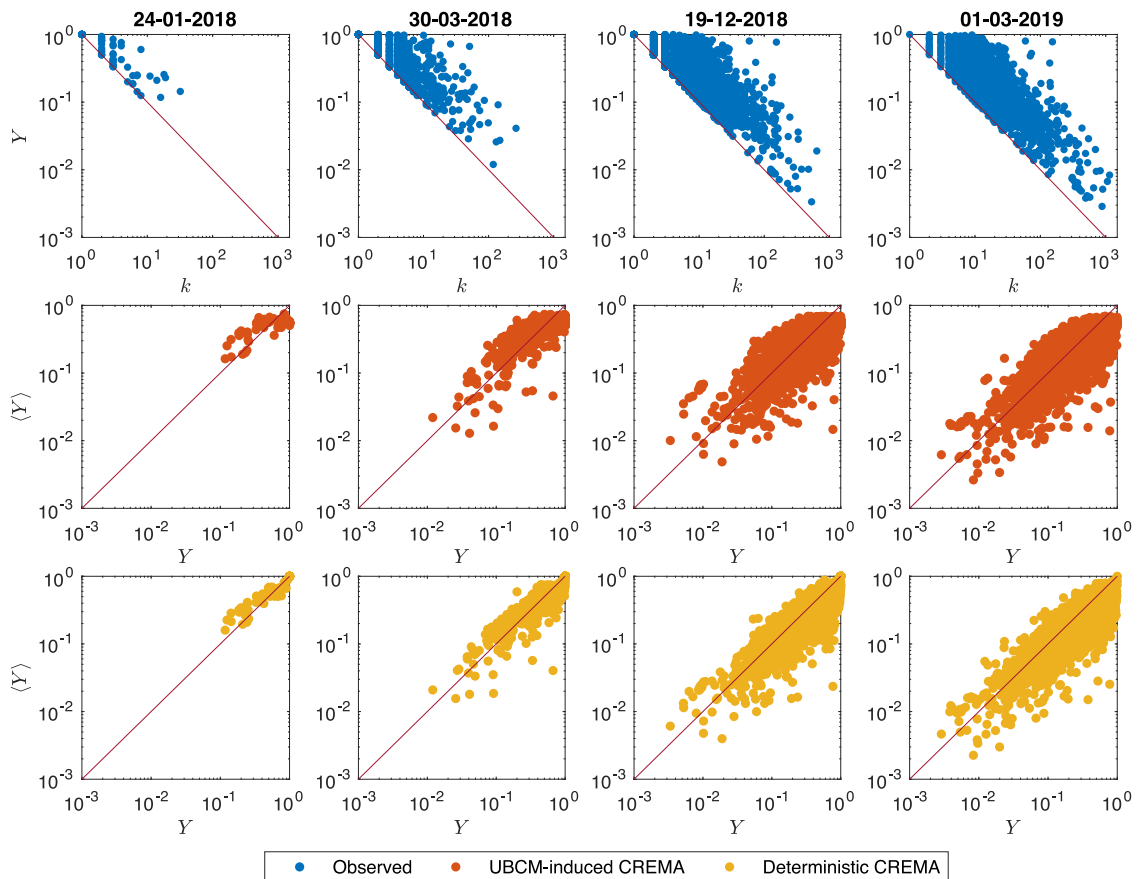


Fig. 10. Upper panels: empirical disparity values scattered versus the degrees, for our usual four snapshots. As the plots reveal, the vast majority of strength values is not evenly distributed across the connections characterizing each node, i.e. $Y_i > 1/k_i$ for the vast majority of nodes. Middle panels: expected disparity values output by the UBCM-induced $CREMA$ model scattered versus the empirical disparity values. Bottom panels: expected disparity values output by the deterministic $CREMA$ model scattered versus the empirical disparity values. The empirical disparity values are, generally speaking, in agreement with our benchmark models; however, the percentage of nodes for which $Q(Y_i \geq Y_i^*) < 0.05$, for our usual four snapshots, amounts at 0%, 3.0%, 9.1%, 11% for the UBCM-induced $CREMA$ model and at 0%, 5.6%, 15%, 17% for the deterministic $CREMA$ model: in other words, the percentage of nodes whose empirical disparity is significantly larger than predicted by one of the two null models considered here is rising throughout the entire BLN history - i.e. its vertices increasingly 'favor' some of the links surrounding them.

times. Lastly, for what concerns the weighted eigenvector centrality, the percentage of snapshots for which $Q(G_{WEC} \leq G_{WEC}^*) < 0.05$ amounts at $\approx 83\%$, i.e. the deterministic $CREMA$ model significantly overestimates the weighted eigenvector centrality for $\approx 83\%$ of the total number of snapshots. The deterministic $CREMA$ model distributes weights more evenly than observed, hence underestimating disparity and letting the strength distribution become wider: this has an interesting consequence, i.e. letting the size of the core under this model become larger than observed - likely, because nodes with relatively low strength become, now, part of the core - while still allowing the unevenness of the WEC distribution rise.

