Trust Amongst Rogues? A Hypergraph Approach for Comparing Clandestine Trust Networks in MMOGs

Muhammad Aurangzeb Ahmad[*]**, Brian Keegan**[†]**, Dmitri Williams**[^]**,**
Jaideep Srivastava[*]**, Noshir Contractor**[†]

[*]Dept. of Computer Science & Engineering, University of Minnesota; [†]School of Communication, Northwestern University;
[^]Annenberg School of Communication, University of Southern California
{mahmad,srivasta}@cs.umn.edu; {bkeegan, nosh}@northwestern.edu; dmitri.williams@usc.edu

Abstract

Gold farming and real money trade refer to a set of illicit practices in massively multiplayer online games (MMOGs) whereby players accumulate virtual resources to sell for "real world" money. Prior work has examined trade relationships formed by gold farmers but not the trust relationships which exist between members of these organizations. We adopt a hypergraph approach to model the multi-modal relationships of gold farmers granting other players permission to use and modify objects they own. We argue these permissions reflect underlying trust relationships which can be analyzed using network analysis methods. We compare farmers' trust networks to the trust networks of both unidentified farmers and typical players. Our results demonstrate that gold farmers' networks are different from trust networks of normal players whereby farmers trust highly-central non-farmer players but not each other. These findings have implications for augmenting detection methods and re-evaluating theories of clandestine behavior.

"A plague upon't when thieves cannot be true one to another!" – Falstaff, *Henry IV, Part 1*, II.ii

Introduction

Gold farming refers to a set of practices that involve the sale of virtual items and currency within massively multiplayer online games (MMOGs) for offline, "real-world" money. Players who lack the time or desire to accumulate their own in-game capital can trade real currency to buy "farmed" currency and items to advance more quickly through the game (Heeks 2008). The size of the market for these gold farming and real-money trade services is not inconsequential: analysts estimate the industry generates between $100 million and $1 billion in revenue annually (Castronova 2008, Lehtiniemi 2007).

However, these exchanges undermine meritocratic norms, upset in-game economic equilibria, and raise complicated legal questions about property, taxes, torts, and labor (Dibbell 2003). Because of these reasons, game administrators attempt to ban gold farmers by observing unusual game activity or investigating reports from other players. However, these detection methods are ad-hoc and —as with criminals in the offline world—many gold farmers escape detection. But the ability to collect exhaustive longitudinal digital trace data on organizations operating under similar motivations and constraints as offline clandestine organization suggests that social behavior in MMOGs can also potentially be mapped back to test and inform theories clandestine social behavior and organization in offline contexts (Williams 2010).

The ability for players to grant other players permission to enter their in-game houses, move objects around in them, or even remove objects from the house is a ready proxy for the level of trust amongst characters. However, these permissions require modeling the relationships among houses, in-game characters, and the user accounts which own each. To capture these complex inter-dependencies, we employ a hypergraph to model triparite relational structures. We define and extract a variety of hypergraph projections for network analysis and compare the graph structures of farmers to typical players and unidentified gold farmers. We extend a label propagation approach developed by Keegan, Ahmad, et al. (2010) to compare the trust network structures of gold farmers, their undetected affiliates, and normal players.

Our findings demonstrate that gold farmers' housing permission behavior has distinct patterns when compared to the general player population as well as farmers who have yet to be detected by the game operator. We conclude by discussing the implication our findings have for augmenting detection methods in MMOGs and evaluating theories of clandestine organization.

Motivation and Background

Traditional analyses of trust networks have mainly focused on trust between people who come together in a certain context to achieve a certain goal or to connect with other people such as recommendation systems, friendship, and resource sharing (Golbeck 2008). In trust-based recommendation networks like FilmTrust (Golbeck 2006) and Epinions (Massa, et al. 2005) trust is measured with respect to the reliability and validity of a recommendation.

However, trust relationships are also prerequisites to a variety of instrumental communication and exchange relationships in clandestine organizations (von Lampe and Johansen 2004). Clandestine organizations assume a variety of forms in response to the shared goals, task demands, and skills of its members as well as trade-offs between resilience and flexibility while avoiding detection (Morselli, Giguere, et al. 2007). Therefore, understanding how trust operates in these organizations can inform the processes which govern other social dynamics in clandestine organizations.

Prior work on detecting gold farming has attempted to use classification techniques to identify cross-sectional and behavioral signatures of gold farmers (Ahmad, Keegan, et al. 2009). Like offline clandestine organizations, gold farmers also rely on peripheral and presumptively legitimate accounts which go unsanctioned by game operators. Research has demonstrated gold farmers' trade networks rely strongly on these undetected intermediaries to support and enable their operations (Keegan, Ahmad, et al. 2010). However, these analyses overlook the role of trust in mediating relationships in MMOGs and clandestine organizations.

