Online Chess Social Networks

Eva Nolan

Hamilton College '22

Valentin Scognamillo

Hamilton College '22

November 2021

Contents

1 Introduction

The goal of this research is to analyze the structure of the network of chess players that play on Lichess.org. We aim to understand the way that Lichess randomizes player pairings and how closely related players are in order to better understand the relationship between ranking and pairing systems. We will also observe the behavior of players with respect to the game types that they play to see if this influences player groupings. Another hope of this project is to do an exploratory analysis of standout players and those that choose their opponents, rather than let the algorithm choose for them.

2 Background

There is little research that we could find on chess social networks online. Most existing research on chess involves networks of moves within games, analyzing the efficacy of different openings, closings, and other move strategies. While this research is valuable and can provide meaningful insight to how the game is played, there are far fewer studies on the structure of chess social networks, especially in an online environment. This environment is of particular interest because players from anywhere in the world have the potential to play together, creating the opportunity for a social network that is very connected. It can be assumed that low level chess players are generally not competing internationally, which would mean the in-person chess network could have very large distances between low level players. There is one existing paper available that compared the structure of online and in person chess social networks (2). Their research studied a longer time span than ours - we speculate they had greater computing power than we had the capacity for. They noted that high-ranked players maintained a more constant ranking overtime than low-ranked players did, and that the high-ranked players tended to be in more dense sections of the graph, more so in the online network than in the in person network. Our research touches on that behavior in online networks and hypothesises explanations for that behavior.

3 Data and Sources

Our data was gathered from the Lichess database (1), which contains the results of all games played online at Lichess.org. Our particular dataset contains all games played during the month of July 2014. This dataset was chosen due to it being the largest dataset that we could reasonably run computations on. The more recent networks contain more players and games, as online chess has grown in popularity. These newer datasets were too large for the some of the software we were using in analysis. Since 2014 there have been additional game types added, but the game types from 2014 remain the primary types. Thus the network we are working with is less complex than recent ones, but should still be a sufficient general representation of more recent networks, as the game of chess and the websites functionality haven't changed since 2014. The original data came in Portable Game Notation (extension .pgn). We reformatted this data using Python to convert the desired data into a CSV, the first few rows of which can be seen in Table 1. This data is organized as a list of edges, the source and target columns indicate the edge's nodes, and the other columns are edge attributes.

Table 1: Sample CSV data

Source	Target	WRating	BRating	Event	Opening	Termination	Result
kalitka_...	ZeruHmyz	2013	2133	Blitz	Frenc...	Time fo...	0-1
PLATIN...	sosisamm	2143	2042	Bullet	Frenc...	Normal	0-1
sosisamm	PLATIN...	2057	2129	Bullet	Semi-...	Normal	1-0
PLATIN...	sosisamm	2116	2070	Bullet	Frenc...	Normal	0-1
JasonVo...	amazingoid	2015	2184	Bullet	Sicilia...	Time fo...	0-1
⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮

- Source/Target: White player and Black player, respectively. Source and Target are used as headers as this is a directed network with edges pointing from the White player, who makes the first move, to the second player, Black.

- WRating/BRating: Player ratings for White and Black. In the original data, player ratings are an edge attribute, given by the players ranking at the start of the game. Rankings are adjusted after each game based on the result of the game. To generalize this information, we averaged each players ranking over the course of the month and assigned that average as a node attribute which was later loaded into analysis software.

- Event: There are several different game modes available on Lichess, which each denote the total amount of time allotted to each player when making moves. Bullet is the fastest event type, with no more than 3 minutes given to each player. The progression then goes to Blitz, which gives 5 to 10 minutes, and Classical gives 11 to 30. There are also tournaments of each game type that are present in the dataset.

- Opening: Refers to the initial moves of the game. It is important to note that this could be the result of moves played by either side.

- Termination: The manner in which the game ended. "Normal" indicates that the game resulted in a checkmate, and "Time forfeit" denotes that it ended as a result of one player running out of time.

- Result: The winner of the game. A White victory is indicated by 1-0, and a Black victory by 0-1.

We decided to only observe players that have a rating above 2000; we are interested in high-level play, and this is also a way to limit the size of the network to allow for faster computations.