Overall, these results point out a behavior that is not reproducible by just enforcing the degree and the strength sequences - irrespectively from the chosen index: in particular, the behavior of the weighted betweenness centrality points out that both null models - even if to a different extent - predict a more-even-than-observed structure.

Small-world -ness. The evidence that the BLN structure is more-centralized-than-expected rises an interesting question, i.e. if the BLN is small-world or not. From a purely empirical perspective, answering this question amounts at checking the behavior of the average path length, \bar{d} , and that of the average clustering coefficient, $\overline{BCC} = \sum_i BCC_i/N$ [34-36].

Fig. 12 shows the results of these two analyses: while the evolution of \bar{d} is described quite accurately by the function $\ln N$ during the first snapshots of the BLN history, its trend has progressively become more and more similar to the smoothest one characterizing the function

$\ln \ln N$ - which has reached the value ≈ 3.5 on the snapshot with 10^4 nodes. For what concerns the average clustering coefficient, one needs to compare it with the value predicted by the URG, i.e. the null model prescribing to link each pair of nodes with the same probability $p = 2L/N(N-1)$: as Fig. 12 shows, the URG significantly underestimates the average clustering coefficient throughout the entire BLN history; taken together, these results indicate that the BLN is indeed small-world. On the other hand, the UBCM overestimates $BCC = \sum_i BCC_i/N$ during the first half of its history (for $\approx 40\%$ of the total number of snapshots), thus signaling a tendency of our system to avoid closing paths among triples of nodes.

An alternative way of testing small-world -ness is that of checking the behavior of efficiency. Overall, the global efficiency amounts at $E_g \approx 0.4$ and it is significantly underestimated by the UBCM throughout the entire history of the BLN. This indicates that the BLN exchanges information more-efficiently-than-predicted by a null model retaining only the information provided by degrees and can be a consequence of the presence of hubs crossed by many paths that shorten the topological distance between (any pair of) nodes.

These results suggest that the BLN has progressively self-organized to 'keep the overall distances low'. What about efficiency from a local point of view? For what concerns the local efficiency, the percentage of nodes for which $Q(E_i \geq E_i^*) < 0.05$ amounts at 75%: hence, the UBCM significantly underestimates it for a large portion of the BLN snapshots - as evident from Fig. 12, the most recent ones. As the local efficiency $E(G_i)$ provides information about how efficient