While there is a large literature on trust in social networks, Golbeck (2009) notes that work comparing different networks in the same study are relatively rare. Ahmad, et al. (2010b) describe the network characteristics of various trust networks for comparative purposes and observed that trust networks which are generated by similar social processes have similar network characteristics. Our research fills in these gaps in three ways. First, we implement a hypergraph model to capture a variety of complex network relations. Second, we use projections of this hypergraph allow us to do multilevel comparisons of structures of various relationships between account owners, their characters, and objects within the game. Third and finally, we employ a label-propagation technique to not only compare network structures but also the behavioral patterns of three classes of users: identified gold farmers, unidentified gold farmer affiliates, and traditional players.

Housing Permissions as Trust in EverQuest II

We use data from EverQuest II (EQII) which is a MMOG which occurs in a fantasy role-playing universe. It is important to make distinction between accounts, characters, and houses. Each account can create several characters, but these cannot be played simultaneously. Each character has the option to buy a virtual house in the game. Thus houses are connected to players which are in turn embedded within accounts. Players can use their houses for a variety of purposes such as displaying valuable items, storing excess inventory, and selling crafted goods.

By default, only the character who buys the house has access to the house. However, a character may grant different levels of access to other characters in the game. In EQII the following access levels, in ascending order of trusted access, are defined:

- None: Has no access and cannot enter the house.
- Visitor: Can enter the house and can interact with objects in the house.
- Friend: Has all the privileges of the Visitor and can move things around the house.
- Trustee: Has all the privileges of the Friend and can add and remove objects in the house. A Trustee can also pay the rent of the house on the behalf of the owner of the house.

From a security perspective, all the access levels except trustee are functionally equivalent because characters who are given that type of access cannot make any change to the value of the house while a character with trustee privileges can make such a change. To simplify our analysis along these functional lines, we dichotomize these three potential types of relations into trustee and non-trustee (visitor and friend).

Hypergraph Model of Housing-Trust Network

The housing-trust network can be modeled in different ways. Previous research using housing-trust networks has looked at the structure of the housing network in terms of access-grants while ignoring the presences of houses or even permissions for multiple characters (Ahmad et al 2010a, 2010b). While these models are sufficient for studying the social networks amongst the gold farmers, they limit the types of inferences that can be made about the larger trust-based social structures and the use of such structures for making inferences about gold farmers.

Our approach follows previous work using hypergraphs to model tagging systems where there is a "natural" distinction between three types of nodes in the networks such as person, tag and object (Zlatic, et al. 2009). We also adopt a hypergraph model to describe the three types of nodes in our data: player account, player character, character house. Multiple models of hypergraphs exist which describe the evolution and generation of such hypergraphs (Ghosal, et al. 2009).
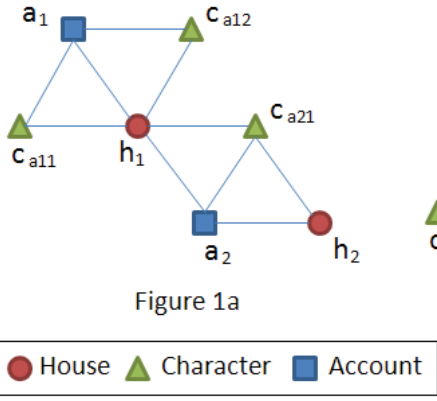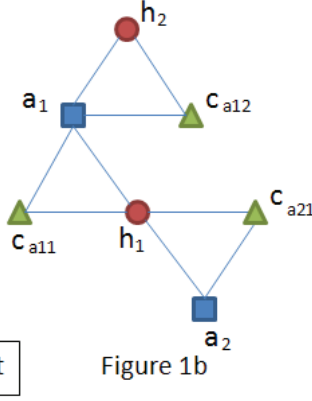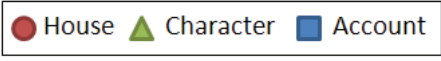
Figure 1: Different scenarios for housing access for characters associated with (a) the same account and (b) with different account

Figure 2: Alternative ways to represent the housing-trust network

The complex game mechanics of EQII which cannot be captured by a traditional graph representation are another motivation for using hypergraphs to model trust relationships. Players at each level not only have these privileges associated with that level, they also have the privilege to grant the same or lesser level of access to other people. Thus consider the situation in Figure 2a and Figure 2b which ignores the player accounts for simplification purposes. In the first case character c_{a11} trusts c_{a22} and c_{a31} trusts. From this representation it is not clear if there is a trust relationship between c_{a11} and c_{a31}. While it could be the case that c_{a11} also trusts c_{a31} but since c_{a22} has already granted permissions to c_{a11} it is not necessary for c_{a11} to grant permissions to c_{a31}. However given that there is still a possibility that c_{a11} instructed c_{a22} to grant access to c_{a31} e.g., c_{a11} is a superior officer of c_{a22}, an important piece of information is lost. One way to remedy would be to add an edge between c_{a11} and c_{a31} but even in this case we will lose information about which players are connected with each other by which house. We use the alternative projection in Figure 2b wherein player nodes are connected by access ties to house nodes. Even in this case some information is also lost such as how the access grants were given but since we are interested in the relationship between houses, players and characters this can be overlooked.

