

Detecting and Mitigating Points of Failure in Community Networks: a Graph-based Approach

Leonardo Maccari *Member, IEEE*

Abstract—A Community Network is a bottom-up network created by a community of people with the goal of gaining control of their communications and overcoming digital divide. Community Networks are blooming, they range from small ones (tens of nodes) to gigantic ones (tens of thousands of nodes). They are made primarily of wireless links but in some cases they mix wired and wireless technologies. Community Networks are generally unplanned and non-layered, and the community tries to mirror the same approach in its governance, avoiding unnecessary management structures and relying on self-organization and spontaneous interactions. Community Networks are Peer Production platforms, a community of people that pools resources and contributes to build a shared value. While this value is generally immaterial (as in Wikipedia) Community Networks instead realize a distributed, peer-to-peer physical communication network.

This paper analyses *ninux.org*, the largest community network in Italy, and one of the eldest in Europe. The goal of the paper is to understand if the spontaneous growth of the network and the community leads to a technically robust network and a socially robust community, or it hides the presence of (potentially interdependent) points of failure. We will show that, in spite of the original motivations of the *ninux* community, the network is fragile under several aspects, and we suggest ways to improve it.

Index Terms—community networks, centrality, social network, robustness, network hierarchy

I. INTRODUCTION

Community Networks (CNs) are distributed mesh networks realized using a grassroots approach with two goals: i) building completely decentralized, bottom-up, open networks, with an alternative model compared to commercial Internet Service Providers (ISPs); ii) connecting otherwise unconnected regions. CNs are not only a low-cost alternative to last mile connections, they support a different vision of connectivity that is not-for-profit, participatory, Peer-to-Peer (P2P), and neutral. People in CNs oppose a model in which a few Internet Service Providers (ISPs) limit the user freedom in a way or another. They perceive the need of a distributed network infrastructure as the first step to build an ecosystem of P2P applications that can locally replace (or at least complement) the current mainstream, centralized, cloud-based web applications. CNs rely on decentralization and redundancy to achieve this goal, as in the original spirit of the Internet [1]. CNs try to apply the same decentralized, non-hierarchical infrastructure also to the governance of the network, which tends to be horizontal and

peer-to-peer. In this sense, CNs are unique techno-social Peer Production platforms with the goal of building an open, decentralized, participatory communication infrastructure. CNs are blooming, networks made of hundreds or even thousands of nodes daily used by thousands of people have been documented [2]. Given the state of the current market-driven Internet model, which introduces failures (approximately 50% of the world population was still disconnected in 2017 [3]) and power asymmetries [4], CNs appear as the building block of a new Internet model based on participation and bottom-up initiatives. For this reason they are gaining attention from different disciplines: the “subversive” potential of CNs is studied by social scientists [5] [6] [7], their peculiar features are interesting for networking researchers [8], [9] and they become the playground for distributed applications, such as P2P live streaming [10] and self-hosted community cloud services [11], [12].

Nevertheless, even systems that have been initially thought to be decentralized and resilient can degenerate in hierarchical and fragile structures. The goal of this paper is to perform a multi-layer analysis of one of these networks in order to understand if both the communication and the social networks are effectively distributed and robust. The research question we answer is: *are the network and the community effectively decentralized and resilient, or their spontaneous growth produces hidden points of failure?* This question is of fundamental importance for a CN; if it grows spontaneously with a robust infrastructure and a balanced participation of people in the project, its growth is sustainable. If instead the CN grows hiding some embedded paths of centralization and single points of failure, it may not survive the failure of one of its components, or the withdrawal of one of its key members. In practice, its bottom-up organization can not offer a level of reliability comparable to the top-down approach of a commercial Internet Provider, thus, it fails to reach its goals.

This paper studies *ninux*, the largest Italian CN. It applies the appropriate metrics to analyse the physical network and the logical communication network. It infers the social ties among the participants from the analysis of the discussion mailing lists and verifies its level of decentralization. It also introduces an algorithm to re-assign the ownership of some network nodes in order to rebalance the control among key participants of the community.

To the best of our knowledge this is the first paper that offers a deep analysis of a CN (or a communication network, in general) including three fundamental layers: its spatial distribution, its communication network, its social network. The three layers with their cross-implications influence the future sustainability of the CN.

Leonardo Maccari is with the Department of Information Engineering and Computer Science, University of Trento, Italy, Email: {maccari}@disi.unitn.it

This work was financed partially by the European Commission, H2020-ICT-2015 Programme, Grant Number 688768 ‘netCommons’ (Network Infrastructure as Commons).

The rest of the paper is organized as follows: Section II introduces CNs, Section III delineates the motivation of this paper, Sections IV to VI singularly analyse the layers of a CN, Section VII does a cross-layer analysis, Section VIII comments on the findings of the previous sections, Section IX introduces the node re-assigning algorithm, Section X reports on the state of the art and Section XI draws conclusions.

II. INTRODUCTION TO COMMUNITY NETWORKS

A CN is a communication network set-up by a community of people with a bottom-up, participatory approach. It is generally a wireless mesh network [13], extended when needed with wired connections. Wireless links are the building blocks of most of the CNs. The participants install devices on their roofs to create links that in turn create a multi-hop wireless mesh network. Today, with a budget of less than 150\$, a wireless link with a capacity up to 400Mb/s (using the IEEE 802.11ac standard) can be realized on a distance up to tens of miles (with line of sight). With this low-cost technology such networks grow to tens, hundreds of even thousands of nodes that cover entire cities with a minimal fraction of the cost of the wired equivalent [2]. Each network node generates traffic, receives traffic and also routes traffic, similarly to an Internet router. Some CNs are closed networks with internal services, some are connected to the Internet. In the latter case, they replace the last-mile connection that is typically offered by a commercial ISP. There are large examples of both systems, like the AWMN network of Athens¹, that is mostly focused on internal services, or the FreiFunk network in Germany, that instead is primarily an access network². Guifi.net is by far the largest network³ documented in the literature [14]. While the technical concepts behind CNs are not new [15] their development in the last period was remarkable. Today, CNs represent an extremely interesting and timely research topic with many communities rapidly growing [2], [16].

From their technical organization derives also a very interesting social aspect: since the technical layer allows the construction of a decentralized network, communities try to keep also the network governance decentralized and cooperative. In CNs, each network node corresponds to a person, a family, an association or a small business and there is no single owner or the network. CNs use a participatory approach which is a key to reduce costs because nodes are installed on private spaces and people collectively contribute to install and maintain them. The network grows “organically” when the underlying community grows and does not require significant capital expenditure to bootstrap. In this sense, CNs are *Peer Production* platforms, i.e. Internet-enabled peer-to-peer organization platforms that allow communities to create social good, like Wikipedia or the many Free Software communities [17]. Community networks extend this concept to the material world, they borrow the principles of Peer Production and apply them to the construction of a communication network.

¹see <http://awmn.net>.

²see <http://freifunk.net>

³Guifi is a mixed wireless/wired network that counts tens of thousands of nodes in eastern Spain, see <http://guifi.net>

<i>Attribute</i>	<i>Value</i>
# nodes	114
# edges	128
average degree	2.246
diameter	13
average path length	6.014
modularity	0.79
density	0.02
average clustering coefficient	0.067

Table I: Main attributes of the ninux communication network, Jan. 2014.

Note that the fact that both the network and the governance are distributed is not by chance. Being distributed alleviates the community from the burden of managing a hierarchical technical and social structure. People in a CN focus on the most important part of the work, setting-up new links, instead of getting carried away by statues, internal rules and time-consuming discussions. As this paper shows, the will to be “distributed” does not guarantee a distributed outcome.

Note that there are CNs that use wired technologies and have a large success in fighting digital divide. Yet in that case, their evolution model and motivation is different from the one we consider in this paper.

A. The Ninux Community Network

As other community networks ninux has strong political motivations: the construction of an independent, robust, decentralized network infrastructure [18] [19]. ninux participants have a critical opinion of ISPs and service providers motivated by the recent discussions about neutrality, privacy and forced disconnections. They identify the root cause of these problems in the centralization (both in the technical and governance sense) of the networks and of the services, and for this reason they build their own decentralized network. The ninux community did not create a formal association, it does not assign formal roles or responsibilities to people. The discussions in the community are primarily carried on in the mailing lists and in weekly face-to-face meetings, and decisions are taken with a consensus-based method. In this sense, ninux is a full do-ocracy [20] that was able to build a network made of hundreds of nodes along the whole Italian peninsula but especially concentrated in Rome [21].

Table I reports the main characteristics of the ninux network in Rome, and Fig. 1 reports a snapshot of the topology of the network. Throughout the paper we refer to ninux as only the nodes of the network concentrated in Rome.