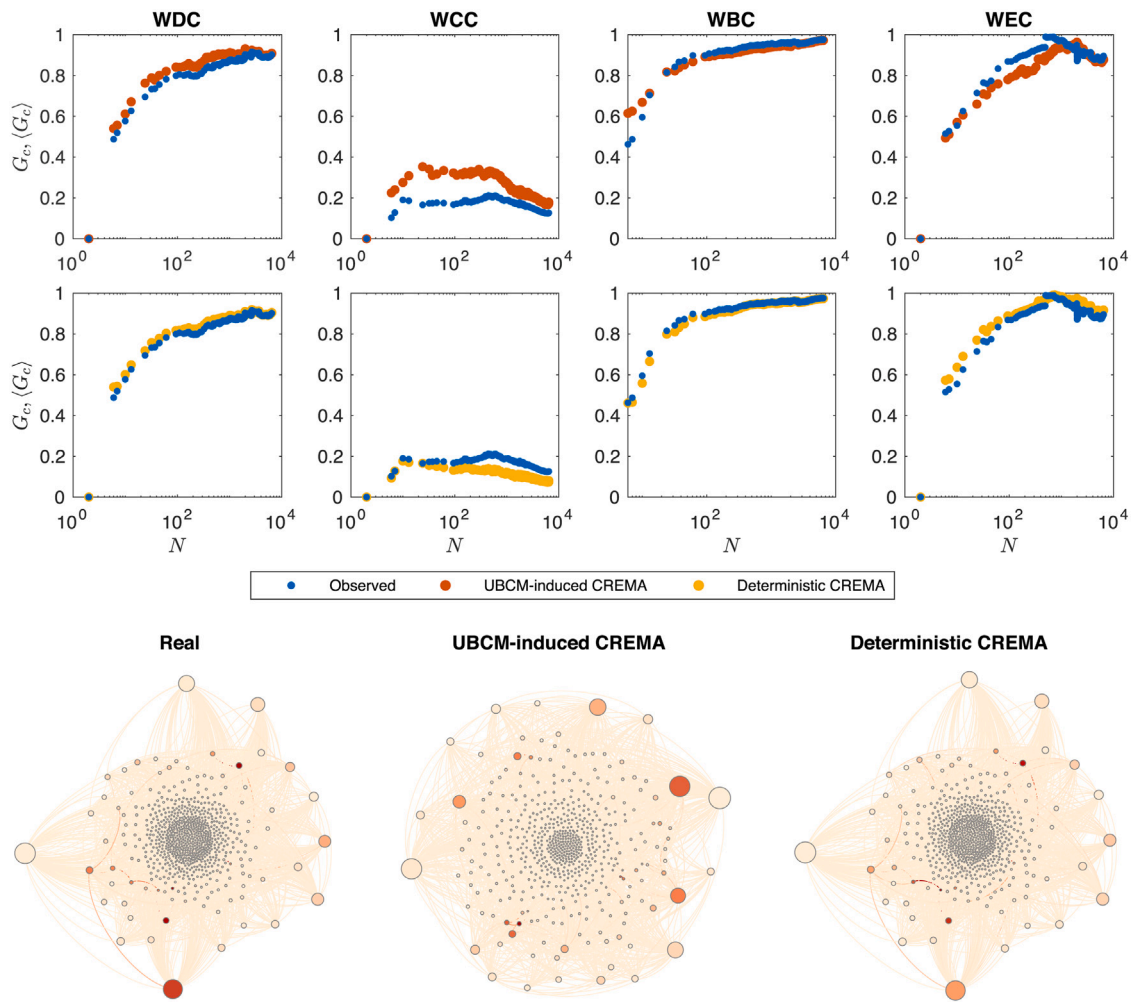


Fig. 11. Upper and middle panels: evolution of the Gini coefficient of the empirical weighted centrality values (blue dots) and of its expected counterpart, under the UBCM-induced CREMA model (red dots, upper panels) and the deterministic CREMA model (yellow dots, bottom panels). The rise of the Gini coefficient for the weighted degree, betweenness and eigenvector centrality points out that the distribution of centrality measures becomes increasingly uneven while the flat trend characterizing the evolution of the closeness centrality confirms what has been observed in the purely binary case: the aforementioned ‘hubs’ ensure the vast majority of nodes to be reachable (hence, to be close to each other) quite easily. Generally speaking, our null models tend to significantly underestimate both the weighted betweenness centrality and the weighted eigenvector centrality, signaling the presence of fewer-nodes-with-heavier-connections than predicted by chance. Bottom panels: comparison between the BLN on day 20-03-2018, a configuration generated by the UBCM-induced CREMA model and a configuration generated by the deterministic CREMA model for the same day. The latter one distributes weights more evenly than observed, hence underestimating disparity and letting the strength distribution become wider: this has an interesting consequence, i.e. letting the size of the core become larger than observed, under this model – it amounts at the 12%, the 27% and the 21% of the total number of nodes on the considered day – while still allowing the unevenness of the WEC distribution rise. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

the communication between the first neighbors of node i is, upon its removal, our results seem to indicate that the BLN is becoming more and more ‘fault tolerant’ than its randomized counterpart (interestingly, it appeared to be much more fragile during the first half of its history). This result can be understood by imagining that a larger number of redundant connections has been established, among nodes, in the more recent snapshots of the BLN history — whence the rise of the average clustering coefficient as well.

As an additional exercise, let us inspect the evolution of the BLN global efficiency as nodes are removed either randomly or sequentially, after they have been sorted in decreasing order of weighted degree, closeness, betweenness and eigenvector centrality. The results of our exercise are shown in Figs. 12 and 13. The depicted trends of Fig. 12 are compatible with a *robust-yet-fragile* architecture, i.e. a topological structure that is robust against a random removal of nodes but fragile against a targeted removal of nodes (e.g. an attack) - or, more correctly, more robust against a random node removal than against a targeted node removal: notice how steeper the decrease of E_g is in the second case; moreover, removing nodes according to their WBC reduces the BLN global efficiency to the largest extent (for the vast majority of

snapshots, larger than removing nodes according to their WDC, WCC and WEC). Fig. 13 sheds light on the reason why such an effect is observed: removing nodes induces a fragmentation of the BLN structure into many (disconnected) components.