4 The Network

The network's nodes represent players and edges represent games between players. As previously mentioned, edges are directed from White to Black. We viewed the network as a directed graph as part of our analysis of the websites randomization. In games played against opponents chosen at random, the players' color is also assigned. It is considered an advantage to play as white (make the first move). We found that in general, each node's in and out degree were fairly similar, though there seems to be a slight preference to play white, as can be seen in Figure 1.

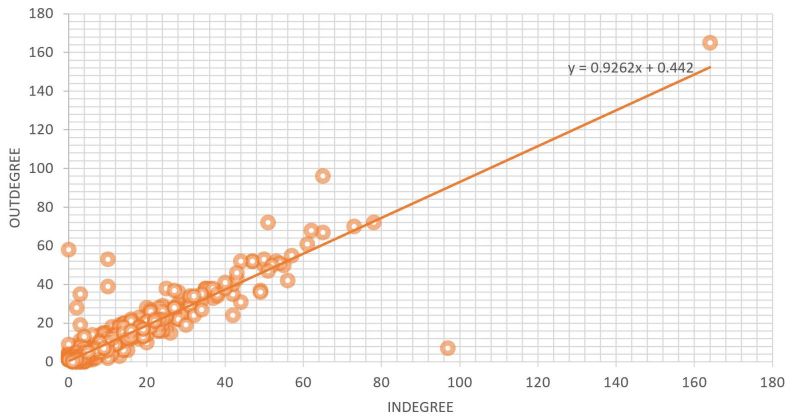

Figure 1: Relationship between in and outdegree of nodes.

The entire network contains 1379 nodes (players) and 21191 edges (games). In our analysis of the network, we do not allow multi edges, thus the weight of each edge indicates how many games were played between two players. In the digraph those games form 8256 edges. Thus a node's degree is the number of opponents, rather than the number of games played. The largest edge weight is 201, the average edge weight is 2.5667, and the median is 1. The few anomalous large edge weights can be attributed to players who repeatedly choose to play against each other instead of being assigned opponents by the site. Additionally, the network contains a giant weakly connected component which contains 97.53% of the network's nodes. We chose to observe weakly connected components because we only cared about people who played each other and did not feel that the player order mattered in components, since piece color is assigned arbitrarily by the site. Note that some players choose their opponents and thus are allowed to choose their piece color. This happens quite infrequently and doesn't significantly sway the data. See Section 5.1 for more on components.

Another trait of the network which was immediately visible in initial analysis was the formation of three distinctly dense areas within the giant component. When viewing

