

Measuring Centralization of Online Platforms Through Size and Interconnection of Communities

Milo Z. Trujillo
University of Vermont
milo.trujillo@uvm.edu

Laurent Hébert-Dufresne
University of Vermont
laurent.hebert-dufresne@uvm.edu

James Bagrow
University of Vermont
james.bagrow@uvm.edu

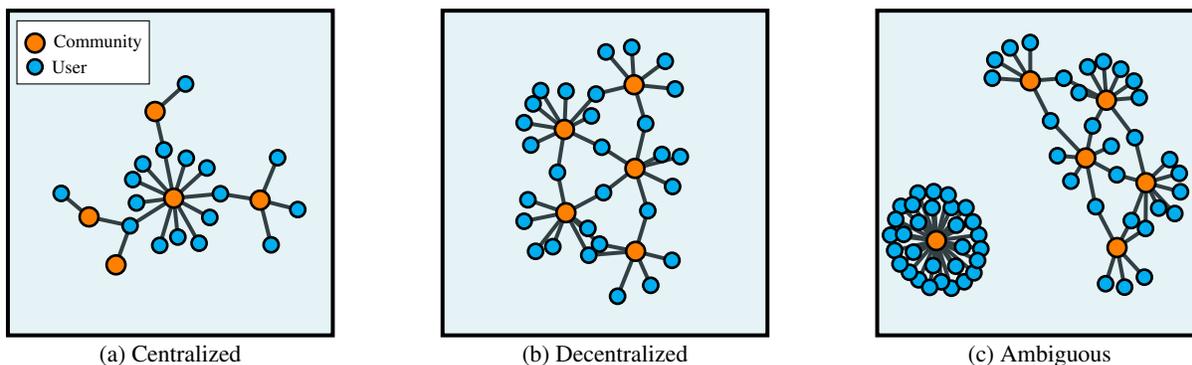


Figure 1: The influence of a community is tied to both its size and topological role in a network. In the centralized network, the orange community at the center both has the largest population of blue users, and serves as a bridge between four other communities. In the decentralized example, communities are of variable size, but none have a pivotal position to influence their peers. In the ambiguous case, one community is much larger, but the remaining network matches the “decentralized” example. Neither a distribution of community sizes nor purely structural measurements like betweenness centrality or graph conductance adequately capture this notion of community-level influence.

Abstract

Decentralized architecture offers a robust and flexible structure for online platforms, since centralized moderation and computation can be easy to disrupt with targeted attacks. However, a platform offering a decentralized architecture does not guarantee that users will use it in a decentralized way, and measuring the centralization of socio-technical networks is not an easy task. In this paper we introduce a method of characterizing community influence in terms of how many edges between communities would be disrupted by a community’s removal. Our approach provides a careful definition of “centralization” appropriate in bipartite user-community socio-technical networks, and demonstrates the inadequacy of more trivial methods for interrogating centralization such as examining the distribution of community sizes. We use this method to compare the structure of multiple socio-technical platforms – Mastodon, git code hosting servers, BitChute, Usenet, and Voat – and find a range of structures, from interconnected but decentralized git servers to an effectively centralized use of Mastodon servers, as well as multiscale hybrid network structures of dis-

connected Voat subverses. As the ecosystem of socio-technical platforms diversifies, it becomes critical to not solely focus on the underlying technologies but also consider the structure of how users interact through the technical infrastructure.

1 Introduction

Online social spaces are vulnerable to centralized authorities making decisions that negatively affect the community. In 2022, the Software Freedom Conservancy recommended that all developers migrate their projects away from GitHub [1], after Microsoft bought the software development collaboration platform and used open source projects as training data for their commercial CoPilot software, in violation of open source licenses and community standards. The same year, users and advertisers departed Twitter after its purchase by Elon Musk and subsequent changes in community policy and staffing, including firing content moderators [2] and reinstating a number of accounts banned for violating the platform’s hateful content and harassment policies [3]. Reddit moderators have historically engaged in black-

outs to protest administrative policies [4], and these trends are ongoing; in June, 2023, Reddit announced plans to begin charging for API access, sparking warnings from scientists [5], outrage among users, and a protest across nearly 9,000 subreddits, the long-term effects of which remain to be seen. As users express dissatisfaction with platform administrators, they have sought alternative platforms without centralized control, leading to the rapid growth of “federated” platforms like Mastodon [6] and Bluesky¹. Alternatively, other users have promoted self-hosted platforms, such as independently operated git servers, or peer-to-peer hosting solutions such as the Interplanetary File System (IPFS) or web-torrent video hosting software PeerTube. Some deplatformed users have also responded by creating close facsimiles of existing centralized platforms with extremely permissive content-policies, frequently called “alt-tech” platforms [7].

What exactly is “centralization” in an online social network? Does it describe ownership of the platform? Its technical infrastructure? The creation and enforcement of community norms? The distribution of activity and reach of content producers? Centralization has long been ill-defined by academics [8], and “decentralization” joins as a widely-used but contextually redefined term today [9]. Of particular interest to us is a notion of group social influence: How much does one community impact information flow across a platform? For example, how independent are subreddits on Reddit, and how closely interlinked are instances on Mastodon, the nascent “decentralized Twitter alternative?” Our goal is to measure the influence of a platform’s sub-communities on their peers by looking at the structure of socio-technical networks, providing a mesoscale metric to quantify centralization at an inter-group level.