Fig. 12 also shows that the nodes whose removal brings the most severe damages to the BLN are those belonging to the core (see the next paragraph), whose size shrinks from $\gtrsim 20\%$ to $\approx 10\%$ of the total number of nodes.

Core-periphery detection. The result concerning the underestimation, by our null models, of the Gini index for the weighted eigenvector centrality - i.e. the presence of well connected nodes which, in turn, are also well connected among them – let us suppose the BLN to be characterized by a statistically significant core-periphery structure: here, however, we are interested in revealing the presence of a weighted core-periphery structure, i.e. a kind of mesoscale organization where core nodes are the ones sharing the ‘heaviest’ connections – and not just those with ‘many’ connections.

To this aim, we adopt a recently proposed approach, based upon the surprise formalism. In particular, we consider the evolution of the

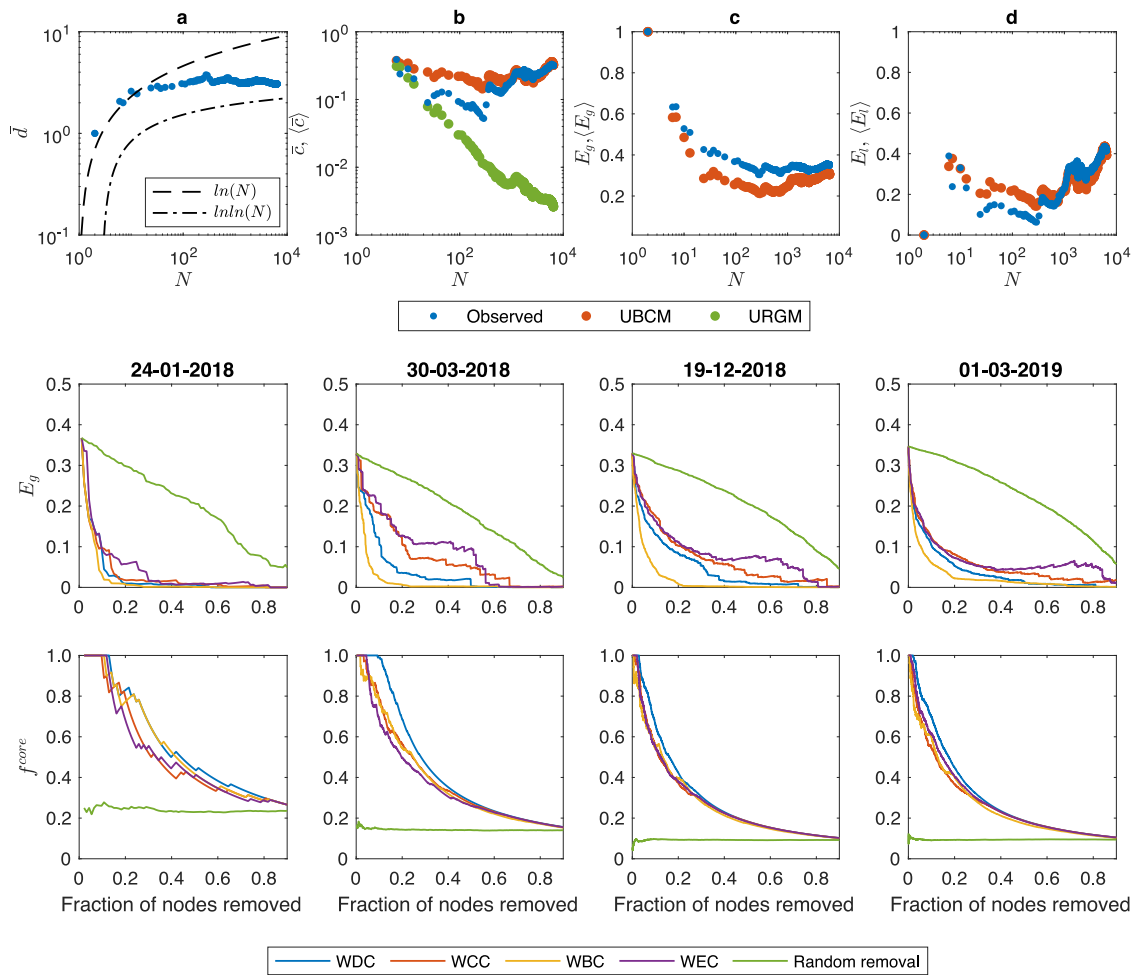


Fig. 12. Upper panels: (a) evolution of the BLN average path length \bar{d} and of the functions $\ln N$ and $\ln \ln N$; (b) evolution of the empirical average clustering coefficient $\overline{C_C}$ (blue dots) and of its expected values $\langle C_C \rangle$ under the URGM (green stars) and the UBCM (red squares); (c) evolution of the empirical global efficiency E_g (blue dots) and of its expected values $\langle E_g \rangle$ under the UBCM (red squares); (d) evolution of the empirical local efficiency E_l and of its expected values $\langle E_l \rangle$ under the UBCM (red squares). The BLN is indeed characterized by a small-world structure; moreover, while it has been always more-globally-efficient-than-expected under the UBCM, it has ‘recently’ become also more-locally-efficient-than-expected under the same null model. Middle panels: evolution of the BLN global efficiency, for our usual four snapshots, when nodes are removed either randomly (green trend) or sequentially, after having been sorted in decreasing order of weighted degree (blue trend), closeness (red trend), betweenness (yellow) and eigenvector (purple) centrality. The trends above characterize a robust-yet-fragile architecture: robust against a random node removal but fragile against a targeted node removal (e.g. an attack). Bottom panels: percentage of core nodes found within the set of nodes removed according to one of the two aforementioned criteria. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)




Fig. 13. Number of connected components (NCC) into which the BLN structure fragments upon removing nodes scattered versus the number of nodes (N_r - ranked by degree, closeness, betweenness and eigenvector centrality) whose removal induces them..

weighted bimodular surprise, $\mathcal{W}_{||}$, across the entire BLN history: it reveals that the statistical significance of the recovered core-periphery structure increases, a result leading to the conclusion that the description of the BLN structure provided by such a model becomes more and

more accurate as the network evolves. As an example, Fig. 13 shows the detected core-periphery structure on the snapshots depicted in the same figure: the nodes identified as belonging to the core and to the periphery are, respectively, colored in orange and purple. Notice also