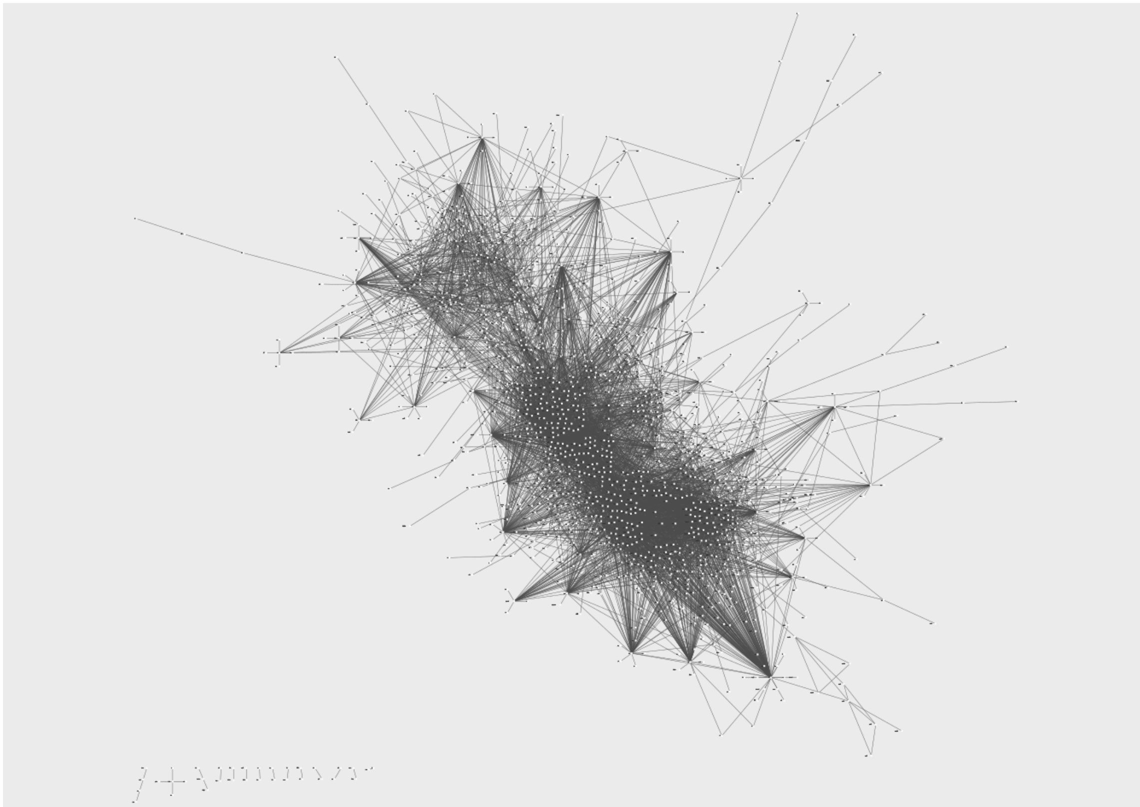

Figure 2: The entire Lichess network for the month of July 2014. Visualization Generated by Cytoscape (6)

the network with forced atlas visualization, it seemed the case that these dense areas were dominated each by one of the main game types offered by Lichess: bullet, blitz, and classical. The network also contains 3 main communities, identified using NetworkX (4). 86.86% of the nodes fall into one of these three communities, each consisting of 30.6%, 29.51%, and 26.76% of nodes respectively. The software identified a handful of very tiny communities, each containing less than 1% of the network's nodes. In our analysis we focus on the three primary communities, as the other small communities are dispersed somewhat randomly about the network and are contain nodes with very few edges that are not well connected to the rest of the network. We explore more formally these dense areas and their behavior in Section 5.2 on Communities.

Table 2: Summary statistics for entire network and game type sub-networks of Lichess games during July 2014. Tournaments of game mode variants are denoted with ($*$).

Network	n	m	$\langle k \rangle$	d	$l(g)$	C	ρ	H	Centralization
Total	1379	21191	8.522	13	3.771	0.142	0.006	1.651	0.133
Blitz	615	6041	7.823	7	3.193	0.143	0.013	1.645	0.157
Bullet	511	13115	8.761	7	2.968	0.23	0.018	1.593	0.33
Classical	413	1529	4.647	7	3.387	0.082	0.012	1.709	0.135
Blitz*	74	139	3.029	7	3.125	0.213	0.045	1.235	0.276
Bullet*	123	347	3.179	7	3.392	0.153	0.029	1.368	0.274
Classical*	18	20	2	7	3	0.091	0.143	0.856	0.495

4.1 Network Statistics

We analyzed the statistics of the entire network and on sub-networks created by only including edges and nodes of one game type, which can be seen in table 2. The network statistics were computed on undirected versions of their respective network as the interpretation of the values made more sense in the context of undirected edges. [1] We also analyzed the statistics for the sub-networks formed by each individual event type. We assigned event type as an edge attribute; for each type we removed all other edges and subsequently removed nodes of degree 0 – players that did not play that event.

We can interpret $\langle k \rangle$ as the number of opponents a player had during July 2014, and it can be observed that bullet players had the most opponents out of any event type; this would naturally arise as a bullet player can play more games per hour than a player of other event types. This also explains why such a large portion of edges fall under the Bullet category. The diameter d shows that the maximum degree of separation between players in the giant component of the total network is 13, while the sub-networks are more connected and all have a smaller diameter of 7 in their giant component. This is likely because paths need to traverse through multiple event type communities to connect nodes of different event types.

[1] n = nodes, m = edges, $\langle k \rangle = avg.degree$, d = diameter, $l(g)$ = characteristic path length, C = clustering coefficient, ρ = density, H = heterogeneity. For an explanation of the statistics and their formulas, please see the appendix.

The characteristic path length $l(g)$ tells us the average shortest distance between any two pairs of nodes. This exemplifies the small-world effect, as even though we have a graph consisting of 1379 players and a diameter of 13, there are on average only 3.77 degrees of separation between any two players.

There are far fewer edges in the classical graph, as games last considerably longer and thus fewer games are possible over the same time span as other events. As a result, classical players are less closely clustered as there are fewer edges and players have fewer opponents, reducing the probability of clusters arising. Low C, ρ, and centralization values indicate that all sub-networks are not very well connected or dense networks. This is similar to the structure of other large social networks.