One common approach to measuring community-level centralization is through community size-distribution. If a small oligarchy of Mastodon instances dwarf the population sizes of their peers, then one could presume that the platform is centralized around these instances. Indeed, several prior studies on Mastodon use community size disparity as a starting point, or presuppose that the largest instances are the most significant and focus their study on the largest communities [10]–[13]. While the community size distribution is related to centralization, assuming they are the same precludes the possibility that a collection of many smaller instances may be more influential than the few largest, or that the largest instances may not represent the platform as a whole.

We reject the assertion that the largest communities must be the most significant, or that their size alone implies centralization, on the grounds that community size

¹Bluesky is still in beta, and while the protocol is federated, only one instance exists at the time of writing.

does not correlate with the number of cross-community links in observed real-world networks. In fact, our results show multiple platforms where the largest communities are *not* well integrated with the platform as a whole (discussed in Sec. 4.1, especially Fig. 4), allowing a more decentralized network of communities to exist outside of the largest groups. Under this view, the largest communities would be the most significant only when they also act as important information bottlenecks for the entire system.

To illustrate this discrepancy, consider Fig. 1. In Fig. 1a the largest community serves as a central hub, connecting several smaller communities together through shared membership. In Fig. 1b, community size is normally distributed, and no community has a pivotal role as a bridge between its peers. Community size-distribution and graph-centric metrics like betweenness-centrality would agree that the former network is centralized, while the latter is decentralized. However, Fig. 1c presents a more complex scenario: the community size distribution is highly unbalanced, but the largest community has no impact on the remainder of the network. Adding one edge to create a complete graph would grant the largest community a high betweenness-centrality because of its pivotal role in connecting so many users to the rest of the graph; but this does not match our intuition that the largest community has a small role in the rest of the network.

We propose a definition of centralization meant to capture the alignment between rankings of community size and information bottlenecks. To do so, we combine theoretical ideas from graph theory on bottlenecks and applied concepts from network science about network resilience. Our metric then measures how removing a community would impact users within remaining communities, based on the number of “bridges” between communities. We study a variety of real and simulated networks with this method to examine platform behavior under a range of conditions, and we compare our metric to existing measurements of centralization and network “bottlenecks.” Finally, we discuss how this work contributes to broader discussions of centralization online, and how techniques like ours can be extended with richer interaction data.

2 Prior Work

In the context of online platforms, centralization is sometimes defined in terms of decision-making power, or who has the authority to make what kinds of decisions about the use of the platform. This definition can be traced to Elinor Ostrom’s work on Institutional Analysis and Development [14], which describes “layers” of decisions, from operational rules (elementary actions any user can perform), to collective rules (the context in which users

operate and interact, such as the Twitter feed or the Amazon marketplace), to constitutional rules (the “meta” rules through which the system changes itself). Modern research on platform design often assesses who has decision-making power, and what levers of change are available to different categories of participants [15], [16].

While qualitative studies examine power structures through analyzing governance and rule sets [17], [18], network science infers structure through the observed interactions between humans [8], [19]. We quantify centralization using attributes that fall into three categories: vertex-level attributes, cluster-level attributes, and graph-level attributes. Vertex-level attributes like betweenness centrality [8] or eigenvector centrality [20] measure the prominence of a particular node in terms of how well it is connected to its peers, or how many paths flow through the node. Cluster-level attributes describe groups of vertices, such as the size of the population that contains a particular attribute, or the assortativity describing how likely vertices with a particular attribute are to be connected to one another. Graph-level attributes describe aspects that span the entire network, including diameter, density, and graph conductance [21]. Quantifiability should not be conflated with objectivity; the modeling choice of what entities are included as vertices and what relationships are represented as edges or attributes presupposes what can be considered influential or centralized [22].

Another thread of research tries to join the social theory of centralization and graph theoretical metrics. Fan and Zhang [18] distinguish between the technical underpinnings of a network and its social layers, focusing on community-run moderation in infrastructure-centralized (Slack, Discord) and self-hosted (Minecraft) services. Prior Mastodon research also bridges this gap, including both geographic and data-center distribution of instances [10], important for understanding resiliency to disruption or power-outage. This approach aligns with notions of network robustness where centralization can be measured by how a network breaks down under targeted pruning of central nodes [23]. Other studies on Mastodon also integrate its social interaction graph [13], important for understanding the influence of sub-communities and their administrators on discourse. Studies on the social structure of Mastodon primarily focus on individual-centralization, such as a “border-index” of what fraction of a user’s neighbors are on a foreign instance [12] and whether some users serve as critical bridges for information flow between instances [24], or community-centralization, such as how clustering coefficients differ between communities (instances) [11]. Our work intends to add to these options, by considering both a community-level centralization metric of how much influence one community has on the broader

platform, and a graph-level centralization score of how quickly a network deteriorates as its largest communities are removed, indicating how much it tends towards monopoly or oligopoly.

Finally, authors like Agre [25] push back against a Boolean or spectrum from “centralized” to “decentralized,” and view platforms as layers of centralization and decentralization. Their primary example is eBay, a centrally-controlled marketplace operated by one company, but containing a decentralized network of buyers and sellers. Other studies have built off of this idea, proposing the creation of decentralized “digital juries” on infrastructure-centralized platforms [18], or creating democratic governance bots with flexible bylaws overlaid on centrally-hosted chat services like Slack [26]. We pose that this layered framework also applies to thinking about complex networks like Fig. 1c, which have both centralized and decentralized components that should be considered independently.

3 Methods and Materials

We introduce our metric in Sec. 3.1, and two data sets: five real world networks that encompass a breadth of configurations (Sec. 3.2), and a set of common synthetic networks for reference (Sec. 3.3).

