

Analyzing Implicit Social Networks in Multiplayer Online Games

Iosup, A; van de Bovenkamp, R; Shen, S; Jia, AL; Kuipers, FA

DOI

[10.1109/MIC.2014.19](https://doi.org/10.1109/MIC.2014.19)

Publication date

2014

Document Version

Accepted author manuscript

Published in

IEEE Internet Computing

Citation (APA)

Iosup, A., van de Bovenkamp, R., Shen, S., Jia, AL., & Kuipers, FA. (2014). Analyzing Implicit Social Networks in Multiplayer Online Games. *IEEE Internet Computing*, (special issue on Networked Games). <https://doi.org/10.1109/MIC.2014.19>

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.

Analyzing Implicit Social Networks in Multiplayer Online Games

Alexandru Iosup, Ruud van de Bovenkamp, Siqi Shen,
Adele Lu Jia, Fernando Kuipers

Delft University of Technology, the Netherlands.

Contact: {A.Iosup,R.vandeBovenkamp,S.Shen,L.Jia,F.A.Kuipers}@tudelft.nl

Abstract—For many networked games, such as the Defense of the Ancients and StarCraft series, the unofficial leagues created by players themselves greatly enhance user-experience, and extend the success of each game. Understanding the social structure that players of these games implicitly form helps to create innovative gaming services to the benefit of both players and game operators. But how to extract and analyse the implicit social structure? We address this question by first proposing a formalism consisting of various ways to map interaction to social structure, and apply this to real-world data collected from three different game genres. We analyse the implications of these mappings for in-game and gaming-related services, ranging from network and socially-aware matchmaking of players, to an investigation of social network robustness against player departure.

I. INTRODUCTION

Networked games are games that use advances in networking and a variety of socio-technical elements to entertain hundreds of millions of people world-wide. Unsurprisingly, such games naturally evolve into Social Networked Games (SNGs): the many people involved organize, often spontaneously and without the help of in-game services, into gaming *communities*. While typical online social networks revolve around friendship relations, new classes of prosocial emotions appear in SNGs. For instance, adversaries motivate each other and together may remain long-term customers in an SNG. Adversarial relationships are one of the implicitly formed in-game relationships we study. Understanding in-game communities and social relationships could help improve existing gaming services such as team formation, planning and scheduling of networking resources, and even retaining the game population.

Few games exhibit a greater need for socially-aware services than the relatively new genre of multiplayer online battle arenas (MOBAs) considered in Section II. Derived from Real-Time Strategy (RTS) games, MOBAs are a class of advanced networked games in which equally-sized teams confront each other on a map. In-game *team-play*, rather than individual heroics, is required from any but the most amateur players. Outside the game, *social relationships* and etiquette are required to be part of the successful clans (self-organized groups of players). Players can find partners for a game instance through the use of community websites, which may include services that matchmake players to a game instance, yet are not affiliated with the game developer.

In the absence of explicitly expressed relationships, understanding the social networks of current SNGs must rely on extracting the *implicit* social structure indicated by regular player activity. However, in contrast to general social networks, a set of meaningful interactions has not yet been defined for SNGs. Moreover, in MOBAs, activities are match- and team-oriented, rather than individual. We address these challenges, in Section III, through a formalism for extracting implicit social structure from a set of SNG-related, meaningful interactions. We extend our previous work [4] by showing that the implicit social structure of SNGs is strong, rather than the result of chance encounters, and that, for MOBAs, the core of the network (the high-degree nodes) is robust over time.

In addition, we apply our formalism to RTS and Massively Multiplayer Online First-Person Shooter (MMOFPS) games, and, in Section IV, show evidence that RTS games exhibit even stronger team structure than MOBAs and indicate that modern

MMOFPSs may require operator-side mechanisms to spurn the formation of meaningful social structure.

Connecting theory to practice, we also show how the extracted implicit social graphs can be useful for improving gameplay experience, and for player and group retention (Section V), for tuning the technological platform on which the games operate, etc.

Last, we identify several challenges and future opportunities for SNGs, in Section VI.