5 Results

5.1 Components

The network contains one giant component and 13 small components. We will first focus on the nature of the giant component, and we will explore small components later in the section.

With 97.93% of the network being in the giant component, it seems that Lichess is effective in connecting players across a singular network. One outcome of this giant network is that player ranking is based on a large number of statistics, as nearly every node can reach another by some path, and rankings comparisons can be in part determined by these paths. However, there were a few branches of the network that stood out in comparison to the shape of the majority of the network. There were a small handful of branches in which players are connected along a path with no alternative paths back to the more connected part of the component. Investigating further statistics on the players in such branches, it appears that they are often involved in manipulating elo. More on this behavior is in section 5.4.

Another structural element of interest in the network was small components. There were a total of 13 small components, each consisting of between 2 and 4 nodes and an average weighted node degree of 1.84. The Lichess website saves all user game history in a public format, which allowed us to look into players involved in small components. While player behavior varied a bit, it was typically the case that the players in small components were often only active on Lichess for a brief period of time and played a very small number of games. Many pairs of nodes from the same small component joined Lichess during July 2014. While we cannot confirm reasoning, it is a reasonable assumption to believe that players within one small component, if they joined at the same time, know each other in real life and joined the site to play each other. Thus there is far less randomization of player pairings in the small components. There were a few other cases in which players are still active, but joined the website late in July 2014, and were not active enough yet to connect to the community. Given the similarity between these players and one time players in the network, a possible explanation is that they joined the website with the intention of playing a specific person, but did not fully develop interest in chess until later on.

5.2 Communities

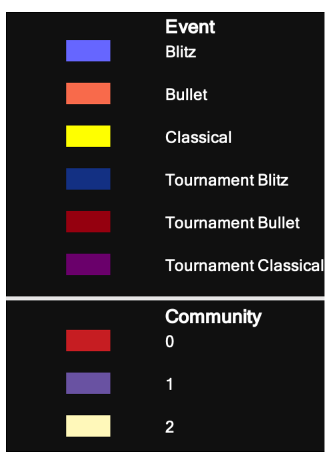

Figure 3: Legend to be used when coloring networks.

In initial network analysis, three very dense sections suggested that there might be com-

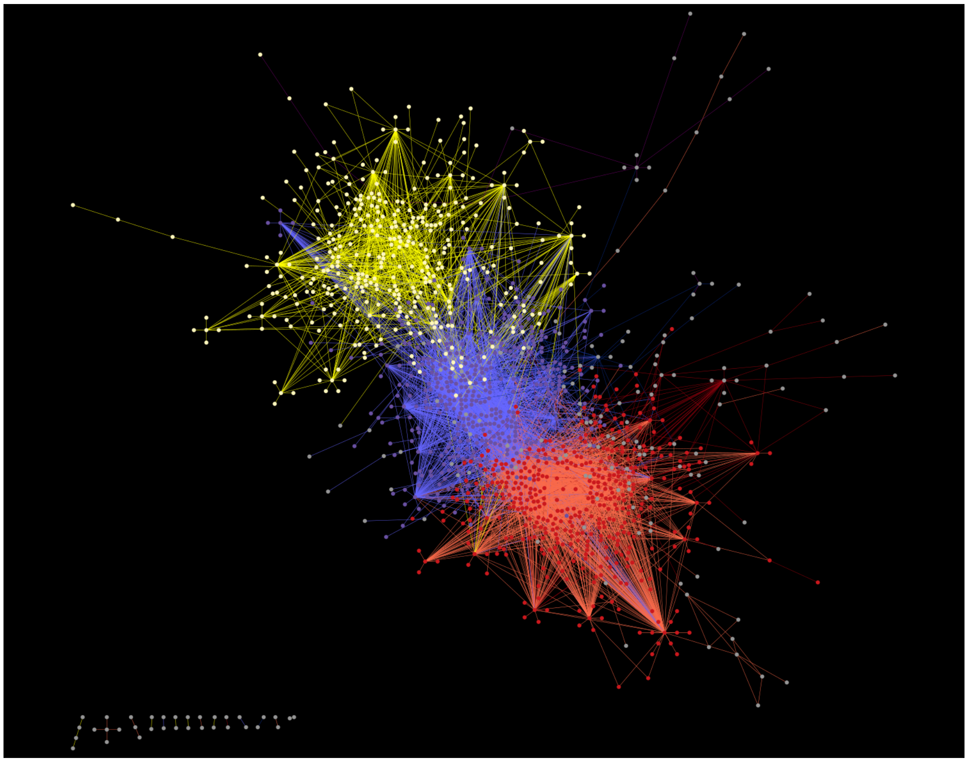

Figure 4: The entire Lichess network for the Month of July 2014.