A hypergraph is a generalization of a graph (Dauber 1969) and can be defined as follows:

Tripartite Hypergraph: A tripartite hypergraph $G = (V,H)$ consists of a set of nodes V and a set of hyperedges V such that the following conditions are satisfied.

1. $V = \{V_h, V_c, V_a \mid V_i \cap V_j = \Phi\}$
2. $H \in \{(v_h \in V_h, v_c \in V_c, v_a \in V_a)\}$

Figure 1a shows a hypergraph which contains hyperedges (a_1, c_{a11}, h_1), (a_1, c_{a11}, h_1), (a_1, c_{a11}, h_1) and (a_1, c_{a11}, h_1). **Node Degree:** The degree of the nodes can be defined in a number of ways. One can define it in terms of how many other nodes is a node connected to. However in this case no di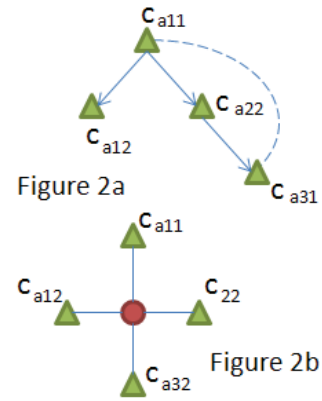stinction is being made between the various types of nodes tha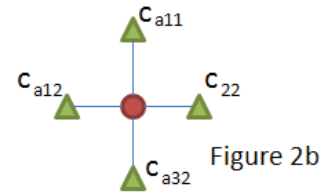t may be present in the hypergraph and in the current domain the semantics of the graph will be lost if such an approach is used. Another approach which is more suited to our present context is to define node degree in terms of the hyperedges that are connected to a node. Thus in Figure 1a the degree of h_1 is 3 and the degree of h_2 is 1.

Edge Degree: In addition to the node degree, it is also possible to describe the edge degree in the hypergraph (Zlaic 2009). The edge degree is defined as the number of hyperedges in which the edge participates in. Consider edge (a_1, h_1) in Figure 1a, it has edge degree two because it participates in two different hyperedges (a_1, c_{a11}, h_1) and (a_1, c_{a11}, h_1). **Projections of a Hypergraph:** There are multiple ways in which hypergraph projections can be formed e.g., one way to create a projection would be to create an edge between two nodes if they share a house, another way to project would to create a node if they share an account. It is also possible to create a double projection by projecting onto a projection (Zlatic 2009).

In order to distinguish between the characteristics of gold farmers and legitimate players we consider the frequent subgraph patterns which are associated with different types of players. We now describe various terms which would be helpful in finding such patterns.

Frequent Tripartite Hypergraph Pattern: Given a tripartite graph H with nodeset N and an edgeset E, a frequent tripartite hypergraph patterns is a sub-hypergraph sub of graph H such that it occurs frequently in H with a support S, confidence C and at least one of the nodes containing a label P. Since the dataset that we are dealing with is not a transaction dataset the definitions of support and confidence are modified accordingly. The support and confidence are defined as follows:

Support of a Hyper-subgraph: Given a sub-hypergraph of size k, sub_P is the pattern of interest containing the label P, sh_P is a pattern of the same size as sub_P and contains the label P, the support is defined as follows:

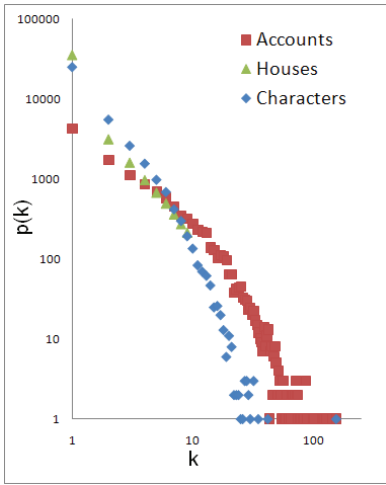$$S = \frac{|sub_P|}{|\{sh_P \mid sh_P \subseteq H, |sh_P| = k\}|}$$