II. SNGS WITHOUT AN EXPLICIT S

Defense of the Ancients (DotA) is an archetypal MOBA game. For DotA, social relationships, such as same-clan membership and friendship, can improve the gameplay experience [1]. DotA is a 5v5-player game. Each player controls an in-game avatar, and teams try to conquer the opposite side's main building. Each game lasts about 40 minutes and includes many strategic elements, ranging from team operation to micro-management of resources.

To examine implicit relationships in DotA, we have collected data for the DotA communities Dota-League and DotAlicious. Both communities, independently from the game developer, run their own game servers, maintain lists of tournaments and results, and publish information such as player rankings. We have obtained from these communities, via their websites, all the unique matches, and for each match the start time, the duration, and the community identifiers of the participating players. After sanitizing the data, we have obtained for Dota-League (DotAlicious) a dataset containing 1,470,786 (617,069) matches that took place between Nov. 2008 and Jul. 2011 (Apr. 2010 and Feb. 2012).

III. A FORMALISM FOR IDENTIFYING IMPLICIT SOCIAL RELATIONSHIPS

A. Social Relationships in SNGs

A *mapping* is a set of rules that define the nodes and links in a graph. Formally, a dataset D is mapped onto a graph G via a mapping function $M(D)$, which maps individual players to nodes (graph vertices) and relationships between players to links (graph edges).

Instead of proposing a graph model, we focus on formalizing mappings that extract graphs from

real data. Because many metrics of social networks only apply to unweighted graphs, relations are often considered as links only if their weight exceeds a threshold. *Thresholding*, therefore, has an important impact on the resulting graph.

Related to our work¹, interaction graphs [5] map users of social applications to nodes, and events involving pairs of users to links via a threshold-based rule.

B. Interaction Graphs in MOBAs

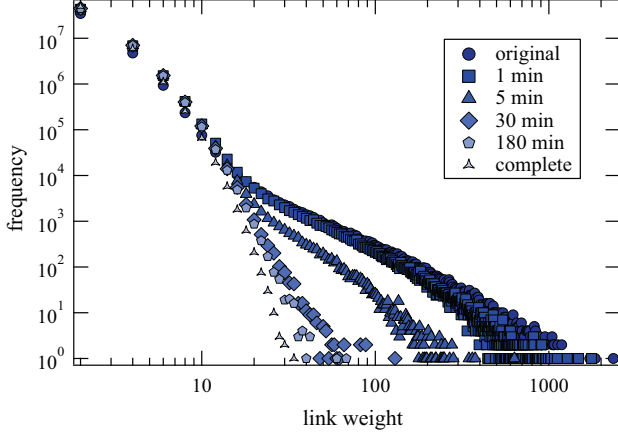
A mapping is *meaningful* if it leads to distinct yet reasonable views of implicit social networks appearing in networked games. We identify six types of player-to-player interactions:

- SM: two players present in the *Same Match*.
- SS: two players present on the *Same Side* of a match.
- OS: two players present on *Opposing Sides* of a match.
- MW: two players who *Won* a match together.
- ML: two players who *Lost* a match together.
- PP: (*directed*) for a player, when present in at least $x\%$ of another player's matches ($x = 10\%$ in this article). This interaction is effectively PP(SM). Similarly, we can define PP(SS), etc.

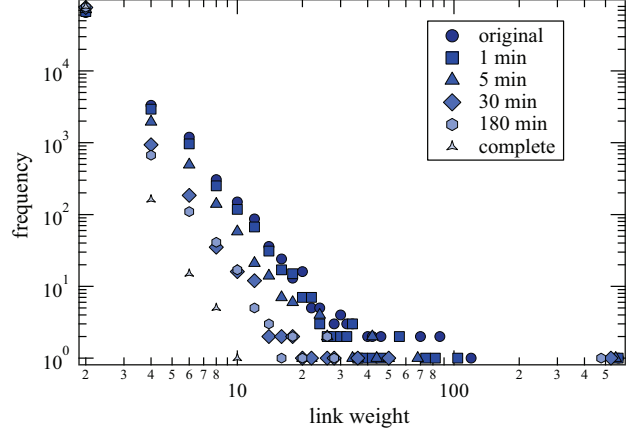
To extract the social networks corresponding to various types of relationships, we extract for each mapping a graph by using a threshold n , which reflects the minimum number of events that need to have occurred between two users for a relationship to exist; e.g., for $SM(n = 2)$, a link exists between a pair of players iff they were both present in at least two matches in the input dataset. A second threshold, t , limiting the duration-of-effect for any interaction, is less relevant, as explained in Section III-C.