III. MOTIVATIONS OF THE PAPER

A strictly hierarchical organization has a few critical elements, if these elements fail, the consequences can be catastrophic. For this reason, when planning such an organization some redundancy must be included. This is true for communication networks that must be designed to avoid single points of failure and provide alternative routes, but also for

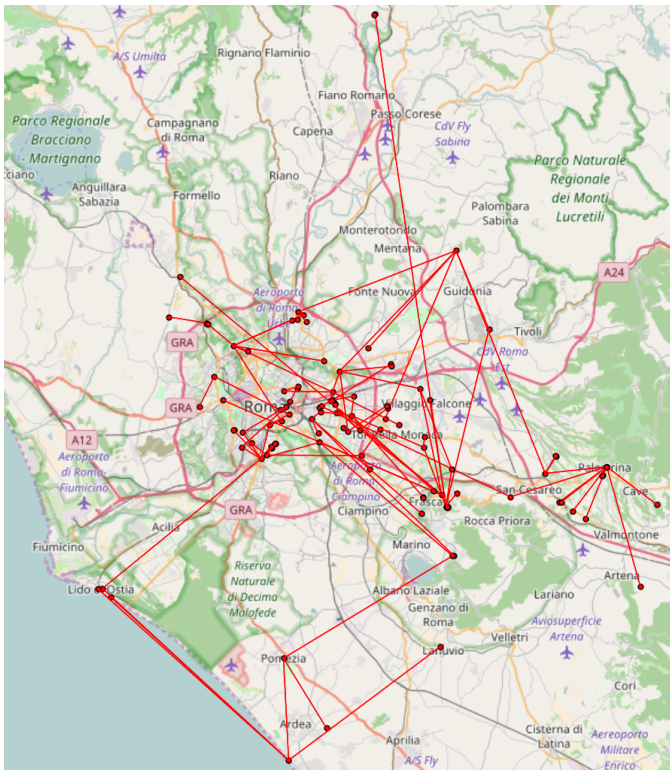


Figure 1: A snapshot of the topology of the ninux CN in Jan 2014.

governance structures, which generally include several bodies that counterbalance each other and offer redundancy (a CEO, a Board of Governors, Vice-Presidents...). When instead an organization evolves with an horizontal “spontaneous” pattern there is no “by design” hierarchy, and thus, people involved develop a naive belief that the network will grow without single points of failure. The objective of this paper is to study if the mix of social and technical approach that generates a CN effectively produces a distributed and resilient network, or it hides embedded patterns of centralization that make the CN fragile.

Intuitively, the more a CN is hierarchical (but people do not realize it) the less it will be resilient to technical failures, or to the departure of a small number of key individuals. To make this intuition concretely measurable we need to analyse the three layers that compose the CN (physical network, communication network and social network) and apply the correct metrics to understand how robust they are. The rest of this paper analyses ninux in order to answer the following questions:

- 1) Is the physical network topology strongly hierarchical?
- 2) Is the control on the communication network distributed fairly among the node owners?
- 3) Is the social network inferred by the analysis of the project mailing list fragile, i.e. a few people play a dominant role in the exchange of information?

Furthermore, a joint analysis of the various layers can be made to answer another question:

- 4) Are the single points of failure in the various layers

related? (i.e. are node owners that control critical portions of the network also dominant in the social network?).

Sections IV to VII analyse each of these questions respectively. Each section first defines the relevant graph-theoretic metrics, then applies them to the specific case of ninux. In Section IV we also compare the results with another network for which we have the exact position of nodes. ninux is the only network for which we can access data to perform the analysis of the other layers. Finally Sections VIII and IX show that the network presents correlated points of failure, and propose an heuristic algorithm to redistribute node ownership.

The inspiration for this work comes mainly from two research areas. The first is the study of scale-free networks, that showed that many networks critically rely on a minimal number of densely connected hubs. In practice, even if there is no intentional design, different networks build an internal hierarchy and produce the well-known “robust-but-fragile” effect [22]. The second source of inspiration comes from social science. History is full of cases in which people in a community join forces to produce some common good with a democratic approach but their outcome is far from being democratic. Sociologist Robert Michels, at the beginning of the 20th century called this the “Iron Law”: the tendency of horizontal and democratic communities to degenerate into oligarchies. Recently, Shaw and Mako Hill showed that the same pattern emerges on some well-known on-line Peer Production platforms in which a few individuals govern large communities and concentrate power, contradicting their original spirit [23].

A. The data-set

Three sources of information were used for this paper, here a brief introduction is given, while Appendix A contains more details, pointer to the data-sets and source code. All the data refer to the period 2013-2016, when the data-collection was carried on.

The first source is the database of the nodes maintained by the community and used to visualize the network topology⁴. The database contains the nodes (active or not), their position, the owner of the node with an e-mail address, and the active links. While the DB contains all the nodes that have ever been activated, it does not maintain a history of the links, it just shows the links available at the time the map is accessed. As the topology changes, and we wanted to obtain the snapshot with the largest size, we matched this information with an archive of dumps of the topology obtained by one of the network users. Each timestamped dump reports the link between active nodes in the period Mar. 2011 - Feb. 2016. From this set of dumps we obtained all the links that were activated in this period of time and derived the snapshot of the largest possible physical network. We call this graph G and we used it to perform the spatial analysis in Section IV. Table I and Figure 1 refer to G . A similar database was accessed to analyse the FunkFeuer network, which is another network used as a comparison in Section IV.

The second source is the topology as exported by the OLSR (Optimized Link-State Routing) routing protocol, a link-state

⁴See <http://map.ninux.org>

routing protocol that makes it possible for each node to be aware of the whole weighted network graph. We collected this information from running nodes and derived the routing graph (each node corresponds to an IP address of a router), which we refer to as $R(V, E)$, and we used it in Section V. Note that $R(V, E)$ differs from G in two aspects. First, $R(V, E)$ is larger than G , as some nodes in $R(V, E)$ are actually a collection of co-located devices and were merged into one single node in G . Second, edges in $R(V, E)$ are weighted using the quality metric used by the routing protocol, so that we can apply Dijkstra's algorithm exactly as the protocol does in real life. Please refer to [24] for more details on how the routing topology was derived.

Finally, the third source is the archive of the mailing lists of the ninux community of Rome that was monitored in the period Jan. 1st 2013 - Mar. 31st 2014. In that period a total of 5139 emails have been sent among which 4351 were answers to previous emails. Aggregating multiple email addresses to unique users is a tricky task that has been carried on with the techniques described in [25] for both node owners in the database and mailing list users. After aggregation, 85 distinct node owners were identified in the database, and 106 distinct senders in the mailing list. A first interesting observation is that only 44 out of the 85 owners participated to the mailing list in the observed period. Direct interaction with the community was necessary to supervise the email address aggregation and in general to give a qualitative interpretation of the quantitative results extracted from this data-set in Section VI.

IV. IS NINUX A HIERARCHICAL SPATIAL NETWORK?

Before we analyse ninux, we have to introduce the notion of spatial networks, and the "separation metric" that is used in the literature to assess their degree of centralization/hierarchy.

A. Separation in Spatial Networks

A spatial network is a graph $G(V, E)$ made of a set V of nodes and a set E of edges in which every node has a "position" attribute. Spatial networks are used to represent physical systems, such as road networks or power-line distribution networks [26]. Their growth can be modelled using a Cost-Benefit Analysis (CBA) framework which seeks a balance between two competing effects [27]. The first is the classical "rich gets richer" effect, in which the probability that a new node v_i is connected to an existent node v_j increases with the number of neighbors that v_j already has (as in the Preferential Attachment model [28]). The second effect is given by a technological constraint: the cost of a physical edge increases with its length, so the probability of adding an edge between v_i and v_j decreases with the length of the edges between v_i and v_j . It has been shown that when the second effect has a non-negligible influence, spatial networks tend to become hierarchical networks [27]. Intuitively, a graph G on which we define a notion of distance between two nodes is hierarchical if three things happen:

- 1) A "root" node can be somehow identified
- 2) If v_i to v_j are at the same distance from the root, to go from v_i and v_j the shortest path climbs up towards the root and then descends again.

- 3) If we assign to each node an area of pertinence (in the physical space) and we navigate on a path that goes from the root towards the fringes of the network, the area assigned to a node is always included in the area of the previous node in the path.