munities at play in the network. Communities are essentially clusterings of nodes that are more connected to each other than they are to the rest of the network. There are multiple different algorithms that can be used to identify communities, and this is still a developing area of network analysis. That being said, we chose the greedy modularity algorithm, which initially considers each node to be its own community. It then iterates through each community, merging it with the community that would most increase its modularity. The algorithm essentially maximizes the modularity of each community, so that communities are individual densely connected subnetworks. Using the greedy_ modularity_communities function in NetworkX (4), we identified a total of 31 communities, only three of which were significant in size. As noted in Section 4, the three large communities contain 30.6%, 29.51%, and 26.76% of nodes respectively. When comparing these communities with node and edge traits, our hypothesis of its relation to game type was confirmed. Coloring edges by game type and nodes by community, it is visually clear that players who primarily play Bullet

Table 3: Event–Community Relationship

Event	nodes in Community 0	nodes in community 1	nodes in community 2
Bullet	75.75%	11.91%	0.2%
Blitz	17.61%	63.97%	11.31%
Classical	1.21%	12.11%	81.11%

events tend to be in the largest community, those who primarily play Blitz fall into the second largest community, and those who play classic fall into the third largest community, as can be seen in Figure 4. Among the smaller communities, players' preferred events are more unpredictable. It was also clear that while there is a fair amount of overlap between the three events, Classical and Bullet have a lower tendency to overlap than other combinations of game type. Given the duration of each game type, this network behavior makes sense. It suggests that players are more willing to play games closer in pace to their preferred event.

When considering that the network is comprised of high-ranked players, it suggests that experienced chess players have a distinct preferred game, which is what creates these communities. For example, if players who prefer Bullet typically elect to play their event of choice, they are likely paired with another player who also tends towards bullet. Thus there is much less interaction between players of different event types. We created sub-graphs that focus individually on the three primary event types, which we will later compare to Configuration Model random graphs to explore the structure of each event's network (See Figures 5, 6, and 7). Note that in these sub-graphs, nodes are still colored according to communities of the original network, but since all edges are of one game type, we did not preserve the edge coloring from the original network. Images of these sub-graphs can be seen in Figures 5, 6, and 7. The breakdown of each sub-graph's community makeup is in Table 3. These statistics confirm that each community is generally composed of players who prefer a particular event type. Further, it emphasises the earlier point that players of Bullet and Classical have much less interaction than observed for other event types.

(a) Bullet Real Network

(b) Random

Figure 5: Real and random graphs for the bullet event network, with nodes colored by community

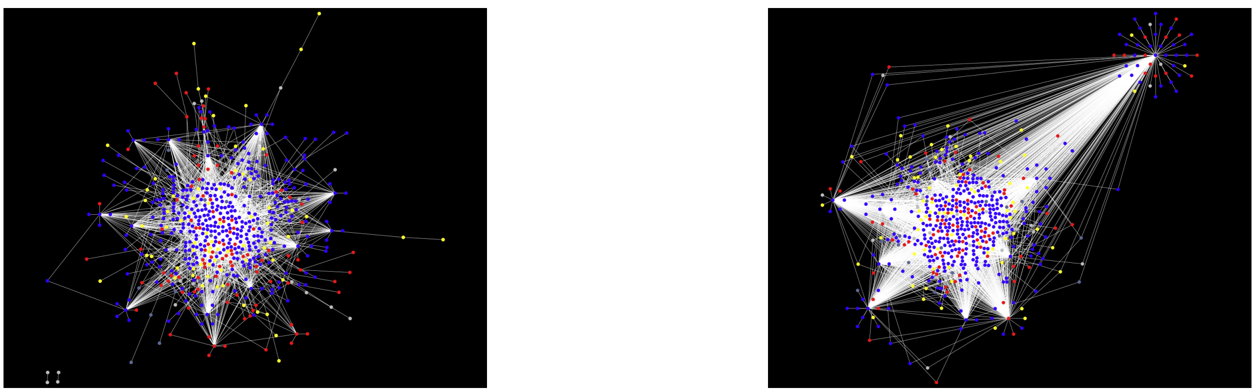

(a) Blitz Real Network

(b) Random

Figure 6: Real and random graphs for the blitz event network, with nodes colored by community