The set of mappings proposed here is not exhaustive. For example, this formalism can support more complex mappings, such as “played against each other at least 10 times, connected through ADSL2, while located in the same country”. The interactions in the set are also not independent. For example, the SS mapping can be seen as a specialization of the SM mapping.

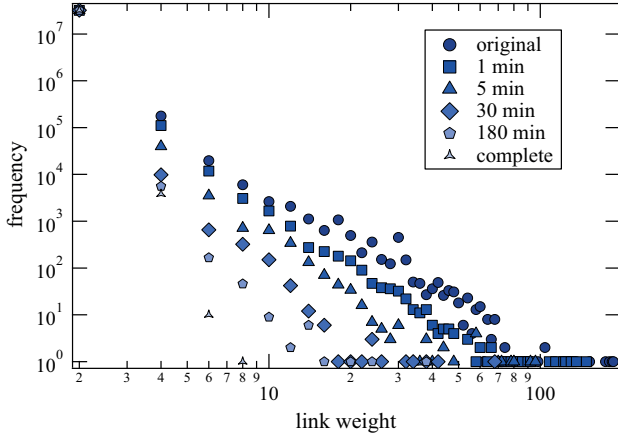
¹Related work is discussed throughout this article and in [4].



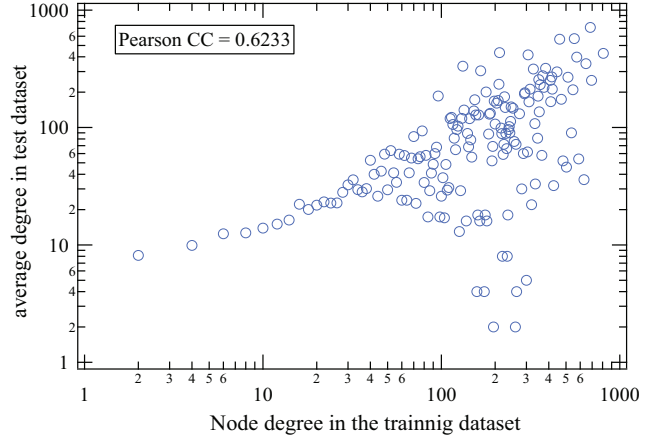
(a) Frequency of occurrence of link weights in the reference model for five different intervals and the original data for DotA. The mapping is SM.



(b) Frequency of occurrence of link weights in the reference model for five different intervals and the original data for SC2. The mapping is SM.



(c) Frequency of occurrence of link weights in the reference model for five different intervals and the original data for League (SM with $n = 10$). WoT. The mapping is SM.



(d) Continued activity for gaming relationships in Dota-model for five different intervals and the original data for League (SM with $n = 10$). WoT. The mapping is SM.

Fig. 1: Results for methodological questions Q1 (a-c) and Q2 (d).

C. Application to the Examples

We focus in this article on three methodological questions:

Q1. Are the relationships we identify the result of players being simultaneously online by chance? To answer this question, we first create a reference model by randomizing, for any window of length w minutes, the interactions observed in the MOBA datasets. The randomization of, for example, the SM mapping is done by taking the players from all matches that started within the current time window and randomly assigning them to matches. Since the SM mapping does not take team information into account, the match assignment comes down

to forming random groups of 10 players from the entire list who were active in the time window. A single player can be in the list multiple times and the random groups have to consist of 10 different players.

We run the parameter w from 1 to ∞ and depict the results, together with the original data, in Figure 1a. Whereas the results for $w = 1$ leave little room for randomization, the results for $w = \infty$ randomize the entire dataset. In Figure 1a, the curves for $w = 1$ and the original data have a powerlaw-like shape. The curves for various values of w follow the $w = 1$ (original) curve for link weights of up to about 15 matches played together, but take afterwards an exponential-like shape, which indicates they are

more likely to be the result of chance than of intended user behavior. The fact that curves are markedly different for small time windows shows that it is very unlikely that players play together often simply because they happen to be online at the same time.