3.1 Measuring centralization: disruption curves

Prior studies on centralization of social networks often focus on graph-level attributes such as detecting components, the size of the giant component, modularity, density, degree distribution [27]. Others may use “bottleneck” metrics like graph conductance [21] to identify bridges and key clusters. These metrics are most appealing in unipartite settings where the structure of the network is not prescribed. However, we focus on bipartite graphs where communities are well defined, such as subreddits, Mastodon instances, or newsgroups. In these contexts, we are not attempting to infer the number or boundaries of communities, but to measure how influential the known communities are on their neighbors. The size distribution of communities tells us how large a subgroup is, but does not capture the overlap between communities. A graph-wide modularity score describes how well-partitioned the graph is into clusters, and so approximates how insular communities are, but cannot provide more nuance as to whether the largest communities are more integrated than smaller ones, whether small communities are well connected to larger peers but not to each other, or other topological features.

We propose that the influence of a community should be measured in terms of how users outside the community would be impacted by its absence. In other words, a community’s influence should be proportional not to its

size, but to the number of bridges between it and other communities. Or, in graph theoretic terms, what fraction of edges would be cut by removing a community, not counting users that do not participate outside the community. More succinctly, “what percentage of edges from surviving vertices would be cut by removing a community?”

We measure disruption cumulatively, rather than discretely per-community. This allows us to answer questions like “how influential are the largest three communities on the rest of the platform?” Since “oligarchies” of large and densely interconnected communities may be common, a cumulative metric is more useful than measuring the influence of a single community on the rest of the oligarchy.

Formally, we define a set of communities that are being cut, A , with associated edges $|A|$. Each user has a set of edges to one or more communities. If users *only* have edges to communities in A , then the user is removed along with A . Surviving users with an edge to at least one remaining community are denoted S , with total edges $|S|$, and edges to cut communities in A denoted ∂S . The disruption curve is calculated as $\partial S/|S|$. This notation was chosen for its similarity to the Cheeger number [21], stressing how our metric measures the alignment of community size and information bottleneck. We additionally outline the algorithm as pseudocode in supplemental Sec. S1.

Our disruption curve metric is intended for bipartite networks, where communities are clearly distinguishable and users can participate in multiple communities. However, some consideration is also given to applying our metric to unipartite settings in Sec. S2.

We plot disruption similarly to a cumulative disruption function (CDF), where the x-axis represents the number of communities removed, cumulatively ordered by degree, and the y-axis represents the fraction of edges from surviving users that have been cut. In other words, the x-axis is the size of A as a fraction of all communities in the graph, and the y-axis is $\partial S/|S|$, where both the numerator and denominator are dependent on $|A|$.

While disruption curves offer insight into the role of the largest communities on a platform, some readers may desire a scalar summary statistic to describe how “centralized” a platform is under our metric. For these scenarios we recommend calculating the area under the curve, as shown in Figs. 2b and 3b. The Disruption AUC (DAUC) does not indicate how much any particular community influences its peers, but summarizes whether a network is prone to disruption if its largest communities are removed. Methodological choices for calculating the DAUC are discussed in Sec. S3.

3.2 Real-World Network Data

We analyze five real-world datasets, each describing online social interactions in bipartite configurations where vertices represent either “users” or “communities.” We utilize a 2021 scrape of the Mastodon follow graph [11]. Mastodon is a Twitter alternative where users are located on one of thousands of “instances,” which are Twitter-like servers with their own administrators and content policies. However, Mastodon users can follow users on other instances, exchanging content between the two communities, so long as the servers are “federated” (willing to exchange content). For a second example of a platform with distributed servers, we include the Penumbra of open-source [28], a data set of independent git servers (not GitHub or GitLab), and users that contribute to repositories on each server. We also include an interaction graph from BitChute [29], an alt-tech YouTube alternative, consisting of users and the channels (video uploaders) whose videos they commented on. We utilize a similar scrape of Voat [30], an alt-tech Reddit alternative active until late 2020, consisting of users and the “subverses” (subreddits) they commented in. We additionally include an archive of Polish Usenet groups [31], providing a much older but similarly structured platform for comparison. Details on the vertex and edge definitions for each network are included in Table 1, and the size of each network is listed in Table 2.

We selected these platforms because they have clear bipartite user and community representations, their data is readily available, and each platform is small enough to obtain a nearly-complete sample. Sub-sampling a larger platform like Reddit is likely to miss lower-population or lower-activity sub-communities, and we are particularly interested in the interactions between smaller communities. The resulting dataset encompasses a variety of approaches to hosting and community governance, providing a spectrum of “centralization.”

3.3 Synthetic Network Data

To understand disruption curves and contextualize our real-world results, we examine a variety of well understood synthetic network topologies.

First we construct a bipartite star network, as a default example of a network centralized around a single hub. In our example plots, we construct a graph with 150 communities and 3000 users, such that every user has an edge to two communities: the central hub, and one other, assigned uniformly. Removing the hub eliminates 50% of all edges, and removing any subsequent communities incurs no additional disruption, because all impacted users will have a degree of zero and be pruned from the graph (see Fig. 3a). This graph type is therefore highly centralized but has a decentralized periphery after the removal of the central community, illustrating

Platform	Community Definition	Edge Definition	Edge weight
Mastodon	Mastodon Instances	Between each user and every instance on which they follow users	The number of users followed on an instance
Penumbra	A git server	Between a user (identified by email) and each server on which they have contributed to a repository	The number of repositories committed to on each server
BitChute	BitChute channels	Between each user and every channel they have commented on videos from	The number of comments made
Voat	A Voat “subverse”	Between each user and subverses they’ve participated in	Number of comments made in a subverse
Usenet	A Usenet newsgroup	Between each user and every newsgroup they have posted in	The number of posts made

Table 1: Definitions of communities and edges for each platform examined

Platform	Comms.	Users	Edges
Mastodon	3,825	479,425	5,649,762
Penumbra	841	41,619	108,038
BitChute	29,686	299,735	11,549,058
Voat	7,515	3,624,486	16,263,309
Usenet	333	2,080,335	58,133,610

Table 2: Population size of each network in terms of community count, user count, and relationship edge count, before compressing duplicate edges into weighted edges

how different topologies can co-exist in the same network, muddying the definition of “centralization.”