In order to quantify this intuition, we introduce the "separation" metric, which is used in the literature to express how much a network can be considered hierarchical [27] (please see Appendix B for details on how we customized the metric to fit our case). First, we need to define a root node. In a network graph with loops, centrality metrics are generally used to rank the importance of a node. We tested several centrality metrics to define the root node, and we finally chose the node with the highest eccentricity. Then we call $l(v_i)$ the distance (in terms of hops on the shortest path) of v_i from the root node. Given $N(i)$, the neighbor set of v_i , we define a subset $N'(i) \subset N(i)$ as:

$$N'(i) = \{v_j \mid v_j \in N(i) \wedge l(v_j) > l(v_i)\}. \quad (1)$$

Given $N'(i)$, an "influence zone" is defined, that is a geographical area that contains $N'(i)$. The definition of the influence zone is context-dependent, we use the convex hull of all the nodes in the subset (again, see Appendix B for more details). Given a level l and a node $v_i \mid l(v_i) = l$ we call \mathcal{I}_l^i the influence zone of v_i , and we call $\mathcal{I}_l = \cup_i \mathcal{I}_l^i$. A spatial network is said to be geographically separated if both these conditions hold:

$$\mathcal{I}_l^i \cap \mathcal{I}_l^j = \emptyset \quad \forall l, i \neq j \quad (2)$$

$$\mathcal{I}_{l+1} \subset \mathcal{I}_l \quad \forall l \quad (3)$$

A network that is fully separated is strongly hierarchical: Eq. (2) says that nodes at the same level can not communicate without first ascending and then descending again in the network tree. In such a network the nodes that are close to the root node are more important than the others, and their failure will dramatically impact the rest of the network. Equation (3) says that the influence of every node is included in the influence of its parent, in practice, the network does not grow if the root node does not enlarge its influence zone (and the same stands in cascade for the other levels).

Real networks are never completely separated, so a metric to measure their degree of separation is needed. The separation of two zones in the same level is defined as:

$$s_l(i, j) = 1 - \frac{\text{Area}(\mathcal{I}_l^i \cap \mathcal{I}_l^j)}{\min(\text{Area}(\mathcal{I}_l^i), \text{Area}(\mathcal{I}_l^j))} \quad (4)$$

We call s_l the intra-level separation index for level l , given by the average of the separation between every couple of zones in the same level. Let V_l be the subset of V containing all the nodes at level l , then:

$$s_l = \frac{2 \sum_{i=0}^{\|V_l\|} \sum_{j=i+1}^{\|V_l\|} s_l(i, j)}{\|V_l\|(\|V_l\| - 1)} \quad (5)$$

Where $\frac{\|V_l\|(\|V_l\| - 1)}{2}$ is the number of all the possible couples (v_i, v_j) of nodes in V_l and $\|\cdot\|$ is the size of a set. If s_l is

close to 1, then the zones inside level l are almost perfectly separated, if it approaches 0, there is overlapping between zones.

We also define a metric to measure inter-level separation:

$$\hat{s}_l = 100 * \frac{\text{Area}(\cup_{i=0}^{l-1} \mathcal{I}_i)}{\text{Area}(\cup_{i=0}^{l_{max}} \mathcal{I}_i)} \quad (6)$$

Where l_{max} is the number of levels in the network. The progression of the values of \hat{s}_l tells if the levels of the network are progressively covering larger portions of the territory or the levels are organized like Russian dolls.

The reason why we are interested in spatial separation is that the literature shows that in spatial networks modelled with a CBA the average value of s_l quickly saturates to 1 as soon as the cost-per-mile of a link becomes non-negligible [27]. We want to verify to what extent the ninux community unwillingly built a hierarchical, not redundant and thus, fragile network.

B. The separation of the ninux network

As a first observation, we note that in a CN both the effects needed in the CBA exist. Consider a node v_i placed in a dominating position such as on the top of a hill. Potentially, many newcomer nodes will have line of sight with v_i , and thus, v_i will probably have many neighbors. As a consequence, the community will improve the node adding radio devices, which will increase the horizontal angle covered by the node (recall that large-scale CNs primarily use directive antennas). This will make it even more likely that new neighbors can be added, so the “rich gets richer” effect takes place. However in a CN the performance of a wireless link decreases with its length. The longest the link, the higher its cost, considering both the higher price of devices with highly directive antennas and the complexity of pointing them, which is a non-trivial operation. Therefore, it is realistic to apply a CBA to CNs.

Table II reports the intra-level separation at each level in the ninux network, together with the number of zones per level. Table II shows that for each level, s_l is fairly high, that means that the nodes of the network in the same level are not densely connected and the chances that there are “horizontal” paths from a node to another are few. If a link going from level l to level $l - 1$ breaks, there is no way to re-route traffic across other zones in the same level, as the interest regions do not overlap. Thus, the lower the level of a node, the higher the number of shortest paths that pass through it, and the harder it is to repair its potential failure. This shows that there is no direct incentive in making the network dense and redundant, while there are incentives in connecting the highest number of nodes with the lowest number of edges. As it happens to other spatial networks, ninux tends to evolve towards a hierarchical organization.

Table III reports the values of the inter-level separation, the corresponding percent of the total area covered by level l , and the percentage of nodes included in each level. Inter-level separation shows that, contrarily to other networks, a CN does not enlarge from its center, but from its edges. This means the network is not strictly hierarchical and allows the fringe of the network to expand without having to add antennas on the

<i>level</i>	<i>zones</i>	<i>s_l</i>
0	1	-
1	1	-
2	3	0.93
3	4	0.95
4	6	0.90
5	9	1.00
6	6	0.80

Table II: Average intra-level separation in ninux.

<i>level</i>	\hat{s}_l	Area (%)	Nodes (%)
0	0.37	0.37	5
1	22.85	22.85	44
2	31.33	7.78	63
3	54.87	12.55	78
4	66.37	45.66	81
5	100.00	31.60	100
6	100.00	9.63	100

Table III: ninux inter-level separation, percent of area covered at each level and corresponding percentage of nodes. Since areas are overlapping the third column does not sum to 100%, and the sixth level is fully included in the previous ones.

nodes in the center of the network. This characteristic is quite evident observing Fig. 2, that shows the covered area for each level.

<i>level</i>	<i>zones</i>	<i>s_l</i>
0	1	-
1	4	0.78
2	10	0.53
3	14	0.82
4	11	0.98
5	4	1.00

Table IV: Average intra-level separation within each level in FFWien.

<i>level</i>	\hat{s}_l	Area (%)	Nodes (%)
0	1.39	1.39	6
1	18.59	15.33	46
2	54.62	45.56	77
3	74.95	48.80	92
4	99.63	41.42	97
5	100.00	0.72	100

Table V: FFWien inter-level separation, percent of area covered at each level and corresponding percentage of nodes.