(a) Classical Real Network

(b) Random

Figure 7: Real and random graphs for the classical event network, with nodes colored by community

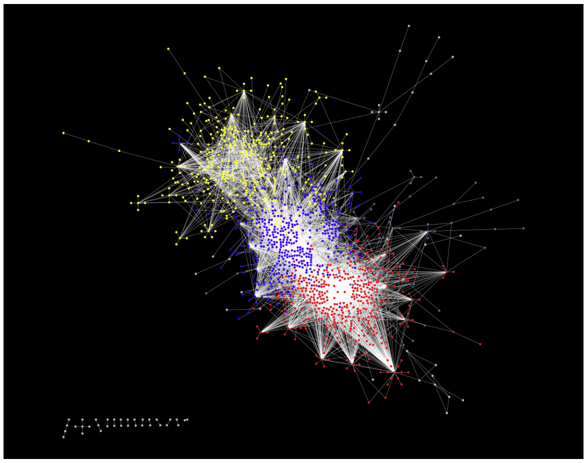

(a) Real

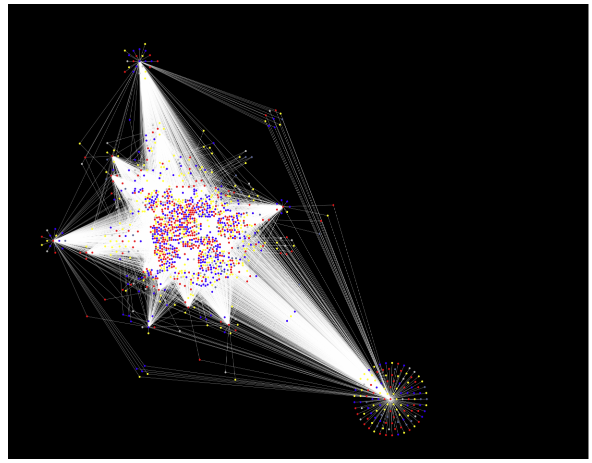

(b) Random

Figure 8: Real and random graphs for the whole network, with nodes colored by community

5.3 Comparisons to Random Graphs

We generated a Configuration Model random graph using the degree sequence of the original network. In Figure 8 it can be observed that the communities formed in the real network become completely meaningless in the random graph. We went on to compare each event type with a configuration model mimicking the events' degree distributions. Visualizations of each event sub-graph with its associated configuration model can be seen in Figures 5, 6, and 7. Note that player Jlindner has more than double the degree of any other player in the network, and plays both Bullet and Blitz games. In the original network he has high-weight edges to multiple other players of high degree, but in the random graphs, he is more connected to very low degree players. The cluster of nodes that juts out from the random graphs in figures 5(b) and 6(b) are centered around this player. We do not observe this behavior in Figure 7(b) because there are no classical players with anomalously high degree. The structure of the Classical network is not remarkably different from that of its related Configuration model, which indicates that the Lichess site effectively randomizes opponents for this event. While the Bullet and Blitz configuration models are visually very different from their corresponding non-random networks, the only significant difference is the behavior of the graphs around Jlindner's node. Since Jlindner is a play that chooses

Table 4: Real and random graph summary statistics.

	n	m	$\langle k \rangle$	d	$l(g)$	C	ρ	H	Centralization
Total	1379	21191	25.653	13	3.771	0.142	0.006	1.651	0.133
Random Total	1379	21191	25.653	4	2.255	0.602	0.019	2.461	0.734
Blitz	615	6041	17.649	7	3.193	0.143	0.013	1.645	0.157
Random Blitz	615	6041	17.649	4	2.248	0.595	0.029	2.106	0.706
Bullet	511	13115	30.2	7	2.968	0.23	0.018	1.593	0.33
Random Bullet	511	13115	30.2	4	2.068	0.699	0.059	1.679	0.797
Classical	413	1529	7.4	7	3.387	.082	.012	1.709	0.135
Random Classical	415	1529	7.4	6	2.772	.297	.018	2.071	.267

opponents, we can ignore this node when considering the efficacy of the sites' random pairing. The structure of the networks excluding this node is in fact fairly similar to the random ones, indicating that on the whole, the Lichess site effectively randomizes opponents.