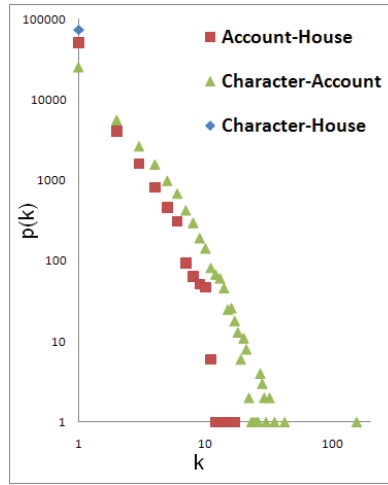**Figure 3a. Distribution of node degree for the trust hyper graph.**

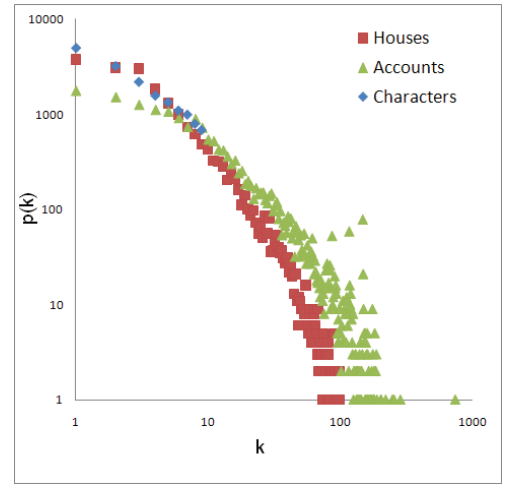**Figure 3b. Distribution of edge degree for the trust hypergraph**

Figure 3c: Distribution of the projection networks

Confidence of a Hyper-Subgraph: Given a sub-hypergraph of size k, sub_P is the pattern of interest containing the label P, sub_G is a pattern which is structurally equivalent but which does not contain the label P, the confidence is defined as follows:

$$S = \frac{|sub_P|}{|\{sub_G | sub_G \subseteq H, |S| = k\}|}$$

Frequent Tripartite Hypergraph Pattern Mining: We now describe a technique which can be used to extract frequent tripartite hypergraph patterns, with and without constraints, from our data. Consider the hypergraphs in Figure 1; it is clear that a hypergraph can be visualized as a graph with a larger number of triads. This implies that there is already implicit structure in the data which can be exploited for pattern mining. The task of mining such patterns can thus be formulated as discovering triads in a 3-Regular graph with certain constraints.

We now describe the problem of discovering the frequent patterns described in the previous discussion. Consider the hypergraph in Figure 1a, if we consider the triads which are connected to h1 then these are (a_1, c_{a11}, h_1), (a_1, c_{a12}, h_1) and (c_{a21}, a_2, h_1). Given that it can be treated as a 3-Regular graph, we know can describe the structure of the neighborhood of h_1 in terms of connectivity of the accounts. For example, account a_1 is connected to h_1 with two characters, account a_2 is connected to h_1 with one character. We can represent the neighborhood of h_1 as (2CH0, 1CH1) where A and C signify accounts and characters respectively. The representation can be further extended by considering the other houses to which a node may have access to. Thus in Figure 1b the neighborhood of h_1 would be represented as (C2H1, C1H0) which show that the representation of the neighborhood of h_2 would be (C2H0). Even with this representation there can be multiple ways to represent the same graph since there are multiple ways to traverse a graph. To address this issue we

represent the subgraphs in the DFS Lexicographical order (Yan, et al. 2002). Of course in this type of representation some information is lost. However with this representation standard association rule mining techniques can be applied to discover useful discriminative patterns in the data as we demonstrate in the analysis section.

Dataset

We use anonymized housing-trust data from EQII provided to us by Sony Online Entertainment. The data consists of more than two million player characters spread across over a dozen servers running parallel game worlds with slightly different rule sets. We use data from a single representative server with a player vs. environment (PvE) rule set encompassing January through September 2006.

The dataset contains 38,217 characters associated with 12,667 accounts, with 43,548 houses and a total of 3,013,741 hyperedges between them. 151 of these accounts were banned by SOE administrators for reasons related to gold farming. A small number of records (105 accounts, 482 characters) were discarded because of incomplete transcription of data. However none of the houses were discarded in this case. The "Trustee" access was granted 20,029 times, the "Friend" access was granted 32,711 times and the "Visitor" access was granted 273,355 times for all the players in the network. Additionally there were 8,295 instances where the trust privileges were revoked. We note that these counts sum up to be greater than the number of edges in the network because there were many redundant instances where the same access was granted to the same person on the same house multiple times.