The results for the other game genres (genres introduced in Section IV, results depicted in Figures 1b and 1c) show similar, yet not so pronounced behavior. Although players do not play nearly as often together in other genres' datasets as in the MOBA datasets, randomization within only small time windows lowers the link weights. We conclude that it is unlikely that the relationships we identify are the result of chance encounters between players and, instead, indicate conscious, possibly out-of-game agreements between players.

Q2. Are players (nodes) preserving their high-degree property over time? If so, then the networks these players form may be robust against natural degradation, with implications for the long-term retention of the most active players. For each MOBA-community, we first divide its last-year's gaming relationships into two parts: the first half year as training data and the second half year as testing data. We only use players who appear in both datasets—about 60% of the training-data players. Then, we plot in Figure 1d, for different degrees of players in the training dataset, the average number of links formed by these players in the testing dataset. From the high-value and positive correlation-coefficient (0.6233 for Dota-League), we derive that players with higher degrees in the training dataset robustly establish more new links in the testing dataset than the other players.

Q3. Are the mappings we propose meaningful for MOBAs? To answer this question, we first extract the interaction graphs for each of our mappings, compute for each a variety of graph metrics, and summarize the results in Table I. We find that:

- Side-specific interactions (SS and OS) are meaningful. For example, playing on the opposing side (OS) is more likely than playing on the same side (SS), in Dota-League (for example, higher N and L in Table I); for DotAlicious, the reverse is true. Game designers could enable OS links by allowing players to explicitly identify their *foes*.
- Outcome-specific interactions (MW and ML) are meaningful. For example, only for

DotAlicious, MW leads to more relationships being formed. Game operators could exploit this in matchmaking services.

- Relative joint participation (PP) is meaningful. For example, for PP(SM), the number of nodes in the graph decreases quickly with the n threshold. Identifying the players who play almost exclusively together can be key to player retention.

IV. APPLICATION TO OTHER GAME GENRES

Among the most popular genres today, RTS games ask players to balance strategic and tactical decisions, often every second, while competing for resources with other players. Although faster-paced, MMOFPS games test the tactical team-work of players disputing a territory. We could expect RTS and MMOFPS games to lead to similar interaction graphs as MOBAs: naturally emerging social structures centered around highly active players. However, these game genres also have different match-scales and team-vs-team balance than MOBAs. Moreover, RTS games can stimulate individualistic gameplay, while MMOFPS games may have teams that are too large to be robust.

We collect then analyze two additional datasets: for the RTS game StarCraft II (SC2) from Mar. 2012 to Aug. 2013, and for the MMOFPS game World of Tanks (WoT) from Aug. 2010 to Jul. 2013. For each of these popular games, we have collected over 75,000 matches, played by over 80,000 SC2 and over 900,000 WoT players. SC2 matches are not generally played in equally-sized teams, and 92% of our dataset's matches are 1v1-player. In contrast, 98% of WoT matches are 15v15-player, but such large teams can be much harder to maintain over time than the teams found in typical MOBAs, due to inevitable player-churn.

Alone or together? For SC2, the mappings lead to small graphs, with many small connected components. The majority of players participate in 1v1-player matches, but the 8% of players who do play in larger groups tend to play against each other more than together ($N = 611$ for the OS mapping, versus 314 for SS). When players do play on the same side, winning tends to strengthen the teams ($N = 212$ for the MW mapping, versus 95 for ML), just as we saw for the DotAlicious dataset. The connected components are strongly connected, yet