5.4 Player Ratings and Cheating

Player ratings are calculated on Lichess using the Glicko-2 rating system (3). In this system, each player is assigned a rating of 1500 upon creating an account. The rating algorithm calculates a new score for a player after each game by taking into consideration the difference in original rankings of the players, as well as a measure of confidence of each player's rank – the more games a player has played, the more confident the algorithm is of a rating. Therefore, if someone with a newly created account were to beat a highly ranked player, their player rank would rise dramatically.

In our exploratory analysis we discovered a some long branches out from the giant component. We picked one of the longer branches to investigate. This branch's root (where it connects to the large component) is player jiktak, then players Blitz177, GNBGHBGH, HJKJKHJK, KSKSKSKSXFNDMVCNMXV, and KIKIKIKIKCVKC lay along the branch and Blitz1777 is the branch's leaf, as can be seen in Figure 9. For brevity, refer to players along the branch as players A, B, C, and D. We looked further into these players' histories on the Lichess site. All 4 players along the branch have usernames that appear to be a random combination of letters. These 4 accounts were all created on July 6, 2014, and all of their games also occurred on that day. In all of these games, each player made only 1 or 2 moves, and the game ended due to running out of time. Note that upon account creation, each of these players were given the starting elo of 1500. Prior to July, Blitz177 had played seven games against Blitz8, winning all of them and achieving an elo of 2255. Blitz177 then played 5 games against A, losing every time, raising A's elo to 2283. Then A played seven games against B, lost every game, raising B's elo to 2442. Then B played 6 losing games against C, raining C's elo to 2561. C played D, raising D's elo to 2622. Finally D played 5 games with Blitz1777, raising Blitz1777's elo to 2596. It seems extremely likely that these accounts were created for the sole purpose of passing elo up to player Blitz1777. We cannot see what happened after these games as both Blitz177 and Blitz1777 have had their accounts deactivated. This branch of the network shows that it in possible for players to manipulate the ranking system. However, there are very few branches like this in the network. From visual inspection of the network, it seems like there are under 10 branches that follow this structure, so we can conclude that elo manipulation of this nature is fairly infrequent in the context of the entire network.

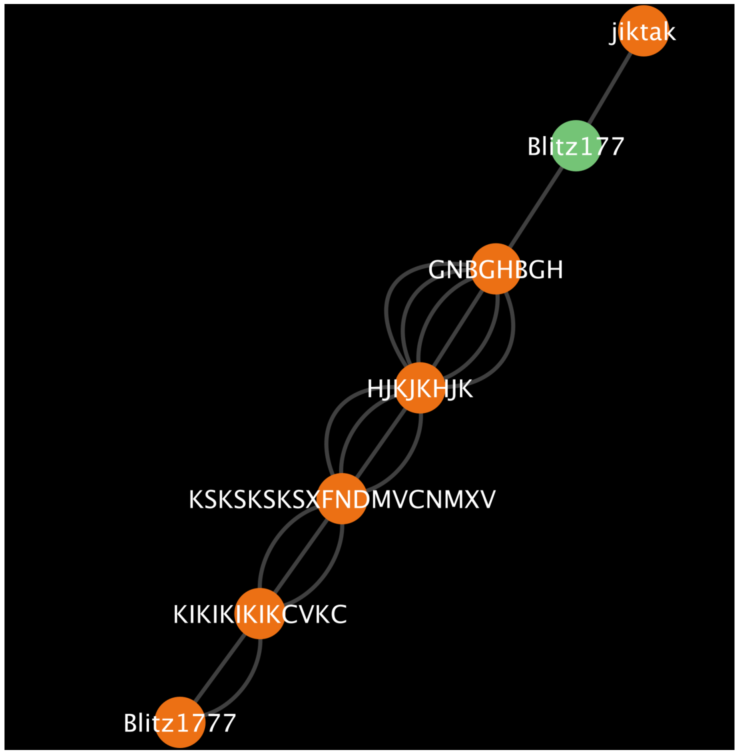

Figure 9: Network of suspected cheaters.

Blitz177 played a total of 74 games before getting banned, 2 of which show up as having an anonymous opponent. Blitz177 had 37 wins, 36 losses, and 1 draw – this is unlikely a bot, or at most it is a user that sometimes switches to a bot. Another hypothesis could be that Blitz177 uses bots on the other accounts in the sub-graph in order to inflate their rating, and then beats them in order to improve his own.