We then test disruption on a variety of bipartite networks with power-law degree distributions. We first adapt the Barabási-Albert preferential attachment model to a bipartite setting, initializing a network with 300 empty communities and introducing users that connect to a given community with probability proportional to their size plus one. We also introduce a range of bipartite configuration models: in each, we assign a degree to each community drawn from a power law with a specified γ exponent. For each community, we create edges according to degree, connecting the community to users uniformly randomly without replacement. Therefore, we do not control for user degree or assortativity, only for the size of communities. Each of these networks produces a curve that slowly decays towards a diagonal, implying that removing the largest communities has some disproportionate impact, after which removing additional communities has a less pronounced result.

For a contrast from power-law degree distributions, we also adapt the Erdős-Rényi model to a bipartite setting, by creating vertices for communities and users, then creating all possible edges with a probability p (in our tests, $p = 0.05$), while preserving the bipartite constraint. These networks produce a disruption curve with a sec-

ond derivative near zero, indicating that most communities have near-equal influence on the population, and so removing the largest communities does not have a much larger impact than removing subsequent communities.

Lastly, we adapt the Watts-Strogatz small-world model to a bipartite setting. To accomplish this, we first produce a *unipartite* network with desired neighborhood size ($n = 5$) and edge density ($p = 0.05$) parameters. Then we apply a clustering algorithm (in our examples we have used weighted community label propagation) to place each user in one community. We create a vertex for each detected community, and replace all user-user edges with user-community edges. This process is described in greater detail in Sec. S2 and an example is illustrated in Fig. S1. These networks have the most uniform community size distribution of any we tested, and their disruption curves are similar to those of Erdős-Rényi networks, with slightly more variability.

For Erdős-Rényi and configuration model networks, it is possible to create a closed-form solution for cumulative disruption. This is not our emphasis, because we are primarily concerned with analyzing real-world networks for which network generating functions are not available, but we detail this formalism in supplemental Sec. S5. We also include further descriptions and reference visualizations for bipartite near-star and bipartite power-law networks in Sec. S4.

4 Results

We plot the cumulative population size, disruption curve, and disruption AUC for real-world networks in Fig. 2, and plot the same results for synthetic network data in Fig. 3. We first focus on discrepancies between the size distribution and disruption curves for real networks in Sec. 4.1, then return attention to synthetic network data when we examine the role of assortativity in Sec. 4.2.

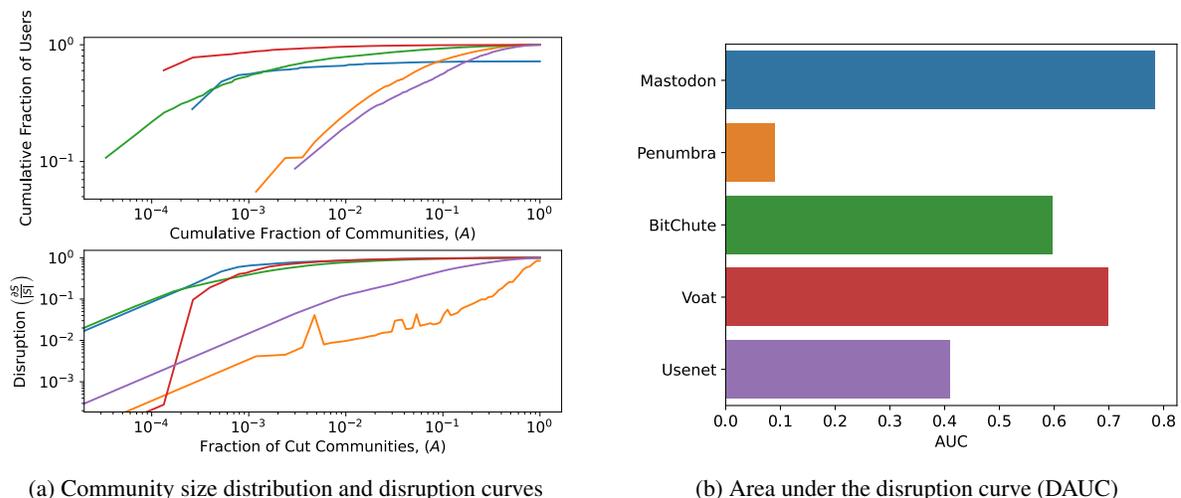


Figure 2: Summary measures of centralization. (a) Our measure of community disruption (bottom) does not correlate with the population distribution of communities (top). (b) The area under the disruption curve (DAUC) provides a summary statistic of the disruption curve that reinforces how network structure combined with community size provide greater insight into centralization (measurement details in Sec. S3). Here, panel (a) consists of cumulative distribution plots of population and disruption, where the top subplot is a CDF of the platform population as smaller communities are included, and the bottom subplot shows how networks are damaged as more of the largest communities are removed. Each line represents a different network, using the color key from panel b.