	DotA-League						DotAlicious					
	SM	OS	SS	ML	MW	PP	SM	OS	SS	ML	MW	PP
N	31,834	26,373	24,119	18,047	18,301	29,500	31,702	11,198	29,377	22,813	21,783	34,523
N_{lc}	27,720	19,814	16,256	6,976	8,078	33	26,810	10,262	20,971	10,795	13,382	3,239
L	202,576	85,581	62,292	30,680	33,289	53,514	327,464	92,010	108,176	43,240	54,009	125,340
L_{lc}	199,316	79,523	54,186	17,686	21,569	120	323,064	91,354	99,063	29,072	44,129	17,213
$d (\times 10^{-4})$	4.00	2.46	2.14	1.88	1.99	0.62	6.52	14.7	0.49	1.66	2.28	1.05
$d_{lc} (\times 10^{-4})$	5.19	4.05	4.10	7.27	6.61	1,100	8.99	17.4	2.51	4.99	4.93	16.4
μ	0.0301	0.0114	0.0060	0.0040	0.0032	-	0.0385	0.0403	0.0120	0.0095	0.0194	-
\bar{h}	4.42	5.40	6.30	8.09	7.67	3.70	4.24	3.97	5.3	6.80	5.95	18.45
D	14	21	24	28	26	9	17	12	19	20	22	74
\bar{C}	0.37	0.40	0.41	0.41	0.41	-	0.43	0.27	0.47	0.47	0.49	-
ρ	0.13	0.26	0.25	0.27	0.28	-0.10	0.08	0.01	0.25	0.27	0.29	0.20
B_m	0.04	0.09	0.09	0.17	0.12	1.21	0.03	0.05	0.04	0.06	0.06	0.37
c_m	85	55	41	19	22	3	131	68	48	16	20	7

	StarCraft II					World of Tanks				
	SM	OS	SS	ML	MW	SM	OS	SS	ML	MW
N	907	611	314	95	212	4,340	477	4,251	561	1,824
N_{lc}	31	22	24	9	14	129	118	122	66	57
L	748	404	327	85	200	9,895	3,253	6,543	1,564	2,923
L_{lc}	58	21	44	13	24	2,329	1,243	1,160	519	473
$d (\times 10^{-4})$	18	22	67	190	89	10.51	286.54	7.24	99.57	17.58
$d_{lc} (\times 10^{-4})$	1,247	909.10	1,594	3,611	2,637	2,821	1,801	1,572	2,420	2,964
D	2	8	3	2	2	6	3	5	4	3
\bar{C}	0.58	0	0.70	0.65	0.65	0.79	0.10	0.78	0.88	0.87
ρ	-0.46	-0.45	-0.42	-0.58	-0.53	-0.10	-0.12	-0.06	-0.03	-0.10
B_m	0.91	0.74	0.84	0.80	0.79	0.09	0.08	0.11	0.20	0.20
c_m	53	11	32	32	32	29	15	14	14	14

TABLE I: Results for methodological question Q3. Metrics [4] for $n = 10$: (top) Data for MOBA games: the Dota-League and DotAlicious datasets [4]. (bottom) Data for other game genres: StarCraft II (RTS) and World of Tanks (MMO and FPS). The metrics we present: number of nodes N , number of nodes in largest connected component N_{lc} , number of links L , number of links in largest connected component L_{lc} , link density d , link density of largest connected component d_{lc} , algebraic connectivity μ , average hop count \bar{h} , diameter D , average clustering coefficient \bar{C} , assortativity ρ , maximum betweenness B_m , and maximum coreness c_m .

small. The connected components of the mappings extracting same-team graphs are highly clustered, whereas the largest component for the OS mapping is even a tree. The clustering coefficients observed in the various RTS networks indicate much stronger team relationships in RTS games than in MOBAs. Because RTS games have not shown a trend of greatly increasing the number of players in the same instance, over the last decade, we hypothesize that RTS games will continue to spawn tightly-coupled teams that always play together; such teams are naturally vulnerable to player departures.

For WoT, the large team-size makes it difficult to organize teams well: the largest connected components for all mappings are not very large. Similarly to SC2 and DotaLicious, in WoT the players who do play often together do so on the same team and, again, players who play together are more likely to win rather than lose together. As modern FPS games tend to be played in increasingly

larger teams, with 32v32-player games now not uncommon, we conclude that MMOFPS games will require additional mechanisms if they are to develop any form of robust social structure. Moreover, even more so than in the SC2 datasets, many players play only one or a few games: 69% of the more than 900,000 players played only once or twice. This is another area where developers could use the emerging social structures among their players to increase the number of players who keep on playing the game.