Figure 3a gives the node degree distribution of the various types of nodes on a log scale. It is clear from the figure that the majority of the accounts have fewer than

	Neighbors' total degree			Neighbors' in-degree			Neighbors' out-degree		
	$\langle n \rangle$	$\langle n_{GF} \rangle$	$\langle n_{Aff} \rangle$	$\langle n_i \rangle$	$\langle n_{i,GF} \rangle$	$\langle n_{i,Aff} \rangle$	$\langle n_o \rangle$	$\langle n_{o,GF} \rangle$	$\langle n_{o,Aff} \rangle$
Farmers	1.82	0.29	1.82	0.89	0.29	0.89	1.07	0.29	1.07
Affiliates	4.03	1.28	0.70	1.55	0.75	0.70	2.88	0.625	0.70
Non-Affiliates	2.73	-	7.77	1.57	-	5.98	1.56	-	2.34

Table 1: Average neighbor connectivity for gold farmers, affiliates and non-affiliates.

four houses and character pairs associated with them. Similarly, the same applies for the characters as well. While the distributions for the accounts and the characters follow a long-tail distribution, the distribution for the houses is linear with a maximum of 8 character-account ties. We note that this is not a constraint in the game. Similarly Figure 3b gives the edge degree distributions for the various edge types. In this case also the account-house and the character-house distributions follow a power law more or less. The character-house edges always have a degree of one because there is a unique mapping from a character to an account in the game.

Analysis of the Housing-Trust Network

Using a label propagation technique derived from Keegan, Ahmad, et al. (2010), we distinguished between three types of players based on their relationship with identified gold farmers in the housing-trust network.

- **Gold farmers:** These are characters who are explicitly labeled as gold farmers in the data.
- **Gold farmer affiliates:** These are characters who have interacted with the gold farmers by either extending housing permissions to gold farmers or are trusted by other gold farmers but they are not labeled as gold farmers themselves. Using our "guilt-by-association" label propagation technique, we assume these characters have a much higher likelihood of being unidentified gold farmers.
- **Non-affiliates:** The rest of the characters who are neither gold farmer nor affiliates.

Table 1 reports the average neighbor connectivity of the three types of players. Here *n* refers to all the neighbors regardless of farmer/affiliate attribute, n_i refers to neighbors with incoming edges and n_o refers to neighbors with outgoing edges. From the table it is clear that gold farmers grant or receive permission from fewer players (1.82) than their affiliates (4.03).

The second column n_{GF} refers to neighbors who are gold farmers. In this case gold farmers also have very low tendency to grant other gold farmers permission (0.29). n_{Aff} refers to the neighbors of affiliates. Here the connectivity patterns of affiliates stand out markedly; on average, non-affiliates have granted housing permission to 7.77 affiliates even though affiliates intra-class connectivity (0.70)

suggests they are unlikely to give other affiliates housing permissions.

On average non-affiliates give 5.98 affiliates housing permission while affiliates only reciprocate by giving permissions to 2.34 affiliates on average. We also see that although gold farmers have relatively low base rates for granting housing permissions to other players, they appear to be strongly averse to granting other gold farmers access. Instead, gold farmers appear to both grant (0.89) and receive (1.07) permissions at a substantially higher rate than they are granted (0.29) or received (0.29) from other gold farmers. As the title of the paper indicates, there appears to be little honor among thieves.

These findings have several important implications. First, housing access appears to serve a non-trivial role in enabling gold farming operations as affiliates and farmers alike avoid granting permissions to characters of the same type. Second, the affiliate players whom gold farmers grant permissions are also players who themselves have high connectivity with the rest of the network. Third, farmers do not grant housing permissions at all to non-affiliates. Clearly the affiliates play a crucial and trusted role in brokering between identified farmers and the general population while isolating themselves from the general player populations. This corroborates previous findings by Keegan, Ahmad, et al. (2010) about differences in centrality between character classes in the trade network. A possible explanation is that these affiliates are gold farmers themselves but they have not been caught by the game administrators and thus the data does not label them as such. However given that affiliates are so strongly trusted by farmers, it could be the case that the gold farmers grant this access as a conduit for distributing their goods via trusted channels. In either case, there is a clear implication that affiliates are an integral part of the gold farming supply chain.

To explore the connectivity of gold farmers in the data, we extracted tripartite hypergraph patterns occurring frequently in the data for the three types of players using standard pattern mining techniques (Agarwal 1994). Most of the patterns which were obtained for gold farmers had a very low support and confidence and only 8 patterns had support and confidence greater than a standard 0.1 threshold. Because of the limitation of space only two most frequently occurring patterns are shown in Figure 4. Figure 4a refers to a pattern where a house is shared by three players two of whom have many characters associated with
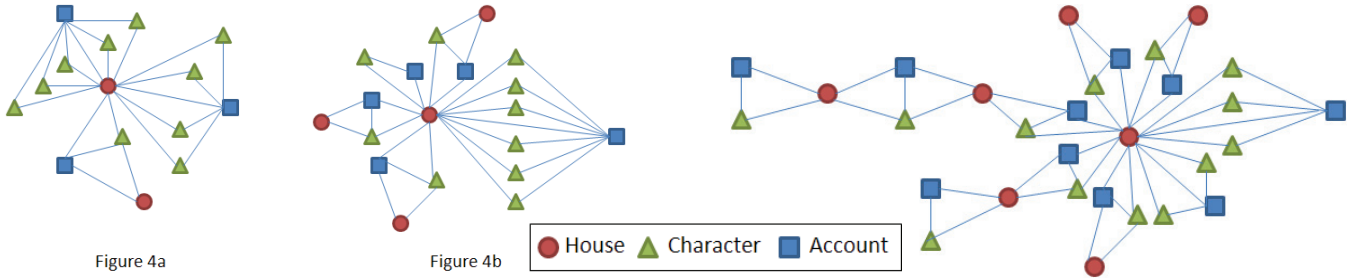
Figure 4a. Gold Farmer Hypergraph Pattern: Support = 0.33, Confidence = 1 Figure 4b. Gold Farmer Hypergraph Pattern: Support = 0.50, Confidence = 1 Figure 4c. Gold Farmer Affiliate Hypergraph Pattern: Support = 0.50, Confidence = 1