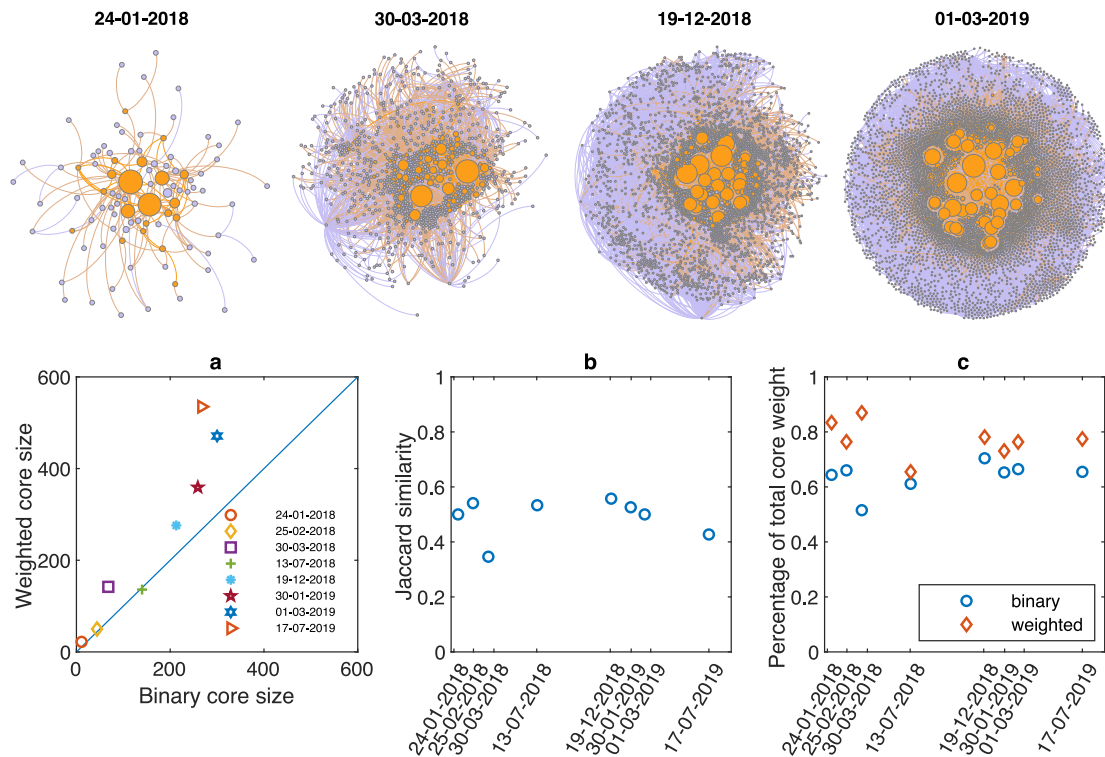


Fig. 14. Upper panels: core–periphery structure, as revealed by the weighted surprise, \mathcal{W}_i , for our four usual snapshots (core nodes are colored in orange and periphery nodes are colored in purple); the size of the nodes is proportional to their strength: hence, the nodes constituting the core of the network are precisely those with a larger strength. Bottom panels: (a) values of the size of the binary core scattered versus the values of the size of the weighted core; (b) evolution of (the percentage of) the overlap between the set of nodes belonging to the binary core and the set of nodes belonging to the weighted one, estimated via the Jaccard similarity, on a selected subset of snapshots of the BLN; evolution of the percentage of the total network weight embodied by ‘core connections’, amounting at $\approx 68\%$ in the binary case (blue circles) and at $\approx 77\%$ in the weighted one (red diamonds). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

that we have drawn the node size proportionally to the node strength: hence, larger nodes, i.e. the ones sharing the ‘heaviest’ connections, are precisely those constituting the core of the network.

First, let us check the correspondence between the nodes in the core (whose size will be indicated as N_{core}) and vertices with large weighted eigenvector centrality by ranking the nodes in decreasing order of WEC and checking the percentage of top N_{core} nodes that also belong to the core: it amounts at 56%, 60%, 57%, 62%, for our usual four snapshots. Then, let us compare the composition of the purely binary core — detected in [7] - with that of the weighted core. As Fig. 14 shows, a nice correspondence between the size of binary core and that of the weighted one indeed exists although, from a certain moment of the BLN history on, the binary core seems to ‘grow slower’ than the weighted one which, instead, enlarges to reach a size of ≈ 600 nodes: this further confirms that the nodes with a ‘large’ strength, revealed by surprise as the most central ones, do not necessarily coincide with those having a ‘large’ degree. The evolution of (the percentage of) the overlap between the set of nodes belonging to the binary core and the set of nodes belonging to the weighted one further confirms that the two sets do not coincide perfectly, although the Jaccard similarity steadily points out a $\approx 60\%$ of overlap: in other words, 60% of the nodes belong to both cores — likely, those hubs whose degree *and* strength are large enough to justify their core-ness in both senses; similarly, the percentage of the total network weight embodied by ‘core connections’ amounts at $\approx 68\%$ in the binary case and at $\approx 77\%$ in the weighted one.

5. Discussion

The analysis of the binary BLN structure carried out in [7] has revealed a system whose topology has become increasingly characterized by star-like structures, whose centers are constituted by ‘hubs’ to which

many nodes having a (much) small(er) degree, in turn, attach. Such a structure – whose disassortativity is confirmed by scattering the ANND values versus the degrees – could explain the more-than-expected level of unevenness characterizing the betweenness and the eigenvector centralization indices, suggesting them to be due to the emergence of channel-switching nodes — apparently, an unavoidable consequence of the way BLN is designed: on the one hand, as longer routes are more expensive, any two BLN users will search for a short(est) path; on the other, nodes have the incentive to become as central as possible, in order to maximize the transaction fees they may earn.

The tendency to centralization is observable even when considering weighted quantities, as the percentage of nodes whose connections embody the 51% of the total weight progressively reduces and the Gini coefficient of several (weighted) centrality measures steadily increases throughout the entire BLN history. This clearly points out the co-existence of nodes playing deeply different ‘structural’ roles, with ‘many’ peripheral vertices co-existing with ‘few’ core ones; if, on the one hand, this structure allows the global efficiency to achieve a large value (i.e. hubs facilitate the global exchange of information, being at the origin of another structural BLN peculiarity, i.e. its small-world -ness), on the other it highlights the tendency of the BLN architecture to become increasingly ‘less distributed’, a process having the undesirable consequence of making it increasingly fragile towards failures and attacks.