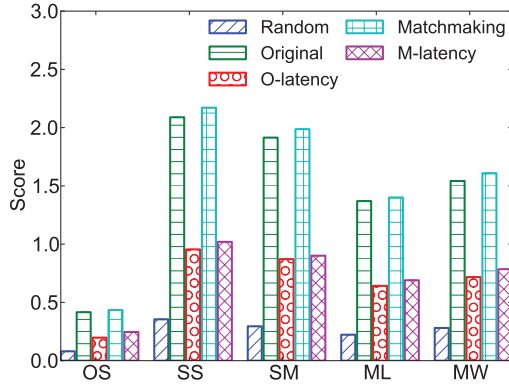
We conclude that our formalism can be applied to other game-genres, for which it leads to *new findings* vs MOBAs, and suggest that even communities of popular networked games could benefit from new mechanisms that foster denser interaction graphs.

V. APPLICATION TO SNG SERVICES

“How can social-networking elements be leveraged to improve gaming services?” We present in this section two exemplary answers.

Team 1		Team 2	
Player	Cluster	Player	Cluster
a	1	f	2
b	2	g	5
c	1	h	3
d	3	i	6
e	4	j	3

(a) Example of scoring for a match. Team 1 consists of players ‘a’ to ‘e’, as can be seen in the column labeled ‘Player’; team 2 consists of players ‘f’ to ‘j’. The column labeled ‘Cluster’ records the cluster identifier for each player. A match receives one point for every same-cluster player present in the match, when at least 2 same-cluster players are present. In this example, 2 points are given for player ‘a’ and ‘c’ (cluster 1), and for players ‘b’ and ‘f’ (cluster 2); 3 points are given for players ‘d’, ‘h’, and ‘j’ (cluster 3). Players ‘e’, ‘g’, and ‘i’ have no fellow cluster-members in the match and will be assigned 0 points. In total, this match is assigned 7 points.



(b) Average match scores for various matchmaking approaches. When considering network latency, matches with players on several continents score 0 points; the others use the scoring exemplified in Figure 2a. “Random” denotes matches obtained via randomly matching players who are online during each time interval. “Original”/“O-latency” denote matches observed in the real (raw) datasets without/with network latency considerations. “Matchmaking”/“M-latency” denote matches obtained with our proposed matchmaking algorithm without/with network latency considerations.

Fig. 2: Matchmaking results for MOBAs.

A. Socially and Network-Aware Matchmaking

Matchmaking players at the start of a game can significantly impact the gameplay experience. Gaming services that perform matchmaking while taking into consideration network latency are already deployed by game operators. In contrast, a *socially-aware* matchmaking service assigns players

to matches, trying to ensure that players in the same social, rather than latency-based, cluster play together. We revisit the example of a socially-aware matchmaking service presented in [4], by also considering network latency.

Socially-aware matchmaking algorithm First, for each sliding window ($\tau = 10$ min. interval), the algorithm builds a list of all the players who are online. Second, from the social graph the algorithm computes the cluster membership for each player. Third, from the largest online players’ cluster to the smallest, all online players from the same cluster are assigned to new matches if size permits; otherwise, the cluster will be divided into two parts and players from one part will be assigned into new scheduled matches. Figure 2a sketches the algorithm for computing the score for an exemplary match. To favour small clusters, which can lead to novel human emotions [3], our scoring system does not consider the largest cluster when assigning points.

We compare our matchmaking algorithm with the algorithms observed in practice in MOBAs in terms of average scores (*utility*), and show selected results in Figure 2.

Expectedly, random matchmaking, which is still employed by many gaming communities, leads to very low utility. Surprisingly, our simple socially-aware matchmaking algorithm also exceeds the performance of the matchmaking algorithm employed by the operators of DotAlicious; this is because the limited community tools available in practice do not make all players aware that some of their friends are online and thus allow them to join other, lower-utility, matches.

Including network characteristics We use the geographical location gleaned from MOBA datasets to estimate possible latency conflicts, e.g., same-match players located in Germany and Asia. We analyze the impact of network latency on the score of our matchmaking algorithm and depict the resulting score in Figure 2b. In this scenario, a significant part of the matchmaking score is lost due to recommendations not taking into account network latency (yet our matchmaking algorithm still outperforms the original matchmaking). We conclude that combining social and network awareness is important for networked gaming services.