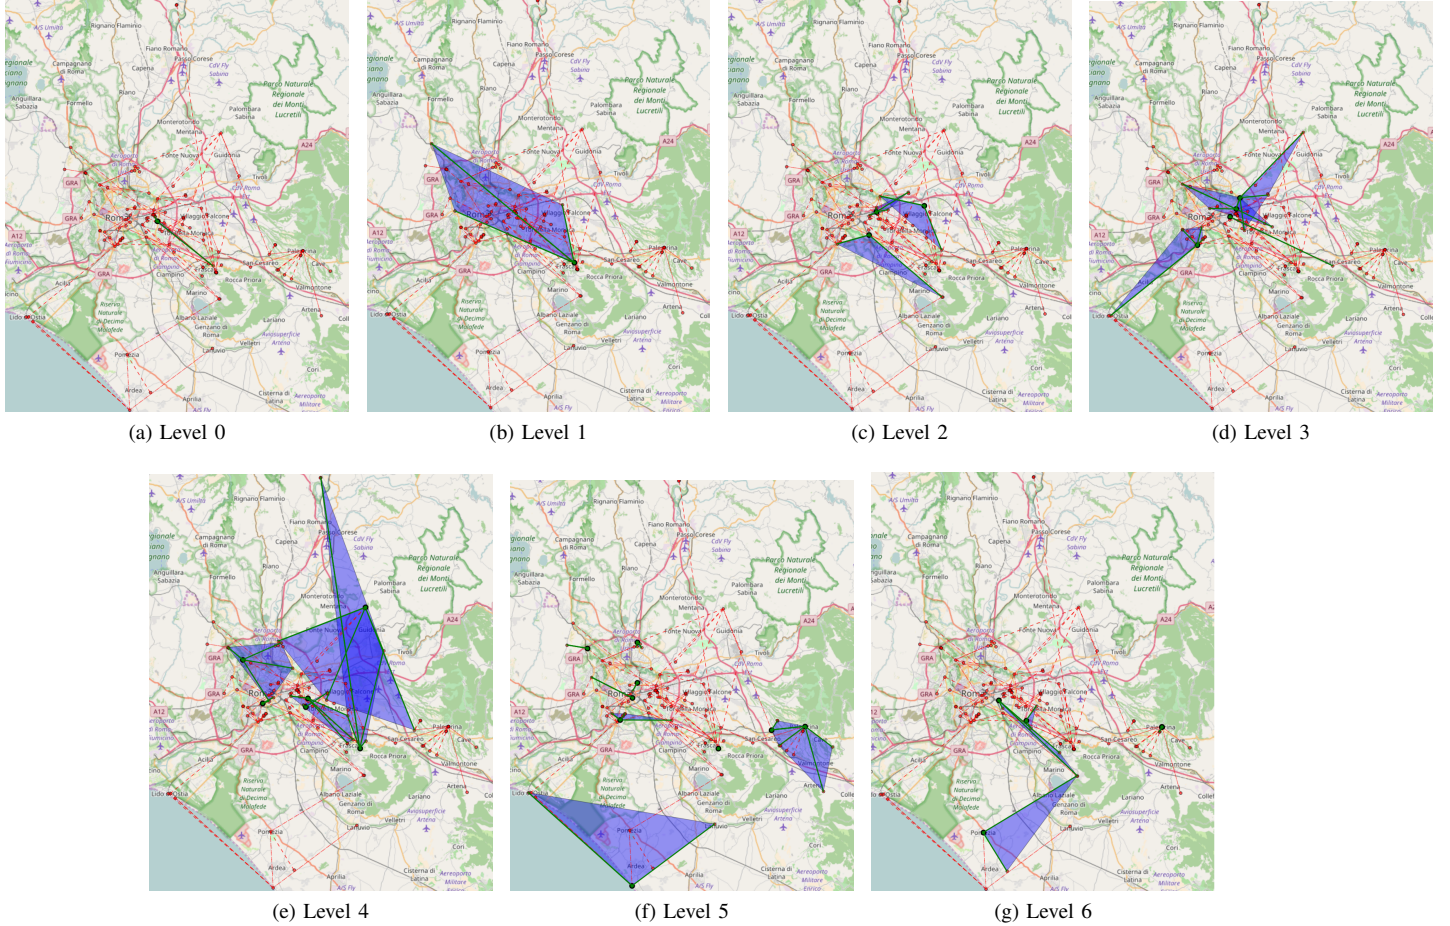


Figure 2: The Interest zone for the levels of the ninux network. Note that level 0 is made of a single zone that is very narrow and thus, barely visible in the picture.

C. Comparison with the FFWien CN

The spatial coverage of ninux depends on a number of factors that are specific to that network. In order to broaden our analysis we were able to collect the geo-referenced data of another network, the FunkFeuer network in Vienna (Austria), which we refer to as FFWien⁵. FFWien is a much denser network, made of 196 nodes and 249 edges, concentrated in an area that is 13 times smaller than the area of ninux. Tables IV and V report the separation metrics for FFWien and show a quite different situation, with lower average values, especially in the first 3 layers, which contain 77% of the nodes.

This difference outlines that a CN is strongly influenced by the external conditions, such as the density of nodes, the capacity and reach of the wireless devices, and the altitude profile of the area.

V. IS NODE OWNERSHIP IN NINUX EVENLY DISTRIBUTED?

This section introduces the association between the nodes in the routing network $R(V, E)$ and their owners, derived from the node database. It also correlates the owners with their importance in the network graph. Owner importance is

⁵See <http://funkfeuer.at>

evaluated with the group centrality metric and with the “owner robustness” defined as follows.

A. Owner Group Centrality and Owner Robustness

a) *Group Centrality*: In a weighted graph $R(V, E)$ where $P_{i,j} = \{v_i \dots v_j\}$ is the set of nodes that constitute the shortest path from node v_i to node v_j , the group centrality of a set of nodes $S = \{v_1 \dots v_n\} \subset V$ is given by:

$$B(S) = \frac{|\{P_{i,j} \mid i, j \in (1 \dots |V|), i \neq j \mid S \cap P_{i,j} \neq \emptyset\}|}{|\{P_{i,j} \mid i, j \in (1 \dots |V|), i \neq j\}|} \quad (7)$$

The group betweenness centrality is the fraction of shortest paths that pass through at least one node in the group. The centrality metric is computed running Dijkstra’s algorithm on the weighted network topology, and, without information on the traffic matrix is the best estimation of the number of traffic flows that a group of nodes can intercept. This definition is functionally equivalent to the original definition by Borgatti and Everett [29] with two marginal differences: first, it includes also the shortest paths that use a node in S as an endpoint, second it assumes there is only one shortest path between any couple of nodes. The second assumption derives from the fact




Figure 3: The number of nodes per user in the ninux network, top-20 users.

that IP networks generally don't support multipath routing, so one shortest path is used at every instant.

b) *Network Robustness*: One classical way to inspect the robustness of network graphs is to remove some nodes in the topology and check the connectedness of the remaining network [30]. A network is fragile if removing a small amount of nodes it is split in many small separated networks. Let S_i be the set of nodes owned by owner i , and let $R_G(S_i)$ be the size of the largest connected component when S_i is removed from the network. $R_G(S_i)$ is a robustness metric, the closer to $\|V\|$, the better. With a little abuse of notation we call $R_G(S_i)$ the robustness of owner i (instead of calling it “the robustness of the network to the failure of all nodes owned by “ i ”). A related metric is the number of disconnected components left in the network when S_i is removed, which we call the fragility of i : $F_G(S_i)$.

B. ninux node Ownership

Fig. 3 presents the number of nodes possessed by the top-20 ninux participants, ordered by nodes owned. Over a total of 85 owners, one user possesses 17% of the nodes and the top-five people own 31% of the nodes, top-13 people own roughly 50% of the nodes, 61 people own just one node. If we exclude the first individual (that we call P_{top}), the ownership distribution is not particularly skewed, reflecting the fact that the number of owned nodes is generally limited by the number of physical locations to which the person has access (home, workplace, houses of relatives etc...). P_{top} owns 24 nodes and is not the owner of all the locations where the nodes are placed, he is simply a technically skilled person that very often offers his help to set up the network for newcomers. As a result, in the database he appears to be the owner and, in practice, he is the technical manager of the nodes.

Fig. 4 shows the group betweenness centrality computed on all the nodes owned by the same person. Fig. 4 shows two metrics, the “node-to-node owner centrality” and the “person-to-person owner centrality”. The first metric is exactly Eq. (7) when S groups all the nodes of a single person. The second metric is a modified version of Eq. (7) when $P_{i,j}$ is not computed on every couple of nodes but only on the shortest path that interconnects nodes belonging to two different people. It expresses the centrality of an owner




Figure 4: The owner centrality for the participants to the ninux network, top 20 users.




Figure 5: The size of the largest connected component remaining after the removal of the set S_i of the nodes belonging to owner i .

between couples of other owners. Both metrics show that P_{top} can potentially control between 80% and 90% of the traffic flows.

High centrality gives to P_{top} an advantage position to control the network. He would be able to spy on a large quantity of the traffic and to filter it. While there is no reason to believe the person was actually enforcing those behaviours, the important observation here is that such a large predominance in the network topology potentially gives to one single person a strong influence and a high decision power.

C. ninux Owner Robustness

To guarantee robustness some level of network design is generally required [31], and we have already shown that community networks are not in general robust to targeted attacks [24]. It is interesting to observe what happens if one person leaves the network or turns off all its nodes. Figs. 5 and 6 show both robustness and fragility of the owners ordered by the number of nodes owned (as in Fig. 3).

Again it is clear that there is one person that represents a single point of failure of the network. When the nodes belonging to P_{top} are removed the main connected component is reduced to less than half the size of the original network and the remaining nodes are distributed on more than 30 isolated components.




Figure 6: The number of disconnected components after the removal of the set S_i of the nodes belonging to owner i .

VI. IS THE NINUX SOCIAL INTERACTION WELL BALANCED?

The analysis of the mailing list messages helps understanding who are the individuals that lead the discussion inside the community. Two metrics defined in the literature have been chosen for this task [25]. The first is the normalized number of answered email per user: given a number X of total messages that reply to some other message, and being x_i the cumulative number of replies to any message sent by the i th person, $R(i) = \frac{x_i}{X}$ is the relevance metric. This is a basic metric that assumes that people that receive a high number of replies are able to generate interesting discussion topics, thus are considered important in the community.

Fig. 7 shows that the relevance to the mailing list is not equally distributed among the participants, a very small number of people lead the discussion. The cumulative distribution in Fig. 8 shows that as little as 6 people receive 50% of all the answers. This is not uncommon, for instance, in open source projects a minority of people leads the discussion [32]. Unfortunately ninux does not represent an exception.