Distinguishing between the two is crucial, in order to properly understand the BLN robustness to ‘damages’. While resilience towards failures can be tested by looking at how the global efficiency ‘reacts’ to random node removal, resilience towards attacks can, instead, be quantified by implementing targeted removals of the ‘most important’ nodes. To this aim, we have ranked nodes in decreasing order of weighted degree, closeness, betweenness and eigenvector centrality and removed them, sequentially: the global efficiency drops rapidly

after few (core) nodes are deleted — in fact, for almost all snapshots, removing just *one top node* (according to any of the aforementioned criteria) is enough to disconnect the graph. Moreover, since top nodes are likely to be part of the core – whose size shrinks from $\geq 20\%$ to $\approx 10\%$ of the total number of nodes – our results indicate that the vertices belonging to it are precisely those whose removal causes the major structural damages. Random failures, instead, cause the decrease of E_g to be much less steep: taken together, the results above seem to indicate that the BLN topology is an example of *robust-yet-fragile* architecture, i.e. a structure that is robust against a random node removal but fragile against a targeted node removal (e.g. an attack).

CRedit authorship contribution statement

Jian-Hong Lin: Formal analysis, Writing – original draft, Writing – review & editing. **Emiliano Marchese:** Formal analysis, Writing – review & editing. **Claudio J. Tessone:** Conceptualization, Methodology, Writing – review & editing. **Tiziano Squartini:** Conceptualization, Methodology, Writing – original draft, Writing – review & editing.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

Data will be made available on request.

Acknowledgments

J.-H. L. and C.J.T. acknowledge financial support from the Swiss National Science Foundation grant no. #200021_182659. J.-H. L. also acknowledges financial support from the China Scholarship Council. T.S. acknowledges financial support from the EU H2020 project ‘So-BigData++’, grant n. 871042.