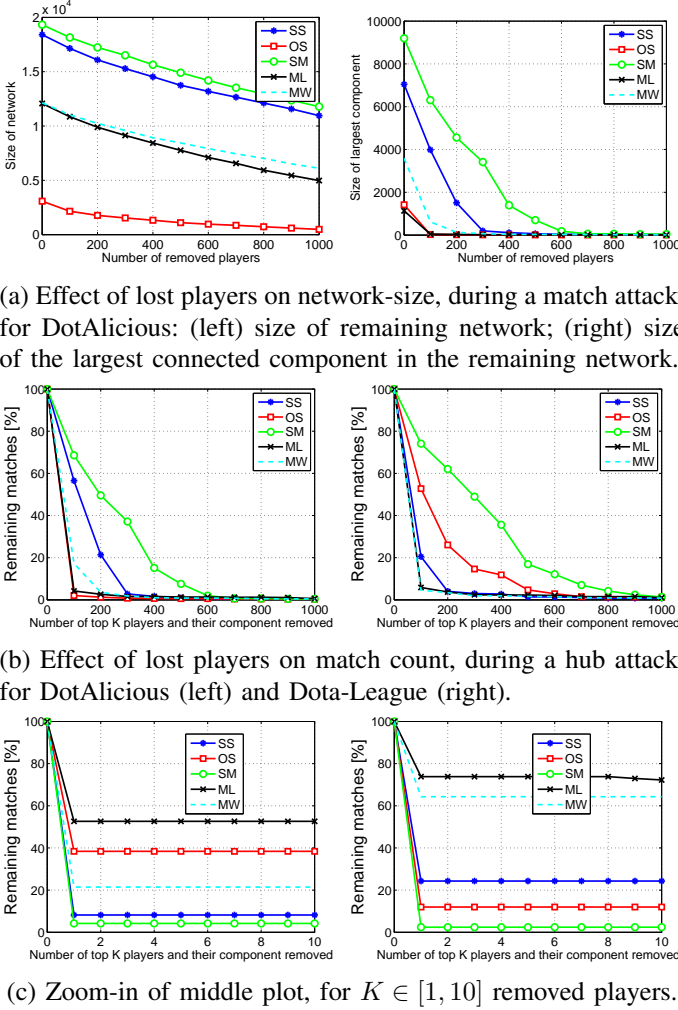


Fig. 3: Results of match and hub attacks on the social network, for $n = 28$ and $K \in [1, 1000]$.

B. Assessing Social Network Robustness

Because social relationships are important in player retention [3], the strength of the social structure may be indicative for the survival chances of the community. If the network starts dismantling, people might lose interest in the game and stop playing. Operators need to assess both the strengths and weaknesses of their games' social structure, to be able to stimulate growth or to prevent a collapse. Conversely, competitors could try to lure away key players (*hubs*), who in turn could sway others.

The anatomy of an attack: To assess the social-network robustness, we conduct a *threshold-based degree attack* on the network: for each mapping, we iteratively remove the *top-K players*, according to their degrees in the extracted graph, in decreasing order. Removing a player either also removes their matches (*match-attack*) or also removes their entire

connected component (*hub-attack*). Then, we re-apply the mapping to the remaining matches to get a new network, and output the size of the new network and largest component. We perform match and hub-attacks on DotAlicious and DotA-League and depict selected results in Figure 3. (We do not conduct experiments in which players form new clans (network rewiring), which represents the opposite of our scenario; in our experience as gamers, when a member of a strongly connected group leaves (for another game), the whole group departs as well.)

The aftermath of an attack: We find that both match and hub-attacks on MOBAs are very efficient. For match-attacks (Figure 3a), removing the top-1,000 players (1.5%) can reduce the size of the network by 15% up to 60% of its initial size, and the size of largest component to below 10. For hub-attacks (Figure 3b), removing only the top-100 players can cause the network to implode. A social-network collapse also implies the collapse of network traffic, which may lead to waste of pre-provisioned networked resources.

We conclude that understanding the social relationships between players can help a game operator improve the social-network robustness, by identifying and motivating the key players. Our formalism provides important tools for the former, but the latter remains open.