their respective accounts and the third player has access to another house. Figure 4b on the other hand shows a situation where a player has many characters and all the characters have access to the house but at the at the same time there are other players who have access to that house but they only have one character and also have access to another house. Both evoke a house being used as a shared, central safehouse shared by many farming character-accounts but also with connections to affiliate character-accounts with access to other houses.

We also extracted the patterns which were associated with the various affiliates and surprisingly a third (15/44) of the sequence patterns with more than 10 nodes were associated with affiliates. We note that these patterns are too long to visualize here, an example of a smaller pattern is given in Figure 4c. Like Figures 4a and 4b, there is a clear star-like structure with several affiliate character-accounts sharing a house, but select few having access to other houses as well. The earlier observation that gold farmer affiliates are highly connected players is borne out here as gold farmers connect to trustworthy affiliates but avoid directly granting trust to each other.

Hypergraph Projection for the Network of Accounts:
As noted earlier, it is possible to create projections of the hypergraph for different node types in the network and determine the prevalence of gold farmers in each network. The characteristics of the various projections are given in Table 2. Here NCC refers to the number of connected components, LCC refers to the size of the largest connected component and %LCC refers to the percentage of the total nodes which are part of LCC. We now describe the various projections of the hypergraph H. The node-degree distributions of these graphs are given in Figure 3c.

If we consider the subgraph which consists of the gold farmers, their affiliates and the neighbors of the affiliates then we observe that the majority (79%) of these accounts are isolates. There are a large number of instances of gold farmers where the gold farmer have exclusive access to the houses without giving access to other players including other gold farmers. On the other hand if we consider the affiliates then again they have a very high connectivity 8.89 as compared to both the gold farmers 0.31 as well as the non-affiliates 3.47. This again reinforces the

observation that gold farmers do not trust one another but they trust other people who are trusted by the population in general.

Hypergraph Projection for the Network of Characters:
The projection of characters is the projection of the accounts and the houses in the networks. The same phenomenon of gold farmers not connecting to other gold farmers is also observed which a large percentage (84%) of gold farmer nodes being isolates. In both the cases of the projection of the accounts as well as the projection of the characters, the degree to which gold farmers are connected to one another is quite low which reinforces the conclusion that sharing houses and thus trust across gold farmers is not very common. The affiliates again have a very high connectivity (10.42) as compared to the rest of the population (3.23).

Network Project.	Nodes	Edges	NCC	LCC	% LCC
Account	18,231	159,676	1,015	14,431	79.16
Character	16,878	119,757	1,070	13,111	77.68
House	19,832	83,715	1,764	14,801	74.63