4.1 Comparison to Size Distribution

Upon comparing the size distribution and disruption curve in Fig. 2a, it is apparent that the community size distribution is insufficient to describe the structure of a network. Voat has the most skewed population distribution: almost all users participate in the largest community, yet the network does not experience significant disruption until the largest *three* communities are removed. Mastodon and BitChute have the next most skewed size distributions, but there is a large distance between the proportional sizes of their largest communities, and almost identical disruption curves as those communities are removed. By population distribution, the Penumbra appears to be more skewed towards its largest git servers than Usenet is towards its largest newsgroups. This is not mirrored in disruption curves, where Usenet has a consistently higher disruption than the Penumbra.

To explain these discrepancies, we examine each network in greater detail. Voat was a Reddit-like platform where users commented and posted in one or more “subverses.” While users chose to subscribe from among 7515 public subverses, new accounts were automatically subscribed to a set of 27 subverses by default. This “default subscription” has no parallel on other platforms we examined. Since these default subverses have an automatic population, they are more likely to receive engagement than subverses that must be discovered according to a user’s area of interest, and we may expect them to

be densely connected with most users on the platform. However, the largest two subverses on Voat by number of unique users were *not* default subverses; v/QRV was a QAnon conspiracy group, and v/8chan was a right-wing news and discussion forum whose name references the white supremacist imageboard 8chan (now “8kun”). Both subverses were highly insular, with little population overlap with the rest of the platform, as illustrated in Fig. 4. Therefore, it is only when we remove the *third-largest* subverse, v/news, that we see a large impact on remaining users on the site.

The Penumbra of open-source represents software development on git servers outside of GitHub and the primary GitLab instance. Each community represents a git servers with one or more public repositories, and edges indicate that a user (identified by email address) contributed to a repository on a server. Servers are often created per-organization; for example, the Debian Linux distribution hosts their own GitLab server at `salsa.debian.org`, and the University of Vermont hosts a GitLab server at `gitlab.uvm.edu`. Users often contribute to multiple repositories on a single server, but connections *between* servers are extremely sparse. This sparsity is responsible for the “spikes” in the Penumbra’s disruption curve; removing a git server may sever an edge to some users, and removing a second, related server may prune all remaining edges to those same users. When the cross-server collaborative users are removed, the im-

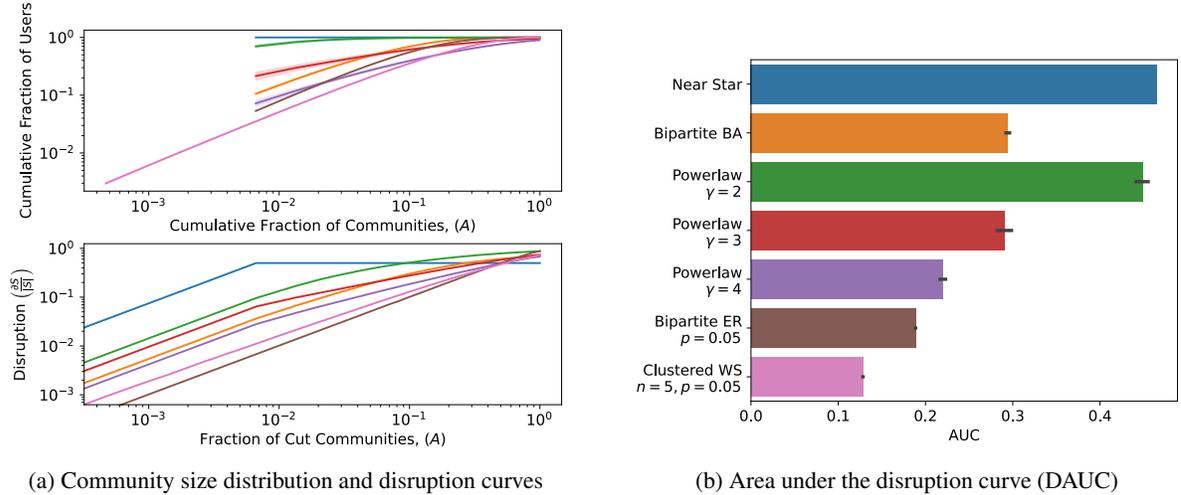


Figure 3: In simulated networks with a variety of degree distributions, the disruption curves for each network much more closely match the population distribution (Fig. 3a), suggesting that non-degree network attributes such as assortativity play a crucial role in determining centralization. As in Fig. 2, the left figure represents cumulative population and disruption as more communities are considered. Each line represents a network sharing the color-key in the right figure. Simulated networks were generated 100 times, and the mean and a 95% confidence interval are shown in both figures.

pact on the remaining less-collaborative community decreases. In all other networks enough users have a sufficiently high cross-community degree that disruption only increases as communities are removed.

Rather than examining the cumulative community size distribution, one could instead examine the size of the giant component of each network. The giant component will shrink as communities are cumulatively removed, providing another means of examining the influence of large communities. This is conceptually similar to our disruption metric, with a critical difference: inclusion of a community in the giant component is Boolean. If a community is completely insulated from the giant component of a network, then its removal will have no effect on the component size; otherwise the component will shrink by the size of the users removed with the community. Our metric provides greater flexibility for measuring how well integrated a community is among its peers. Measurement of the giant component size is discussed further in supplementary Sec. S7.

In summary, the community size distribution cannot adequately describe the topology of a platform because it does not account for features including assortativity or sparsity.