The second metric is the centrality of a person in the mailing list social graph. The social graph is an undirected graph $G(V, E)$ in which every node v_i is a person in the mailing list and there is an unweighted edge between two nodes v_i, v_j if person v_j ever answered to person v_i (or vice-versa). Mailing list centrality for v_i is computed on the social graph as in Eq. (7) when $S = \{v_i\}$. Betweenness centrality on mailing lists is used to understand who is able to make other people join the same discussion, so that he/she can facilitate the flow of information in the community. Again, Fig. 9 shows that there is a small number of people connecting all the other participants, and one in particular whose centrality is at least the double of the others.

Another layer of analysis is given by the identification of communities in the mailing list graph. We applied the Louvain community detection algorithm [33] to the ninux mailing list. The algorithm identified 9 communities, among which 4 made of a single user, the partition modularity is 0.156. Fig. 10a shows the interaction graph between the communities of more than one person. The size of each community and the strength of each link is reported in Figs. 10b and 10c. The graph shows that the ninux mailing list is quite “compact”,




Figure 7: The fraction of answered emails per person over the total.




Figure 8: The cumulative distribution of the fraction of answered emails.

meaning that there are only five communities, three of which include 80 users and are very well connected with each other. The modularity is not high. Finally, apart from a small set of 4 source email addresses, everybody belongs to some community.

VII. MATCHING THE COMMUNICATION AND THE SOCIAL NETWORK

Fig. 11 reports the percentage overlap on the two betweenness rankings from Fig. 4 and Fig. 9. The percentage overlap gives a measure of the correlation between the two rankings. Given a family of sets B_i and the respective ordering functions $o_i(v)$ on their elements, we call B_i^k the first k element of B_i





Figure 9: The ranked centrality of the top 20 participants in the ninux mailing list.



(a)

Com.	Size		0	1	2	3	4
0	9	0	47	4	129	155	143
1	13	1	4	51	30	38	22
2	21	2	129	30	541	682	634
3	29	3	155	38	682	627	679
4	30	4	143	22	634	679	569

(b)

(c)

Figure 10: (a) The communities identified in the ninux mailing list. Circle size reflects community size, edges gradient represent the relative strength (% of exchanged emails) of the connection: deep-blue strong connection, light blue weak connection. (b) The size of each community. (c) Number of emails exchanged between communities.

ordered by $o_i(v)$: $B_i^k = \{v | v \in B_i, o_i(v) \leq k\}$. Given two sets B_1 and B_2 the percentage overlap $p(k)$ is a function of k that shows the percentage of elements present in both sets when considering only the first k elements:

$$p(k) = \frac{100}{k} \times ||B_1^k \cap B_2^k|| \quad (8)$$

Fig. 11 shows two fundamental points: the first is that P_{top} , the person that owns more nodes and has the highest person network centrality is the same one that has the highest centrality in the social graph. P_{top} contributes to the growth of the network and to the mailing list discussion in a way that gives him a tremendous power to steer the direction of the community. The second point evidences a different, and more encouraging trend. If we exclude P_{top} the correlation between the communication and the social network centrality is not strong since $p(10) = p(20) = 30\%$. Therefore it seems that there is diversity between the owners of the most critical nodes and the leaders of the discussion in the mailing list. The general idea that the use of open communication and discussion tools guarantees plurality and participation is only partly matched by reality.

Another encouraging element raises from Fig. 12 that reports the number of owners grouped by the community they belong to among the top 5, 10, 15, and 20 node owners. Communities are ordered for their size from the largest to the smallest. The figure shows that the distribution of the top owners per community is not particularly skewed towards one




Figure 11: The percentage overlap metric computed on the ranked mailing-list and group node centrality.

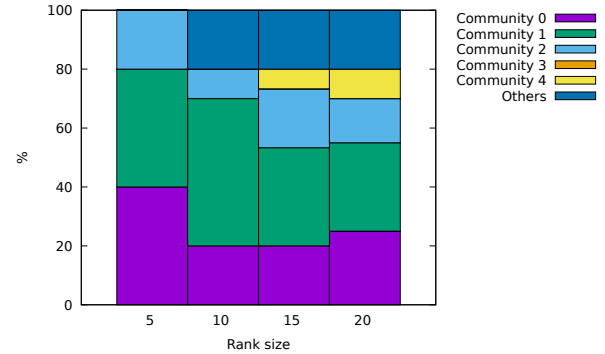


Figure 12: For the top 5, 10, 15, and 20 owners of network nodes, the percentage of nodes belonging to each community.

community. At least three communities are present in all the bars so there is not a single clique of users that dominates the communication network *and* the mailing list discussion. Another key element is that the owners with more nodes actually participate to the mailing list. Even if only 44 owners over 85 are present in the mailing list, only 20% of the top 10, 15, and 20 owners do not participate to the mailing list.

VIII. INTERPRETATION OF THE RESULTS

The spatial analysis, the distribution of the ownership, and the person centrality show that, albeit the goal of the ninux community is to build a technically and socially decentralized network, the results diverge from the goal. The network is spatially hierarchical, so nodes in the low levels have more importance than nodes in the high levels. Moreover, one person in ninux managed a sufficient number of nodes to be able to control the network and to be a single point of failure. The same person, given his technical skills was a central person in the social network of the community, so he had an influential voice in the email interactions.

Indeed, direct discussion engaged with people in the community revealed that this person left the community in 2015 and the nodes he managed started to fail and disconnect entire areas. At the time of writing the main component of the network is made of 87 nodes, and the other nodes were disconnected from the main component. In conclusion, ninux does not seem to be a robust do-ocracy, since the network at

the time of data collection had a huge single point of failure represented by P_{top} .

However the situation changes excluding P_{top} from the analysis. Fig. 3 shows that the maximum number of owned nodes is generally capped by the amount of physical locations that the users have access to, which intrinsically limits the chances of some individual to take over the network. Also, even if the social network metrics show that the relevance of the participants to the mailing list is not evenly distributed, the correlation between the most relevant node owners and the most relevant members of the mailing list is low (see Fig. 11). This means that people participate to the community in diverse ways, with the construction of new nodes or rising discussion topics.

A simple solution to this problem would be to prevent people to manage nodes in physical locations they do not own. This way, Wi-Fi range limitation would not allow a single person to be too central, and thus too critical for the network economy. This would change the nature of a CN which instead, to grow, must be participated not only by individuals but also by associations and small businesses that can be physically located in several places and thus, may own several nodes. A better solution is to reassign the ownership of nodes or to share the management credentials of nodes among several people as next section proposes.

A possible mean to reduce network separation would be to increment the density of links per node. At the current state of things, a new node v_i is connected only to the closest node in line of sight, and then, a new device is added to it only when some other new node v_j needs to connect to v_i to enter the network. One way of increasing the density is to mount an additional device on v_i at its creation, even if there is no other node to connect to at that time. Adding spare devices pointing to an uncovered direction will make it more likely that in the future new nodes will join, or that existing nodes will be connected to more than one other node, in order to increase the density. With an even stricter approach this could be translated into avoiding leaf nodes, so that a new node is added to the network only if it can connect to at least two other nodes.

A. A Path Towards Generalization

The analysis we carried on ninux is specific to this network and depends on the physical, routing and community graphs, which in turn depend on the internal organization of ninux. The only document that determines the “governance” of ninux is the so-called picopeering agreement, which mandates a few rules to be accepted by new node owners. The picopeering agreement has been formulated by a group of activists and is in use by many community networks in Europe⁶, which suggests that other communities could have developed with a pattern similar to ninux. To confirm this intuition the same analysis should be performed on more networks, which is a daunting task. While it is relatively easy to compare network topologies extracted by the routing layer, it is very hard to access to data on node position and node ownership. Even harder would be

to access data on the on-line interactions and match them with node owners due to the need of personal interactions with the community and of course, language barriers.

An alternative approach would be to use synthetic algorithms for the generation of the physical network topology, the routing topology and for the corresponding communities. While there is a large body of literature on community detection [34] and studies that describe the participation to community projects [32], the availability of large-scale mesh network routing topologies is very scarce. Moreover, it was shown that synthetic topologies created with the goal of achieving a specific topological feature (i.e. degree distribution) are not sufficiently detailed to catch the interesting features of real communication networks [35]. Finally, examples of the physical topologies of real mesh networks are not available. To overcome these limitations a possible option is to design a network topology generator that takes into account the real terrain conditions of a certain location. To realize such instrument we need precise data on the terrain elevation, including building elevation and building shapes. Many public administrations in Europe and North America publish open data sets on buildings altitude obtained with Lidar (light detection and ranging) aerial surveying campaigns. From these data sets a connectivity matrix can be built, expressing the presence of line-of-sight between any couple of two buildings and possibly the expected pathloss in dB (and thus, the quality on the link in the routing graph). Matching this connectivity matrix with data from the national census, the generator can create realistic topologies applicable to zones with low Internet penetration, or low-income. From these topologies, both the physical and network layer can be analysed.