Table 2: Global Characteristics of the Projection Networks of the Hypergraph H

Hypergraph Projection for the Network of Houses:
Another way to project the hypergraph H is to project the accounts and the characters so that we get a η of the houses in the network. In the projected House network η there are 43,548 nodes and 83,715 edges. There are 521 gold farmer houses which we define to be a house having a direct connection with a gold farmer. However many houses associated with gold farmers are isolated nodes. Table 2 shows that there are a large number of components (1,764) but a single giant component contains three-quarters of the nodes. The rest of the components are relatively small – the second largest connected component has 30 nodes. Thus the smaller components in Figure 6 are indeed isolated components and the large component is part of the largest connected from the original component. It is clear that there are many cases where the gold farmers's houses form isolated groups. The most prominent examples are the two components in the upper right side of Figure 8 with
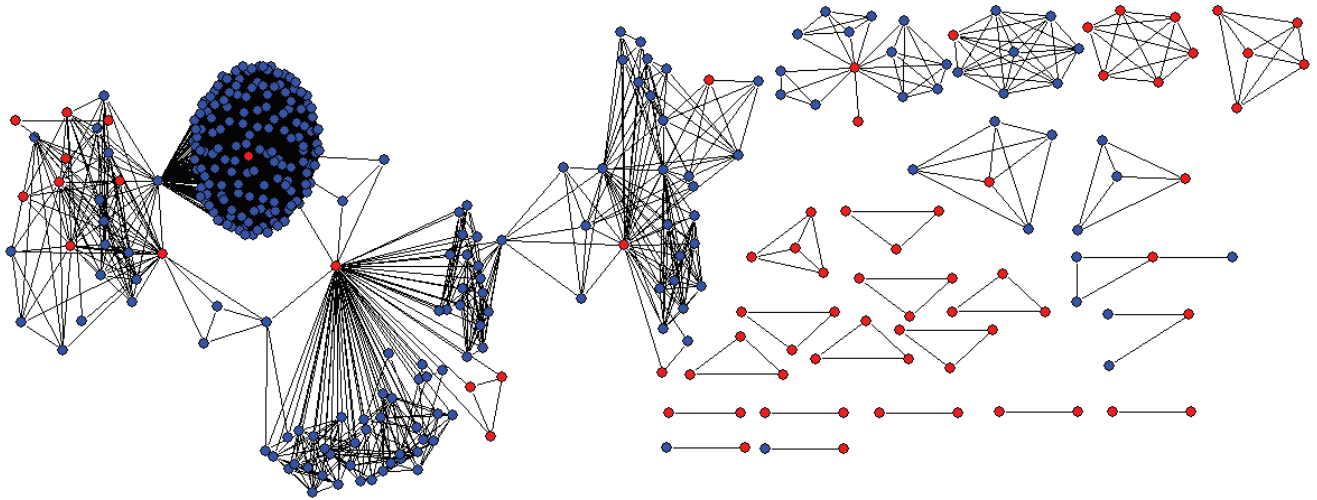
15

Figure 6: The house projection displaying houses associated with gold farmers (in red) and houses associated with affiliates (in blue). The network is obtained by the projection of the Hypergraph *H*

farmers' houses having access to other farmers' houses. In the larger component, at least four main clusters are easily identifiable. In there are cases where the gold farmers' houses are almost at the site of cut vertices and join a large number of other houses on the either side. These are promising candidates for gold farming distribution centers.

In terms of connectivity the house projection network is much more highly connected. The average connectivity for gold farmers is 7.56, non-affiliates is 7.09 and for affiliates it is extremely high: 84.02. This implies affiliates' houses are connected to a large number of other houses. On average gold farmer houses are connected to 5.86 other gold farmer houses but the average connectivity with non-affiliates is 21.88. This again reinforces the idea that gold farmers tend to trust only the individuals who are trusted in general but not other gold farmers.

Discussion

Our results provide novel insights into the trust networks which exist among players engaged in clandestine behavior in an online game. Using a hypergraph model to capture the complex dependencies and relationships between accounts, characters, and houses, we performed network analyses on projections of this hypergraph's to identify behavioral patterns of granting and receiving trusted access among farmers, affiliates, and general player population.

We showed that the distribution of links in the hypergraph is very heterogeneous and follows a long-tailed distribution such that most of links in the housing network are concentrated in a few nodes. These distributions arise in a variety of other complex networks and suggest an underlying preferential attachment process (Newman 2003).

Examining this topology based upon the types of accounts, characters, and houses, we found that gold farmers preferentially grant trusted housing access to affiliates who remain undetected rather than to other farmers. These affiliates, in turn, are strongly connected to the rest of the network. The strong disparities between farmers and affiliates' housing permissions behavior compared with the general player population suggests these selective patterns capture trust-based relationships. Permissions appear to serve an instrumental purpose in enabling farming operations and avoiding detection.

Using frequent subgraph mining techniques, we also identified structural patterns in the hypergraph associated with farmers To the extent that they capture underlying trust among members of these clandestine organization, these frequent subgraphs reveal the strategies adopted to conceal their operations. It may be possible to develop detection algorithms to identify these patterns and improve predictive models.

To Sir Falstaff's lament referenced in the introduction, because gold farmers avoid granting trust permissions to other gold farmers, our results seem to suggest that our "thieves" are in fact rogues among themselves. However the absence of trust ties among these players may not reflect amoral opportunism on the part of this type of players but rather a principled survival instinct evolved and honed from prior encounters with authorities. Or, it could be a combination of both.

Nevertheless, gold farmers do not represent a monolithic behavioral class of players; like other criminal organizations, the dividends of comparative advantage lead to a division of labor and skill specialization. We expect that gold farming operations should in many ways resemble drug trafficking operation which need farmers to generate the raw material, distributors to package and deliver the goods, and dealers to interact with customers. Farming operations may exploit administrator heuristics—which only detect certain behaviors—to concentrate

essential but easily-identified behavior into expendable characters. These identified farmers may be "sacrificial lambs" serving an instrumental but easily replaced role in the operation as well as distracting administrators from identifying the latent organizational patterns we observed. The dissortative or heterophilic mixing we observed among player types could be a strategy employed by farmers to increase survivability of the organization by routing goods and services produced by farmers through complex relationships with other co-conspirators whom they trust will remain unidentified.