4.2 Assortativity and Centralization

We expect that degree assortativity (or degree-degree correlation) plays a significant role in the differences between observed community disruption (Fig. 2a) and network behavior under controlled degree distributions

(Fig. 3a). In a purely random setting, users are likely to have edges to multiple large communities, because most edge stubs in a configuration model come from high-degree communities. In real social settings, the content of communities may inhibit assortativity, as in Voat, where the largest two communities are highly insular (see Fig. 4), creating a large disparity between the community size distribution and disruption metric.

To explore this hypothesis, we randomly rewired each social network to increase assortativity. We select pairs of edges uniformly without replacement, and swap the communities of the edges if it would increase user-community degree assortativity. We continue this process until we have rewired a desired percentage of edges; if we exhaust the edge supply before finding sufficient valid swaps, we re-shuffle the edge list and continue drawing. For each rewired network we calculate its disruption and the area under the disruption curve, as in Fig. 2b, and plot the change in AUC during rewiring in Fig. 5.

This experiment is useful in distinguishing the idea of network centralization from classic ideas of monopoly. These are two different, but related, problems that are easy to confuse when focusing solely on summary statistics like community size distributions. When a network consists of disconnected communities, it is decentralized under the disruption metric regardless of the size distribution of these communities. This conclusion follows from our definition of centralization since removing a community in a sparse (or disconnected) network,

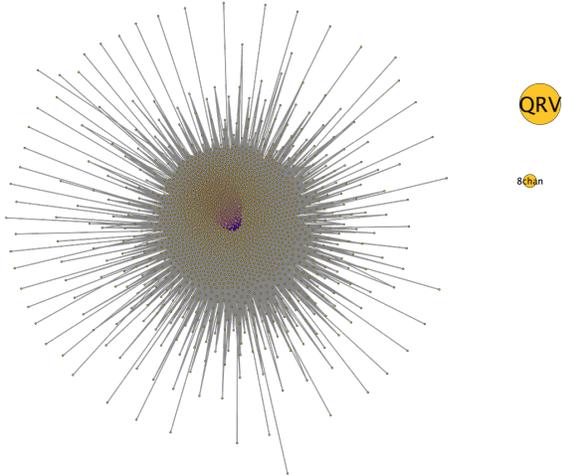


Figure 4: The two largest Voat communities (‘QRV’ and ‘8chan’) are dramatically larger than their peers, but have almost no overlap in population, making community size a poor proxy for platform-wide influence or centralization. In this network visualization, nodes represent Voat “subverses,” and edges represent at least thirty shared users active in two communities. Node size correlates with user count, and color correlates with strength; i.e. the level of overlap with neighboring communities. The purple communities at the center are default subverses all new users are subscribed to (“news,” “whatever,” etc), surrounding pink and orange communities are popular with lots of user overlap. The largest two communities, “QRV” and “8chan,” have almost no user overlap with other communities and are rendered to the right.

has little (or no) impact on other communities. This rewiring experiment highlights this logic: As networks get rewired to increase correlations, we increase the likelihood of having all the activity of a user focused on a single community and therefore progressively disconnect the community and decentralize the network. The only exception is Voat, whose initial state contains large disconnected communities that can get coupled to the rest of the network by rewiring, before being re-disconnected as we rewire more and more. Small correlations in large networks can therefore increase centralization, since large communities can broker more bridges when they contain well-connected users; while strong correlations in smaller networks can decrease centralization by focusing user activity on single communities.

We analyze other assortativity metrics in additional detail in the supplemental material, Sec. S6, as well as a comparison to the Cheeger number “bottleneck” metric [21], in Sec. S8. We also confirm our intuition about the role of weak correlations with mathematical analysis

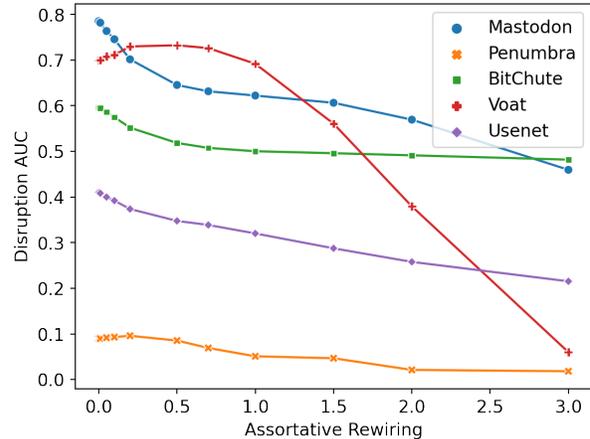


Figure 5: Increasing user-community degree assortativity through edge-rewiring increases the influence of the largest communities in highly insular (Voat) or sparse settings (Penumbra), but decreases disruption in all networks as increased rewirings eliminate cross-community edges and yield insular and sparse networks. Y-axis represents disruption AUC (see Fig. 2b), so that the slope shows change in disruption AUC as networks are rewired to increase user-community degree assortativity.

of random infinite networks in supplemental Sec. S5.

5 Conclusion and Future Work

We have added to the wealth of centralization metrics by proposing a mesoscale measurement that indicates how much influence one sub-community has over a broader network, by accounting for how many edges to remaining users would be severed if a community were removed. This metric allows us to differentiate between networks with a substantial community size-imbalance, and networks where the largest communities play a core structural role in their smaller peers. We extend our metric to create a graph-level measurement that indicates how “oligopic” a network is, or how well-integrated its largest communities are with the population at large.

We have utilized our disruption metric to examine a range of real-world social networks, comparing their network topology, distribution of community sizes, and the influence of those communities. We find that some platforms, like Voat, are much less centralized than their skewed community-size distribution would suggest, while others, like Usenet and the Penumbra of Open-Source, have similar size distributions and widely divergent disruption curves. Mastodon, while vocally supportive of decentralization, has a disruption curve mostly characterized by the skewed population distribution of its sub-communities and is therefore relatively centralized.