References

- Poon J, Dryja T. The bitcoin lightning network: Scalable off-chain instant payments. 2016.
- Tasca P, Tessone CJ. Taxonomy of blockchain technologies. principles of identification and classification. Ledger 2019;4:1–39.
- Spychiger F, Tasca P, Tessone CJ. Unveiling the importance and evolution of design components through the “tree of blockchain”. Frontiers in Blockchain 2021;3:613476.
- Brânzei S, Segal-Halevi E, Zohar A. How to charge lightning. 2017, ArXiv preprint arXiv:1712.10222.
- Béres F, Seres IA, Benczúr A, et al. A cryptoeconomic traffic analysis of Bitcoin’s Lightning Network. Cryptoecon Syst 2021;1(1):1–46.
- Lee S, Kim H. On the robustness of lightning network in bitcoin. Pervasive Mob Comput 2020;61:101108.
- Lin J-H, Primicerio K, Squartini T, Decker C, Tessone CJ. Lightning network: a second path towards centralisation of the bitcoin economy. New J Phys 2020;22(8):083022.
- Martinazzi S, Flori A. The evolving topology of the Lightning Network: Centralization, efficiency, robustness, synchronization, and anonymity. PLoS ONE 2020;15(1):e0225966.
- Motamed AP, Bahrak B. Quantitative analysis of cryptocurrencies transaction graph. Appl Netw Sci 2019;4(1):1–21.
- Campajola C, Cristodaro R, De Collibus FM, Yan T, Vallarano N, Tessone CJ. The evolution of centralisation on cryptocurrency platforms. 2022, ArXiv preprint arXiv:2206.05081.
- De Collibus FM, Partida A, Piškorec M, Tessone CJ. Heterogeneous preferential attachment in key ethereum-based cryptoassets. Frontiers in Physics 2021;5:68.
- De Collibus FM, Piškorec M, Partida A, Tessone CJ. The structural role of smart contracts and exchanges in the centralisation of ethereum-based cryptoassets. Entropy 2022;24(8):1048.
- Guo Y, Tong J, Feng C. A measurement study of bitcoin lightning network. In: 2019 IEEE international conference on blockchain. IEEE; 2019, p. 202–11.
- Rohrer E, Malliaris J, Tschorsch F. Discharged payment channels: Quantifying the lightning network’s resilience to topology-based attacks. In: 2019 IEEE European symposium on security and privacy workshops. IEEE; 2019, p. 347–56.
- Mizrahi A, Zohar A. Congestion attacks in payment channel networks. In: International conference on financial cryptography and data security. Springer; 2021, p. 170–88.
- Conoscenti M, Vetrò A, De Martin JC. Hubs, rebalancing and service providers in the lightning network. IEEE Access 2019;7:132828–40.
- Bartolucci S, Caccioli F, Vivo P. A percolation model for the emergence of the Bitcoin Lightning Network. Sci Rep 2020;10(1):1–14.
- Javarone MA, Wright CS. From bitcoin to bitcoin cash: a network analysis. In: Proceedings of the 1st workshop on cryptocurrencies and blockchains for distributed systems. 2018, p. 77–81.
- Wu J, Liu J, Zhao Y, Zheng Z. Analysis of cryptocurrency transactions from a network perspective: An overview. J Netw Comput Appl 2021;190:103139.
- Decker C. Lightning network research; topology datasets. <https://github.com/lrresearch/topology>. [Accessed 1 October 2020].
- Newman M. The structure and function of complex networks. SIAM Rev 2003;45(2):167–256.