The generalizability of our findings and the extent to which they map to offline clandestine contexts crucially depends on the extent to which both contexts share the same affordances and constraints. On one hand, the costs of identification for gold farmers are largely pecuniary (re-creating a character) rather than physical (violent reprisal, imprisonment, etc.). On the other hand, previous work (e.g., Keegan, Ahmad, et al. 2010) has established striking similarities between online and offline clandestine networks which suggests the need for further comparative and situated research on how gold farmers operate.

Future research examining trust networks among clandestine organizations in MMOGs should emphasize generative rather than the descriptive models of behavior we employed. Agent based models, exponential random graph approaches, and stochastic actor-oriented models are all methods for generating graph structures based on local behavioral properties. Future work employing these methods permit the statistical testing of multilevel, multitheoretical hypotheses about processes governing the evolution of networks (Monge and Contractor 2003).

Acknowledgements

References

Agrawal, R. and Srikant. 1994. Fast algorithms for mining association rules in large databases. In *Proc. VLDB'94* Santiago, Chile: 487-499.

Ahmad, M.A.; Poole, M.S.; and Srivastava, J. 2010a, Network Exchange in Trust Networks. In *Proc. SocialComp'10*. Minneapolis, MN, USA. IEEE

Ahmad, M.A.; Borbora, Z.; Srivastava, J.; and Contractor, N. 2010b. Link Prediction Across Multiple Social Networks. In *Proc. ICDM'10*, Sydney, Australia.

Ahmad, M.A.; Keegan, B.; Srivastava, J.; Williams, D.; and Contractor, N. 2009. Mining for gold farmers: Automatic Detection of Deviant Players in MMOGs. In *Proc. SocialComp'09*, Vancouver, Canada. IEEE

Castronova, E. 2006. A cost-benefit analysis of real-money trade in the products of synthetic economies. *Info* 8(6): 51-68.

Dauber, E. 1969. *Graph theory*, ed. F. Harary. Addison Wesley.

Dibbell, J. 2003. "Owned! Intellectual property in the age of dupers, gold farmers, eBayers, and other enemies of the virtual state." *The State of Play: Law, Games, and Virtual Worlds*, eds. Balckin, J.M. and Noveck, B.S. New York University Press, New York.

Golbeck, J. and Hendler, J. 2006. Inferring binary trust relationships in web-based social networks. *ACM Trans. Internet Technology*, 6:497-529.

Golbeck, J. 2009. Computing with Social Trust. Springer.

Ghoshal, G.; Zlatic, V.; Caldarelli, G.; and M.E.J. Newman. 2009. Random hypergraphs and their applications. *Physical Review E*, 79:066118, 2009.

Heeks, R. 2008. Current Analysis and Future Research Agenda on Gold Farming: Real World Production in Developing Countries for the Virtual Economies of Online Games. Institute for Development Policy & Management, University of Manchester.

Keegan, B. Ahmad, M.A.; Srivastava, J.; Williams, D.; and Contractor, N. 2010. Dark Gold: Statistical Properties of Clandestine Networks in Massively-Multiplayer Online Games. In *Proc. SocialComp'10*, Minneapolis, MN, USA.

Lehtiniemi, T. 2007. "How big is the RMT market anyway?" Virtual Economy Research Network.

Massa, P. Avesani, P. 2005. Controversial Users Demand Local Trust Metrics: An Experimental Study on Epinions.com Community. In *Proc. AAAI'05*. 121-126

Monge, P. and Contractor, N. 2003. *Theories of Communication Networks*. Oxford University Press.

Morselli, C.; Giguere, C.; and Petit, K. 2007. The efficiency/security trade-off in criminal networks. *Social Networks*, 29(1): 143-153.

Newman, M.E.J. 2003. Structure and function of complex networks. *SIAM Review*, 45: 167-256.

von Lampe, K. and Johansen, P.O. 2004. On the conceptualization and empirical relevance of trust in the context of criminal networks. *Global Crime*, 6(2).

Williams, D. The Mapping Principle and a Research Frameowrk for Virtual Worlds. *Communication Theory*, 20(4): 451-470.

Yan, X. and Han, J. 2002. gspan: Graph-based substructure pattern mining. In *Proc. ICDM'02*, Maebashi City, Japan. IEEE: 721–724.

Zlatic, V.; Ghoshal, G.; and Caldarelli. G. 2009. Hypergraph topological quantities for tagged social networks. *Phys. Rev. E* 80:036118, 2009.