Using simulated networks with a range of degree dis-

tributions, and rewiring techniques to adjust assortativity, we have begun to explore the interplay between community size, structure, and community-level centralization. However, we limited ourselves to traditional network generative models like Erdős-Rényi and power-law configuration model networks. Future research could directly simulate networks with chimeric centralization which combine decentralized and centralized components to more realistically represent the diversity observed in social networks.

Our network representations are oversimplified in that we assume that each edge on a network represents a path of information flow. However, one user following another represents *potential* information flow; a bridge between two communities is only realized if the following user is online and chooses to propagate information from the edge to their own followers and instance.

More thorough research should examine how many potential bridges are utilized by, for example, monitoring the number of “boosts” (Mastodon’s equivalent to “retweets”) across instance boundaries on Mastodon. Observed information spread, and examining the reception of cross-pollinated ideas in non-originating communities, would provide much greater insight into how multi-community platforms function in practice.

Competing Interest Declaration

The authors declare no competing interest.

Data, Documentation, and Code Availability

Real-world networks used in this study are available from [11], and [28]–[31]. Code used in this research, including scripts for generating our synthetic networks, will be submitted to a public archive before publication.

Acknowledgements

All authors were supported by Google Open Source under the Open-Source Complex Ecosystems And Networks (OCEAN) project. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of Google Open Source.

References

- [1] D. Gingerich and B. M. Kuhn. “Give up github: The time has come!” (Jun. 30, 2022), [Online]. Available: <https://sfconservancy.org/blog/2022/jun/30/give-up-github-launch/> (visited on 07/19/2023).
- [2] K. Kupferschmidt, “As musk reshapes twitter, academics ponder taking flight,” *Science (New York, NY)*, vol. 378, no. 6620, pp. 583–584, 2022.
- [3] D. Hickey, M. Schmitz, D. Fessler, P. E. Smaldino, G. Muric, and K. Burghardt, “Auditing elon musk’s impact on hate speech and bots,” in *Proceedings of the International AAAI Conference on Web and Social Media*, vol. 17, 2023, pp. 1133–1137.
- [4] J. A. Jiang, P. Nie, J. R. Brubaker, and C. Fiesler, “A trade-off-centered framework of content moderation,” *ACM Trans. Comput.-Hum. Interact.*, vol. 30, no. 1, Mar. 2023, ISSN: 1073-0516. DOI: 10.1145/3534929. [Online]. Available: <https://doi.org/10.1145/3534929>.
- [5] B. I. Davidson, D. Wischerath, D. Racek, *et al.*, “Social media apis: A quiet threat to the advancement of science,” *PsyArXiv*, 2023.
- [6] H. B. Zia, J. He, A. Raman, I. Castro, N. Sastri, and G. Tyson, “Flocking to mastodon: Tracking the great twitter migration,” *arXiv preprint arXiv:2302.14294*, 2023.
- [7] J. Donovan, B. Lewis, and B. Friedberg, “Parallel ports: Sociotechnical change from the alt-right to alt-tech,” *en*, 2019. DOI: 10.25969/MEDIAREP/12374. [Online]. Available: <https://mediarep.org/handle/doc/13283>.
- [8] L. C. Freeman, “Centrality in social networks conceptual clarification,” *Social Networks*, vol. 1, no. 3, pp. 215–239, Jan. 1978, ISSN: 03788733. DOI: 10.1016/0378-8733(78)90021-7. [Online]. Available: <https://linkinghub.elsevier.com/retrieve/pii/0378873378900217> (visited on 08/08/2022).
- [9] G. Di Bona, A. Bracci, N. Perra, V. Latora, and A. Baronchelli, *The decentralized evolution of decentralization across fields: From governance to blockchain*, Number: arXiv:2207.14260, Jul. 28, 2022. arXiv: 2207.14260[physics]. [Online]. Available: <http://arxiv.org/abs/2207.14260> (visited on 08/08/2022).
- [10] A. Raman, S. Joglekar, E. D. Cristofaro, N. Sastri, and G. Tyson, “Challenges in the decentralised web: The mastodon case,” in *Proceedings of the Internet Measurement Conference*, Amsterdam Netherlands: ACM, Oct. 21, 2019, pp. 217–229, ISBN: 978-1-4503-6948-0. DOI: 10.1145/3355369.3355572. [Online]. Available: <https://dl.acm.org/doi/10.1145/3355369.3355572> (visited on 08/13/2021).
- [11] M. Zignani, S. Gaito, and G. P. Rossi, “Follow the “mastodon”: Structure and evolution of a decentralized online social network,” in *Proceedings of the International AAAI Conference on Web and Social Media*, vol. 12, 2018, pp. 541–550.