This development would make it possible to analyze the robustness of realistic physical and routing layer in order to generalize the observations made on ninux, and also to fine tune the node re-assigning algorithm. It would help characterize any mesh network that evolves with an “organic” model as the one described so far. This activity is undergoing but requires an amount of work that is out of the scope of this paper⁷.

IX. NODE RE-ASSIGNMENT PROCEDURE

Once the community accepts that a person can administer one node without owning the corresponding physical location, we can leverage on this to redistribute the responsibility of some nodes from the original owner to somebody else, in order to reduce the importance of each single person. This section introduces an algorithm whose goal is to raise the minimum $R_G(S_i)$ beyond a certain threshold.

Re-assigning the control of nodes can be done with any node-labelling algorithm that maximises the fair distribution of nodes among the people in the community. Of course, that algorithm would not consider practical constraints, such as the chances that the old owner would not trust the new owner, or that the new owner must have physical access to the node. In practice, some real-world constraints must be introduced to make the resulting re-assignment practically useful.

⁶See <https://picopeer.net/>.

⁷See <https://github.com/AdvancedNetworkingSystems/TerrainAnalysis>.

Let $C(i)$ be the community of node i in the social graph, and S_i the nodes owned by owner i . Let S be the set of all the nodes in the communication network and C be the set of all the node owners. When re-assigning a node v from owner i to node j the proposed assignment scheme is subject to the following constraints:

- $S_j \neq \emptyset$: In order for j to have enough technical skills he must be the owner of at least one node before re-assignment.
- $j \in C(i)$: If this is not possible due to the previous condition, then j will be chosen iteratively from the next community that has the strongest link with $C(i)$.
- The probability of owner j of being assigned the management of node v decreases linearly with the minimum distance (in number of hops) from v to any node in S_j .

The last point can be expressed formally as: consider owners j and k that are two candidates for being assigned v . We call $d_{ij} = \min\{\|P_{ij}\| \mid \forall i, j \mid v_i \in S_i \wedge v_j \in S_j\}$. Then, for the probabilities $P(j)$ and $P(k)$ of owner j and k respectively, to be assigned v holds: $P(j) = P(k) \frac{d_{ik}}{d_{ij}}$. In other words, there is a bias in reassigning a node to owner j if owner j owns a node that is topologically close to v . The rationale behind this choice is twofold, first is that we assume that being topologically close implies also being physically close, which is important when physical maintenance is needed. Second, when v needs to be re-configured or maintained, being “topologically” close means that it is more likely that owner j can access node v with a better connection than a person that is topologically far away.

Listing 1 describes an heuristic algorithm that takes the least robust owner and reassigns his/her nodes up to when his/her robustness is higher than a threshold T . The procedure must be repeated up to when $R_G(S_i) > T \forall i$.

Before we describe the details of the algorithm, it is important to stress on a feature that makes it suitable for our context. The algorithm is an heuristic, iterative one which does not require a global optimization that would probably trigger a large number of re-assignments. It can be used on an existing network and it can be periodically run using the CN data as the network and the community evolves. Another key feature is to be *explainable*, meaning that people can easily understand the way it works and modify some parameters (i.e. modify T or add exceptions for some users) to fit their case. This is extremely important since people in a CN do not easily accept modifications to their network that are somehow “imposed” by some external factor they do not understand. They are instead available to share responsibilities for the good of the community, if they clearly see the reason. Albeit a global optimization would produce a better solution, it would likely make it hard to explain and accept. Finally, the algorithm increases the robustness of each individual one after the other, which makes it suitable to be used in CNs (like ninux) in which the average behaviour is correct, except for a few outliers.

Summing up, we do not claim this to be the best algorithm to solve this problem, but it is one that can be realistically accepted by the community, and, as next section shows it alleviates the unfair allocation of nodes.




Figure 13: People robustness after the node re-assignment process.

A. Algorithm Details

The first challenge is to define T , which can not be arbitrarily high, but depends on the network structure. To achieve this, the algorithm computes the robustness $R_G(\{v\})$ for every node v in the network and finds the node v_l with the lowest $R_G(\{v_l\})$. It is intuitive that if $v_l \in S_i$ then $R_G(S_i) \leq R_G(\{v_l\})$, so initially we set $T = R_G(\{v_l\}) - 1$. If at the end of the execution a solution can not be found, T is decremented and the algorithm is run again. Another important observation is that when computing $R_G(S_i)$ the contribution of leaf nodes must be omitted. To understand this let L_i be the number of leaf nodes owned by node i , with $L_i \subseteq S_i$. In the extreme case in which for the least robust owner i we have $L_i = S_i$ and thus $v_l \notin S_i$, the algorithm will re-assign all the nodes of S_i . It makes no sense to redistribute the ownership of a leaf node since its failure only affects the owner of the node, so, when comparing $R_G(S_i)$ with T , $R_G(S_i)$ is increased of the size of $\|L_i\|$.

Given these premises the algorithm in Listing 1 does the following:

- Lines 1-4: identify the owner i and the node n_l with the lowest robustness.
- Lines 5-7: define S_i and L_i .
- Line 9: start the re-assignment of non-leaf nodes.
- Line 15: pick a random person in $C(i)$, with a bias on close-by people.
- Line 18-19: if no person can be chosen, break (jumps to line 38).
- Line 20-24: test if re-assignment is feasible. If the new owner after re-assignment has robustness below threshold, blacklist him/her.
- Line 26-34: check if after the re-assignment of v , $R_G(S_i)$ is still below the threshold. If not, exit from main loop.
- Line 38-42: if the re-assigning process in unsuccessful, decrement the threshold, undo changes and loop again.

The main loop starting at line 9 is executed once every time T is decremented, so at most $T < \|V\|$ times. The loop starting at line 11 runs at most $\|S_i\| < \|S\|$ times, while the inner loop starting at line 14 runs at most $\|C(i)\| < \|V\|$

```

1  R_o = sort_owners_robustness()
2  # returns a sorted list of R(C)
3  R_n = sort_nodes_robustness()
4  # returns a sorted list of (node, R(node))
5  T = R_n[0] # set T to the lowest node robustness
6  least_robust_owner = R_o[0][0]
7  sorted_owned_nodes =
8      get_sorted_nodes_by_owner(least_robust_owner)
9  # returns a list of nodes for an owner,
10 # sorted by their robustness
11 leaf_nodes = get_leaf_nodes(sorted_owned_nodes)
12 exit_loop = False
13 while not exit_loop: # main loop
14     reassigned_nodes = []
15     for node in sorted_owned_nodes.remove(leaf_nodes):
16         # loop on non-leaf nodes
17         black_list = []
18         while True: # inner while loop
19             new_friend = get_random_friend(node, black_list)
20             # return a random person in the community
21             # of the owner of node, excluding the black_list
22             if not new_friend:
23                 break # no one can receive this node.
24                 # break inner while loop
25             if not test_reassign(node, new_friend, T):
26                 # temporarily reassigning the node to new_friend, recompute
27                 # its robustness, return False if new_friend
28                 # is himself breaking the T, keep looping
29                 black_list.append(new_friend)
30             else:
31                 reassign_node(node, new_friend)
32                 # reassign the ownership to new_friend
33                 reassigned_nodes.append(node)
34                 # keep track of reassigned nodes
35                 new_owned_nodes = sorted_owned_nodes.remove(node)
36                 new_robustness = compute_robustness(new_owned_nodes)
37                 # recompute the robustness
38                 if new_robustness + len(leaf_nodes) > T:
39                     exit_loop = True # will exit main loop
40                     break # exit inner while loop, stay in the for loop
41             if exit_loop:
42                 break # exit the for loop
43         # Failed to reassign nodes: must decrement T,
44         T.decrement(1)
45         # reset all done and stay in the main loop and try again
46         for node in reassigned_nodes:
47             reassign_node(node, least_robust_owner)
48     return reassigned_nodes

```

Listing 1: Re-Assigning Heuristic

times. The most complex operation in the loop is computing $R_G(S_i)$ that requires the computation of all Dijkstra’s trees which is an operation with polynomial complexity on $\|S\|$. The algorithm complexity thus remains polynomial on $\|V\|$ and $\|S\|$, and it instantly finds a solution for the ninux network on a standard PC.