VI. ON CURRENT AND FUTURE SNGs

Many current networked games provide limited social-networking features, yet rely on their players to self-organize. For example, games in the popular class of MOBA-networked games have fostered the creation of many communities of players. In this work, we have shown how a general formalism can be used to extract social relationships from the interactions that occur between networked-game players. We have investigated their implicit social structures based on six types of interactions, using community traces that characterize the operation of four popular MOBA, RTS, and MMOFPS games, and provided hints on improving gaming-experience through two socially-aware services.

The field of social-networks research applied to networked games is rich and could lead to important improvements in gameplay, with direct repercussions to networked-resource consumption and quality-of-experience. We identify several challenges and opportunities related to our study:

- 1) *Expanding the formalism*: The mappings-set could be expanded, to provide a richer framework for implicit relationships. The framework could focus on temporal aspects such as loose (dense) interactions over long (short) periods of time.
- 2) *Complementing our work with social/other theory*: The pro-social emotions appearing in games may have important implications. It would be beneficial to explore them. From our datasets one could infer finer-grained relationships from the combination of explicit friendship relationships and implicit interaction graphs, and to test them against theories developed for complex networks, sociology and psychology.
- 3) *Applying the formalism to networked-game services*: The main purpose of this work is to provide support for (future) social game-services. We anticipate use in: player management and retention, through matchmaking recommendations and identification of key players; the design and tuning of capacity planning and management systems, through prediction of graph evolution; etc.

VII. DATASETS

Our datasets are available through the Game Trace Archive [2].

VIII. ACKNOWLEDGMENTS

This research has been partly supported by the EU FP7 Network of Excellence in Internet Science EINS (project no. 288021), the STW/NWO Veni grant @larGe (11881), a joint China Scholarship Council and TU Delft grant, and by the National Basic Research Program of China (973) under Grant No.2011CB302603.

REFERENCES

- [1] EDGE Team. Rise of the Ancients. *EDGE*, 257:14–15, 2013.
- [2] Y. Guo and A. Iosup. The game trace archive. In *NETGAMES*, 2012.
- [3] J. McGonigal. *Reality is Broken: Why Games Make us Better and How They Can Change the World*. Jonathan Cape, 2011.
- [4] R. van de Bovenkamp, S. Shen, A. Iosup, and F. A. Kuipers. Understanding and recommending play relationships in online social gaming. In *COMSNETS*, 2013.
- [5] C. Wilson, B. Boe, A. Sala, K. P. N. Puttaswamy, and B. Y. Zhao. User interactions in social networks and their implications. In *EuroSys*, 2009.

AUTHOR BIOS

Dr. Alexandru Iosup is currently Assistant Professor with the Parallel and Distributed Systems (PDS) group at TU Delft, where he received his Ph.D. in 2009. As personal grants, he was awarded a TU Delft Talent Award (2008) and a Dutch NWO/STW Veni grant-equivalent to NSF Career for Young Researchers (2011). He is a member of the ACM, IEEE, and SPEC, and Chair of the SPEC Cloud Working Group. His research on large-scale distributed systems has received several awards. He is also interested in research on undergraduate and graduate higher education for computer science.

Ir. drs. Ruud van de Bovenkamp is currently on a Ph.D. track with the Network Architectures and Services (NAS) group at TU Delft. His research focuses on complex networks, including those forming in multiplayer online games.

Siqi Shen, Eng., is currently on a Ph.D. track with the PDS group at TU Delft. His research focuses on massivizing online games, including through the use of the social networks forming in multiplayer online games to better provision and allocate servers and other resources.

Dr. Adele Lu Jia is currently a post-doc with the Parallel and Distributed Systems Group at TU Delft, where she received her Ph.D. in 2013. Her research focuses on the social aspects of large-scale distributed systems, applied to peer-to-peer data-sharing networks.

Dr.ir. Fernando Kuipers (IEEE SM) is currently Associate Professor with the NAS group at TU Delft, where he received his Ph.D. (cum laude) in 2004. His research on routing and network algorithms, and on complex networks has received several awards. He is also interested in quality of service and of experience. He has been TPC member of several conferences, among which IEEE INFOCOM, and is member of the executive committee of the IEEE Benelux chapter on communications and vehicular technology.