- [12] M. Zignani, C. Quadri, S. Gaito, H. Cherifi, and G. P. Rossi, “The footprints of a “mastodon”: How a decentralized architecture influences online social relationships,” in *IEEE INFOCOM 2019 - IEEE Conference on Computer Communications Workshops (INFOCOM WKSHPS)*, Apr. 2019, pp. 472–477. DOI: 10.1109/INFCOMW.2019.8845221.
- [13] L. La Cava, S. Greco, and A. Tagarelli, “Understanding the growth of the fediverse through the lens of mastodon,” *arXiv:2106.15473 [physics]*, Jun. 29, 2021. arXiv: 2106.15473. [Online]. Available: <http://arxiv.org/abs/2106.15473> (visited on 08/13/2021).
- [14] E. Ostrom, “Background on the institutional analysis and development framework: Ostrom: Institutional analysis and development framework,” *Policy Studies Journal*, vol. 39, no. 1, pp. 7–27, Feb. 2011, ISSN: 0190292X. DOI: 10.1111/j.1541-0072.2010.00394.x. [Online]. Available: <https://onlinelibrary.wiley.com/doi/10.1111/j.1541-0072.2010.00394.x> (visited on 02/05/2022).
- [15] R. E. Kraut and P. Resnick, *Building successful online communities: Evidence-based social design*. Mit Press, 2012.
- [16] S. Frey, P. M. Krafft, and B. C. Keegan, ““this place does what it was built for”: Designing digital institutions for participatory change,” *Proceedings of the ACM on Human-Computer Interaction*, vol. 3, pp. 1–31, CSCW Nov. 7, 2019, ISSN: 2573-0142. DOI: 10.1145/3359134. [Online]. Available: <https://dl.acm.org/doi/10.1145/3359134> (visited on 02/04/2022).
- [17] N. Schneider, P. De Filippi, S. Frey, J. Z. Tan, and A. X. Zhang, “Modular politics: Toward a governance layer for online communities,” *Proceedings of the ACM on Human-Computer Interaction*, vol. 5, pp. 1–26, CSCW1 Apr. 13, 2021, ISSN: 2573-0142. DOI: 10.1145/3449090. [Online]. Available: <https://dl.acm.org/doi/10.1145/3449090> (visited on 02/05/2022).
- [18] J. Fan and A. X. Zhang, “Digital juries: A civics-oriented approach to platform governance,” in *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, Honolulu HI USA: ACM, Apr. 21, 2020, pp. 1–14, ISBN: 978-1-4503-6708-0. DOI: 10.1145/3313831.3376293. [Online]. Available: <https://dl.acm.org/doi/10.1145/3313831.3376293> (visited on 02/05/2022).
- [19] M. Nadini, L. Alessandretti, F. Di Giacinto, M. Martino, L. M. Aiello, and A. Baronchelli, “Mapping the NFT revolution: Market trends, trade networks, and visual features,” *Scientific Reports*, vol. 11, no. 1, p. 20902, Oct. 22, 2021, ISSN: 2045-2322. DOI: 10.1038/s41598-021-00053-8. [Online]. Available: <https://doi.org/10.1038/s41598-021-00053-8>.
- [20] P. Bonacich, “Power and centrality: A family of measures,” *American Journal of Sociology*, vol. 92, no. 5, pp. 1170–1182, 1987.
- [21] B. Mohar, “Isoperimetric numbers of graphs,” *Journal of Combinatorial Theory, Series B*, vol. 47, no. 3, pp. 274–291, Dec. 1, 1989, ISSN: 0095-8956. DOI: 10.1016/0095-8956(89)90029-4. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/0095895689900294> (visited on 04/28/2023).
- [22] C. T. Butts, “Revisiting the foundations of network analysis,” *Science*, vol. 325, no. 5939, pp. 414–416, 2009.
- [23] R. Albert, H. Jeong, and A.-L. Barabási, “Error and attack tolerance of complex networks,” *Nature*, vol. 406, no. 6794, pp. 378–382, 2000.
- [24] L. La Cava and A. Tagarelli, “Information consumption and boundary spanning in decentralized online social networks: The case of mastodon users,” *arXiv:2203.15752 [physics]*, Mar. 29, 2022. arXiv: 2203.15752. [Online]. Available: <http://arxiv.org/abs/2203.15752> (visited on 04/01/2022).
- [25] P. E. Agre, “P2P and the promise of internet equality,” *Communications of the ACM*, vol. 46, no. 2, pp. 39–42, Feb. 2003, ISSN: 0001-0782, 1557-7317. DOI: 10.1145/606272.606298. [Online]. Available: <https://dl.acm.org/doi/10.1145/606272.606298> (visited on 01/21/2022).
- [26] A. X. Zhang, G. Hugh, and M. S. Bernstein, “PolicyKit: Building governance in online communities,” in *Proceedings of the 33rd Annual ACM Symposium on User Interface Software and Technology*, Virtual Event USA: ACM, Oct. 20, 2020, pp. 365–378, ISBN: 978-1-4503-7514-6. DOI: 10.1145/3379337.3415858. [Online]. Available: <https://dl.acm.org/doi/10.1145/3379337.3415858> (visited on 02/05/2022).
- [27] Z. Ao, G. Horvath, and L. Zhang, “Are decentralized finance really decentralized? a social network analysis of the aave protocol on the ethereum blockchain,” *arXiv preprint arXiv:2206.08401*, 2022.

- [28] M. Z. Trujillo, L. Hébert-Dufresne, and J. Bagrow, “The penumbra of open source: Projects outside of centralized platforms are longer maintained, more academic and more collaborative,” *EPJ Data Science*, vol. 11, no. 1, p. 31, 2022.
- [29] M. Z. Trujillo, M. Gruppi, C. Buntain, and B. D. Horne, “The mela bitchute dataset,” in *Proceedings of the International AAAI Conference on Web and Social Media*, vol. 16, 2022, pp. 1342–1351.
- [30] A. Mekacher and A. Papasavva, ““I can’t keep it up.” a dataset from the defunct voat. co news aggregator,” in *Proceedings of the International AAAI Conference on Web and Social Media*, vol. 16, 2022, pp. 1302–1311.
- [31] B. Taudul. “Archiwum polskiego usenetu.” (), [Online]. Available: <https://usenet.nereid.pl/> (visited on 05/27/2022).