For ninux, the algorithm produced the reassignment of 6 nodes from P_{top} to 6 different people, Fig. 13 reports the corresponding graph of $R_G(S_i)$ for the top-20 owners after reassignment and shows that the minimum robustness is strongly increased.

X. RELATED WORKS

Mesh networks have been a very active area of research in the first decade of the 2000s, with many implications regarding their performance, [13], routing [36], security [37], [38], support for mobility [39], [40]. In the recent years, this

technology matured, and the focus of research shifted to their application as CNs. CNs have been the subject of a series of works in the past years that had the goal of analysing their topological features [41] [21] [24] [8], their routing solutions [9] [42], and their social and management aspects [14] [43]. The only paper that deals with community networks and uses a similar approach is from Vega et. al. and analyses the Guifi.net community [44]. Guifi is probably the largest community network in the world, and the analysis of the mailing lists and interactions is hard to perform to the level of detail adopted in this paper. In a more controlled environment it is easier to draw solid conclusions on the techno-social dynamics of the CN, moreover, this paper uses the social analysis of the community to propose a way to remove the points of failure which is a completely new contribution. This paper extends what done in a previous work [45] in which the initial analysis on the ninux social network was done, but the spatial analysis and the re-assigning algorithm were not present.

XI. CONCLUSIONS

CNs are socio-technical networks spontaneously developed by communities of people. Some of the networks have a very clear social vision and propose a networking model different from commercial ISPs. At the light of the growing debate on network neutrality and network access, CNs represent a promising alternative and/or complementary model. The novelty of CNs lays both in the technical organization as a mesh network, and in the governance of the network, that is horizontal and participated on the model of Peer Production platforms. Nevertheless, a community organized around a distributed network and open social networking instruments can anyway develop in an unbalanced way and hide single points of failure. The paper showed that the network is spatially hierarchical, meaning that, even if realized without planning, it tends to be organized in a tree-like structure that favours “vertical” communications compared to “horizontal” communications. While this is a characteristic of ninux, we also showed that another network (FFWien) shows a more robust organization, thus, the spatial characteristics of a CN are largely influenced by the external conditions. We also showed that ninux let one person become the owner of many critical nodes and a broker in the social network. Thus, spontaneity must be helped with a set of instruments that let the community understand the direction it is taking, in order to avoid pitfalls. CNs, that do their best to grow with the limited resources that they voluntarily share (and achieve excellent results) can embed the metrics and the methodology presented in this paper in the instruments they use to monitor the network status, and detect the emergence of problematic situations before they become critical. This work is actually undergoing together with the ninux community⁸

ACKNOWLEDGEMENT

This work was financed partially by the European Commission, H2020-ICT-2015 Programme, Grant Number 688768 ‘netCommons’ (Network Infrastructure as Commons).

⁸See the netCommons Project results: <https://www.netcommons.eu/?q=content/monitoring-cns-report-experimentations-cns-v2>

REFERENCES

- [1] P. Baran, "On distributed communications networks," *IEEE transactions on Communications Systems*, vol. 12, no. 1, pp. 1–9, March 1964.
- [2] L. Navarro, R. Baig, F. Freitag, E. Dimogerontakis, F. Treguer, M. Dulong de Rosnay, L. Maccari, P. Micholia, and P. Antoniadis, "Report on the Existing CNs and their Organization (v2)," Sept. 2016. [Online]. Available: <http://netcommons.eu/?q=content/report-existing-cn-and-their-organization-v2>
- [3] International Telecommunication Union (ITU), "Ict facts and figures 2017," 2017. [Online]. Available: <https://www.itu.int/en/ITU-D/Statistics/Pages/facts/default.aspx>
- [4] L. Belli and P. De Filippi, *Net Neutrality Compendium*, Springer, Ed., 2017.
- [5] P. De Filippi and F. Tréguer, "Expanding the Internet commons: The subversive potential of wireless community networks," *Journal of Peer Production*, no. 6, Jan. 2015.
- [6] C. Fuchs, "Sustainability and community networks," *Telematics and Informatics*, vol. 34, no. 2, pp. 628 – 639, May 2017.
- [7] S. Crabu and P. Magaudda, "Bottom-up infrastructures: Aligning politics and technology in building a wireless community network," *Computer Supported Cooperative Work (CSCW)*, vol. 27, no. 2, pp. 149–176, 2018.
- [8] D. Vega, R. Baig, L. Cerda-Alabern, E. Medina, R. Meseguer, and L. Navarro, "A technological overview of the guifi.net community network," *Computer Networks*, vol. 93, pp. 260–278, Dec 2015.
- [9] C. Barz, C. Fuchs, J. Kirchhoff, J. Niewiejska, and H. Rogge, "OL-SRV2 for Community Networks: Using Directional Airtime Metric with external radios," *Computer Networks*, vol. 93, Part 2, pp. 324–341, Dec. 2015.
- [10] L. Baldesi, L. Maccari, and R. Lo Cigno, "Improving P2P streaming in wireless community networks," *Computer Networks*, vol. 93, Part 2, pp. 389 – 403, Dec. 2015.
- [11] F. Hao, G. Min, J. Chen, F. Wang, M. Lin, C. Luo, and L. T. Yang, "An Optimized Computational Model for Multi-Community-Cloud Social Collaboration," *IEEE Transactions on Services Computing*, vol. 7, no. 3, pp. 346–358, Jul. 2014.
- [12] J. Jiménez, R. Baig, P. Escrich, A. Khan, F. Freitag, L. Navarro, E. Pietrosemoli, M. Zennaro, A. Payberah, and V. Vlassov, "Supporting cloud deployment in the Guifi.net community network," in *IEEE Global Information Infrastructure Symposium*, 2013.
- [13] I. F. Akyildiz and X. Wang, "A survey on wireless mesh networks," *IEEE Communications Magazine*, vol. 43, no. 9, pp. S23–S30, Sept. 2005.
- [14] R. Baig, R. Roca, L. Navarro, and F. Freitag, "Guifi.Net: A Network Infrastructure Commons," in *ACM International Conference on Information and Communication Technologies and Development, ACMDev*, 2015.
- [15] S. Jain and D. Agrawal, "Wireless community networks," *Computer*, vol. 36, no. 8, pp. 90–92, Aug. 2003.
- [16] B. Braem, C. Blondia, C. Barz, H. Rogge, F. Freitag, L. Navarro, J. Bonicioli, S. Papataniasiou, P. Escrich, R. Baig Vias, A. L. Kaplan, A. Neumann, I. Vilata i Balaguer, B. Tatum, and M. Matson, "A Case for Research with and on Community Networks," *ACM SIGCOMM Comput. Commun. Rev.*, vol. 43, no. 3, pp. 68–73, Jul. 2013.
- [17] Y. Benkler and H. Nissenbaum, "Commons-based peer production and virtue," *Journal of Political Philosophy*, vol. 14, no. 4, pp. 394–419, Nov. 2006.
- [18] S. Crabu, F. Giovannella, L. Maccari, and P. Magaudda, "A Transdisciplinary Gaze on Wireless Community Networks," *TECNOSCIENZA: Italian Journal of Science & Technology Studies*, vol. 6, no. 2, pp. 113–134, Jan. 2016.
- [19] J. Sderberg, "Free Space Optics in the Czech Wireless Community: Shedding Some Light on the Role of Normativity for User-Initiated Innovations," *Science, Technology & Human Values*, vol. 36, no. 4, pp. 423–450, Jan. 2011.
- [20] I. Verhoeven, T. Metz, and T. Van de Wijdeven, "Do-ocracy's democratic anchorage," in *Systematising Comparison of Democratic Innovations: Advanced explanations of the emergence, sustenance and failure of participatory institutions, 42nd ECPR Joint Sessions*, Apr. 2016.
- [21] L. Maccari, "An analysis of the Ninux wireless community network," in *International Workshop on Community Networks and Bottom-up Broadband (CNBuB)*, 2013.
- [22] R. Albert, H. Jeong, and A.-L. Barabasi, "Error and attack tolerance of complex networks," *Nature*, vol. 406, no. 6794, pp. 378–382, Jul. 2000.
- [23] A. Shaw and B. Mako Hill, "Laboratories of oligarchy? how the iron law extends to peer production," *Journal of Communication*, vol. 64, no. 2, pp. 215–238, Apr. 2014.
- [24] L. Maccari and R. L. Cigno, "A week in the life of three large wireless community networks," *Ad Hoc Networks*, vol. 24, Part B, no. 0, pp. 175 – 190, Jan. 2015.
- [25] C. Bird, A. Gourley, P. Devanbu, M. Gertz, and A. Swaminathan, "Mining Email Social Networks," in *ACM International Workshop on Mining Software Repositories*, 2006.
- [26] M. Barthlemy, "Spatial networks," *Physics Reports*, vol. 499, no. 1, pp. 1–101, Feb. 2011.
- [27] R. Louf, P. Jensen, and M. Barthelemy, "Emergence of hierarchy in cost-driven growth of spatial networks," *Proceedings of the National Academy of Sciences (PNAS)*, vol. 110, no. 22, pp. 8824–8829, May 2013.
- [28] A.-L. Barabási and R. Albert, "Emergence of scaling in random networks," *Science*, vol. 286, no. 5439, pp. 509–512, Oct. 1999.
- [29] M. G. Everett and S. P. Borgatti, "The centrality of groups and classes," *The Journal of mathematical sociology*, vol. 23, no. 3, pp. 181–201, Aug. 1999.
- [30] A.-L. Barabási, *Network science*. Cambridge University Press, 2016.
- [31] M. Abd-El-Barr, "Topological network design: A survey," *Journal of Network and Computer Applications*, vol. 32, no. 3, pp. 501–509, May 2009.
- [32] S. L. Toral, M. R. Martínez-Torres, and F. Barrero, "Analysis of virtual communities supporting OSS projects using social network analysis," *Information and Software Technology*, vol. 52, no. 3, pp. 296–303, Mar. 2010.
- [33] V. D. Blondel, J.-L. Guillaume, R. Lambiotte, and E. Lefebvre, "Fast unfolding of communities in large networks," *Journal of statistical mechanics: theory and experiment*, vol. 2008, no. 10, Oct.
- [34] A. Lancichinetti, S. Fortunato, and F. Radicchi, "Benchmark graphs for testing community detection algorithms," *Physical review E*, vol. 78, no. 4, p. 046110, 2008.
- [35] L. Li, D. Alderson, W. Willinger, and J. Doyle, "A First-principles Approach to Understanding the Internet's Router-level Topology," in *Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications*, ser. SIGCOMM '04. New York, NY, USA: ACM, 2004.
- [36] E. Alotaibi and B. Mukherjee, "A survey on routing algorithms for wireless ad-hoc and mesh networks," *Computer Networks*, vol. 56, no. 2, pp. 940 – 965, 2012.
- [37] S. M. S and C. Seon, "Security issues in wireless mesh networks," in *2007 International Conference on Multimedia and Ubiquitous Engineering (MUE'07)*, April 2007, pp. 717–722.
- [38] R. Fantacci, P. Neira Ayuso, L. Maccari, and R. Martínez Gasca, "Efficient packet filtering in wireless ad-hoc networks," *IEEE Communication Magazine*, vol. 46, no. 2, pp. 104–110, 2008.
- [39] J. Xie and X. Wang, "A survey of mobility management in hybrid wireless mesh networks," *IEEE Network*, vol. 22, no. 6, pp. 34–40, November 2008.
- [40] R. Fantacci, L. Maccari, T. Pecorella, and F. Frosali, "Analysis of secure handover for ieee 802.1x-based wireless ad hoc networks," vol. 14, no. 5, pp. 21–29, 2007.
- [41] L. Cerda-Alabern, "On the topology characterization of Guifi.net," in *IEEE 8th International Conference on Wireless and Mobile Computing, Networking and Communications (WiMob)*, Oct. 2012.
- [42] L. Cerda-Alabern, A. Neumann, and L. Maccari, "Experimental Evaluation of BMX6 Routing Metrics in a 802.11an Wireless-Community Mesh Network," in *3rd International Conference on Future Internet of Things and Cloud (FiCloud)*, 2015.
- [43] J. Kos, M. Milutinovic, and L. Cehovin, "nodewatcher: A substrate for growing your own community network," *Computer Networks*, vol. 93, pp. 279–296, Dec. 2015.
- [44] D. Vega, R. Meseguer, and F. Freitag, "Analysis of the Social Effort in Multiplex Participatory Networks," in *International Conference on Grid Economics and Business Models*, 2014.
- [45] L. Maccari, "On the Technical and Social Structure of Community Networks," in *IFIP Internet of People Workshop,IoP*, 2016.