- Newman M, Barabási A-L, Watts DJ. The structure and dynamics of networks. Princeton University Press; 2006.
- Srinivasan BS, Lee L. Quantifying decentralization. 2017, <https://News.Earn.Com/Quantifying-Decentralization-E39db233c28e>.
- Bonacich P. Power and centrality: A family of measures. Am J Sociol 1987;92(5):1170–82.
- Borgatti SP. Centrality and network flow. Social Networks 2005;27(1):55–71.
- Newman M. Networks. An introduction. Oxford University Press; 2018.
- Rodrigues FA. Network centrality: an introduction. In: A mathematical modeling approach from nonlinear dynamics to complex systems. Springer; 2019, p. 177–96.
- Newman M. A measure of betweenness centrality based on random walks. Social Networks 2005;27(1):39–54.
- Pfritzer R, Scholtes I, Garas A, Tessone CJ, Schweitzer F. Betweenness preference: Quantifying correlations in the topological dynamics of temporal networks. Phys Rev Lett 2013;110(19):198701.
- Bonacich P. Some unique properties of eigenvector centrality. Social Networks 2007;29(4):555–64.
- Liu J-G, Lin J-H, Guo Q, Zhou T. Locating influential nodes via dynamics-sensitive centrality. Sci Rep 2016;6:21380.
- Morgan J. The anatomy of income distribution. Rev Econ Stat 1962;270–83.
- Crucitti P, Latora V, Porta S. Centrality measures in spatial networks of urban streets. Phys Rev E 2006;73(3):036125.
- Watts DJ, Strogatz SH. Collective dynamics of ‘small-world’ networks. Nature 1998;393(6684):440–2.
- Latora V, Marchiori M. Efficient behavior of small-world networks. Phys Rev Lett 2001;87(19):198701.
- Amara L, Scala A, Barthelemy M, Stanley HE. Classes of small-world networks. In: The structure and dynamics of networks. Princeton University Press; 2011, p. 207–10.
- Mariani MS, Ren Z-M, Bascompte J, Tessone CJ. Nestedness in complex networks: observation, emergence, and implications. Phys Rep 2019;813:1–90.
- Lin J-H, Tessone CJ, Mariani MS. Nestedness maximization in complex networks through the fitness-complexity algorithm. Entropy 2018;20(10):768.
- Solé-Ribalta A, Tessone CJ, Mariani MS, Borge-Holthoefer J. Revealing in-block nestedness: detection and benchmarking. Phys Rev E 2018;97(6):062302.
- König MD, Tessone CJ, Zenou Y. Nestedness in networks: a theoretical model and some applications. Theor Econ 2014;9(3):695–752.
- Bardoscia M, De Luca G, Livan G, Marsili M, Tessone CJ. The social climbing game. Journal of statistical physics 2013;151(3):440–57.
- Marchese E, Caldarelli G, Squartini T. Detecting mesoscale structures by surprise. Commun Phys 2022;5(1):1–16.
- van Lidth de Jeude J, Caldarelli G, Squartini T. Detecting core-periphery structures by surprise. Europhys Lett 2019;125:68001.
- Parisi F, Squartini T, Garlaschelli D. A faster horse on a safer trail: Generalized inference for the efficient reconstruction of weighted networks. New J Phys 2020;22:053053.
- Vallarano N, Bruno M, Marchese E, Trapani G, Saracco F, Cimini G, Zanon M, Squartini T. Fast and scalable likelihood maximization for exponential random graph models with local constraints. Sci Rep 2021;11(1):1–33.
- Clauset A, Shalizi CR, Newman ME. Power-law distributions in empirical data. SIAM Rev 2009;51(4):661–703.