6 Conclusion

Our analysis establishes that, generally speaking, Lichess is extremely effective at randomizing opponent pairings and maintaining a well connected network. Since the average path length and diameter are so low in each event's subnetwork, this indicates that the network effectively bridges geographical barriers and serves as a playing platform who's rating system spans a range of demographics. Thus its rating system is able to compare players across the world very effectively, meaning that the sites rating system is likely more thorough than that of in-person chess. Figure 10 shows that there is a normal distribution for games played vs. win rate for each player – this is further proof of the efficacy of the matching system, as a player over time plays against opponents of similar skill levels.

While online chess introduces different means of cheating than in-person chess might, these problems seem to make up a very small portion of the network, and it can still be concluded that the Lichess system is able to fairly compare players on a scale that is more comprehensive and far-reaching. The lack of distinct communities within each game type's network indicates that geographical/time zone factors have little to no sway on opponent matchings.

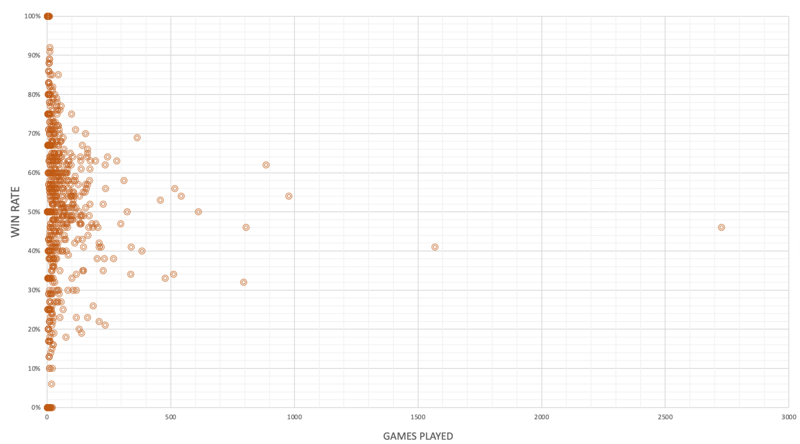

Figure 10: Games played vs. win rate for players of the Lichess network.

7 Appendix

7.1 Definitions of statistics

Here we detail the calculations used by the Cytoscape (6) application for measures that differ from those in Mark Newman's Networks (5)

- Characteristic path length: The characteristic path length $l(G)$ of a graph $G = (V, E)$ is defined as the average number of edges in the shortest paths between all vertex pairs given by

$$l(G) = \frac{1}{|V| \cdot (|V| - 1)} \sum_{v \in V} \sum_{v' \in V \setminus \{v\}} spl(v, v'),$$

where $spl(v, v')$ gives the number of edges in a shortest path between vertices v and v'.

- Centralization: The normalized connectivity centralization is given by

  $$Centralization = \frac{n}{n-2}(\frac{max(\mathbf{k})}{n-1} - Density) = \frac{max(\mathbf{k})}{n} - Density,$$

  where $max(\mathbf{k})$ is the max degree.

- Heterogeneity: The coefficient of variation of the connectivity distribution i.e.

  $$Heterogeneity = \frac{\sqrt{variance(\mathbf{k})}}{mean(\mathbf{k})}.$$

References

[1] *Standard Chess Data*, Lichess.org public database, July 2014. Accessed 2021-10.

[2] N. Almeira, A. Schaigorodsky, and J. Perotti et al, *Structure constrained by metadata in networks of chess players*, Scientific Reports, 7 (2017).

[3] M. E. Glickman, *Example of the Glicko-2 system*, (2013).

[4] A. A. Hagberg, D. A. Schult, and P. J. Swart, *Exploring network structure, dynamics, and function using NetworkX*, in Proceedings of the 7th Python in Science Conference (SciPy2008), (2008), pp. 11–15.

[5] M. Newman, *Networks*, Oxford University Press, 2 ed., 2018.

[6] P. Shannon, A. Markiel, O. Ozier, N. S. Baliga, J. T. Wang, D. Ramage, N. Amin, B. Schwikowski, and T. Ideker, *Cytoscape: a software environment for integrated models of biomolecular interaction networks*, Genome research, 13 (2003), pp. 2498–2504.