APPENDIX A
DETAILS ON THE DATA-SET

All the source code realized for the paper is released as open source and is available online⁹.

The data-set for the ownership analysis in Section V was derived from a number of sources. The first is the data-set obtained and published in a previous work [24] that describes the ninux network in early 2014. This is a set of topology dumps extracted from the ninux mapserver and integrated with the metrics exported by the OLSR protocol in the same period. From the thousands of graphs collected, here we analyse the one with the highest number of geo-located nodes.

APPENDIX B
SPATIAL SEPARATION

The spatial separation metrics were partly modified from their original definition ([27]) due to the different context in which we use them. The first modification is due to the fact that CNs are not tree-shaped but they are undirected graphs (a functional Wi-Fi link must be able to transmit in both directions, so there is no need to use directed edges to represent the network graph, albeit, link performance can be asymmetric). As such, in the ninux graph there is no root node that is naturally identifiable. We repeated the measures with three choices for the root node, selecting the node that maximises the following metrics: closeness centrality, betweenness centrality, and eccentricity. All the choices yield qualitatively similar results even if they identify different root nodes. Results are reported in Table VI.

The second difference is in the definition of Eq. (1). In a tree v_i has by definition only a neighbor with a lower level (its parent in the tree) and all the other neighbors have a higher level (its descendants in the tree). In a graph, among the neighbors there can be also nodes with the same level, and it is important to consider these links, or else separation is artificially increased. Thus we modified 1 as follows:

$$N'(i) = \{v_j \mid v_j \in N(i) \wedge l(v_j) \geq l(v_i)\}. \quad (9)$$

Finally, we modified the original definition of the influence zone in order to better reflect the behaviour of a CN. In the original definition the zone for node i is defined as “the circle centered on the barycenter of i ’s neighbours that belong to the next level, of radius the maximum distance between the barycenter and those points” [27]. In CN long links are realized with directional antennas, thus, there is not a well-defined concept or “radius” of the influence zone. Consider

⁹See the fromdiff repository <https://github.com/leonardomaccari/fromdiff> for the email-parsing function, and the fairgraph repository <https://github.com/leonardomaccari/fairgraph> for the implementation of the algorithm in Section IX. Some of the network analysis functions are included in https://github.com/leonardomaccari/community_networks_analysis and will be better documented in the future. Finally, the code for the spatial analysis can be found at this DOI 10.5281/zenodo.1218746. The topology files are published at this link: <https://ans.disi.unitn.it/redmine/projects/ninux-temporal-evolution-analysis/wiki>, the position of the nodes are not public, they were disclosed by the community for this research but since they identify the location of private houses they are not going to be released to the public. The geo-located data of the FFWien network are published by the community at this URL: <https://map.funkfeuer.at/wien/data.php>

level	s_l (e)	s_l (c)	s_l (b)
0	-	-	-
1	-	0.93	-
2	0.93	0.86	0.93
3	0.95	0.99	0.95
4	0.90	0.88	0.90
5	1.00	-	1.00
6	0.80	0.99	0.80
7	-	-	1.0

Table VI: Average separation within each level with various choices of root node. Dash means that there are less than 2 zones in the level (e: eccentricity; c: closeness; b:betweenness).




Figure 14: An example definition of influence zone.

Fig. 14 in which the node for which we compute the interest zone is the red one (v_i), and the yellow nodes are the nodes in $N'(i)$. The dashed circle is the influence zone as per the original definition, and the dotted polygon is the hull envelop of the points. The area marked with the letter B and C are areas for which there is no assurance that v_i has any coverage, since there is potentially no antenna pointed in their direction. Area A extends beyond an existing node, and in that direction the line of sight may be obstructed by obstacles. Consequently, for its application to CNs, the convex hull of the area including $N'(i)$ is a more realistic choice than the original definition.



Leonardo Maccari (M) received a Ph.D in the Faculty of Computer Science Engineering in the University of Florence in 2010 and he is currently an Assistant Professor at the Department of Computer Science and Information Engineering of the University of Trento (Italy). He has been awarded with a Marie Curie grant for the PAF-FPE project for the period 2011-2014, he is the Work-Package coordinator of the netCommons.eu H2020 project. He is an IEEE and ACM member and co-authored more than 50 publications in refereed conferences, journals and book chapters. He is also among the authors of European and US patents. Among his research interest there are network protocols and privacy in large-scale wireless mesh networks with special focus on Community Networks.