

# An Analysis of Underground Forums

Marti Motoyama, Damon McCoy, Kirill Levchenko, Stefan Savage and Geoffrey M. Voelker

Department of Computer Science and Engineering  
University of California, San Diego

## ABSTRACT

Underground forums, where participants exchange information on abusive tactics and engage in the sale of illegal goods and services, are a form of online social network (OSN). However, unlike traditional OSNs such as Facebook, in underground forums the pattern of communications does not simply encode pre-existing social relationships, but instead captures the dynamic trust relationships forged between mutually distrustful parties. In this paper, we empirically characterize six different underground forums — BlackHatWorld, Carders, HackSector, HackElite, Freehack, and L33tCrew — examining the properties of the social networks formed within, the content of the goods and services being exchanged, and lastly, how individuals gain and lose trust in this setting.

## Categories and Subject Descriptors

H.3.5 [Information Storage and Retrieval]: Online Information Services; J.4 [Computer Applications]: Social and Behavioral Sciences; K.4.4 [Computers and Society]: Electronic Commerce

## General Terms

Human factors, Measurement, Security

## Keywords

Underground forums, Online social networks

## 1. INTRODUCTION

Online social networks (OSNs) capture, in a concrete form, the character and dynamics of human social relationships. Consequently, the popularity of such services (e.g., Facebook, Twitter, etc.) has been followed closely by researchers using the explicit nature of these networks to characterize social graph properties and how they inform user interaction [6, 11]. While less well explored, there are also a range of *implicit* social networks defined via interaction on other shared interaction sites (e.g., Web forums, blogs, etc.). In this paper, we focus on a particular sub-population of such activities: *underground forums*.

Users of underground forums participate in many activities similar to those found on traditional online social networks: they maintain profiles, add fellow users to buddy lists, and engage in conversations via private messaging. However, the “raison d’etre” for

such forums is not simply for social contact, but to support criminal (or at best “grey hat”) activities. Thus, users of these forums regularly engage in the buying, selling and trading of abusive services and illegally obtained goods such as credit card numbers, online currencies, compromised accounts and even drugs. However, since underground users frequently *only* know each other online (and via pseudonyms even there), they must develop new means to establish trust among themselves.

In this paper, we examine these implicit social networks and how they are used in the context of six underground forums — BlackHatWorld, Carders, FreeHack, HackElite, HackSector, and L33tCrew — for which we have complete activity records. Our analysis is organized into three parts: first, we analyze the structure of the underlying social networks present on the forums, followed by an examination of the commercial aspects of the sites (e.g., what types of products are being sold, who are the most active players in the market, etc.) and finally we look at how different reputational factors impact behavior. We believe our work is the first analysis of this type and provides valuable insight into how online criminal actors create and develop social relationships in support of their goals.

## 2. BACKGROUND

Online underground markets have existed in various forms for decades. Early markets used Internet Relay Chat (IRC), documented by Thomas *et al.* [10] and Franklin *et al.* [4], to provide a public medium for sharing information about the availability and pricing of goods and services (e.g., stolen credit cards, accounts, botnets, cash out services, etc.).<sup>1</sup> Over time, many of these markets moved to using persistent Web forums and expanded to cover a broader range of information sharing. Zhuge *et al.* first documented the use of such forums in China [12] and contemporary analyses have been published by Holt *et al.* [5], Radianti [7] and Fallmann *et al.* [3]. Over time, some of these forums have specialized and many have moved to “closed” models (i.e., in which new members must be explicitly vouched for by existing members); for example, Stone-Gross *et al.* [9] recently documented the membership and goods on offer on the private Spamdott.biz forum, which specialized in support for email spammers. Ultimately, the goal of all such forums is to expand the knowledge base of the participants (e.g., which registrars will “look the other way”, how to best manipulate Google ranking results, etc.) as well as to expand the set of potential trading partners. However, there is little public research that empirically examines the social networks formed in such forums or the mechanisms employed to manage trust. Indeed, such analyses can be difficult since modern forums combine public sections, restricted sections (requiring higher status) and person-to-person private messages (PMs) that may not be externally visible.

<sup>1</sup>For a brief overview of the how this credentials market operates today, see Shilman [8].

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

IMC’11, November 2–4, 2011, Berlin, Germany.

Copyright 2011 ACM 978-1-4503-1013-0/11/11 ...\$10.00.

Forum	Abbrv	Dates Covered	Subforums	Threads	Posts	Pvt Msgs	Users	Lurkers
L33tCrew	LC*	May 07 – Nov 09 (30 mo.)	239	120,560	861,459	501,915	18,834	46.5%
HackSector	HS	Nov 01 – Nov 07 (72 mo.)	147	72,734	724,820	78,777	33,986	55.0%
FreeHack	FH	Jul 04 – Dec 10 (77 mo.)	152	62,972	499,736	112,318	38,377	62.9%
Carders	CC*	Sep 08 – Dec 10 (27 mo.)	121	52,188	373,143	197,067	8,425	35.0%
BlackHatWorld	BH	Oct 05 – Mar 08 (29 mo.)	38	7,270	65,572	20,849	8,718	47.9%
HackeL1te	HL	Mar 10 – Apr 11 (13 mo.)	43	5,501	9,018	541	2,431	66.7%

Table 1: Summary of the data from the six forums, ranked by number of posts (\* denotes forums geared toward commerce).

Forum	Buddy				Private Message				Thread			
	Partic.	Links	L/P	WCC	Partic.	Links	L/P	WCC	Partic.	Links	L/P	WCC
LC	2,587	4,448	1.7	214	7,898	170,954	21.6	7	9,124	3,791,330	415.5	1
HS	1,282	1,057	0.8	352	9,562	43,807	4.6	69	13,502	2,582,513	191.3	1
FH	1,921	5,944	3.1	100	10,294	55,945	5.4	21	11,833	1,473,824	124.6	3
CC	1,000	1,834	1.8	154	5,065	63,409	12.5	12	4,593	792,962	172.6	2
BH	199	205	1.0	37	3,438	11,183	3.3	3	2,940	320,028	108.9	2
HL	38	65	1.7	4	174	271	1.6	5	690	10,477	15.2	3

Table 2: Summary of the basic social networking statistics for each forum. *Partic.* means participants, or users who have links of the specified type; *L/P* represents the number of links divided by the number of participants. *WCC* means weakly connected components.

### 3. DATA OVERVIEW

In this study we have the luxury of “ground truth” — complete records of six underground forums via SQL dumps of their underlying databases. We do not claim that these six are representative of all underground forums, but they provide us with a starting point for understanding the dynamics of underground forums. Each of these datasets has been acquired by unknown outside parties and made public (“leaked”) via various methods. Each forum contains a wealth of information: user registration data, private messages exchanged, forum posts, member status changes, banned user logs, etc. For a more comprehensive list of the available data, please refer to the Invision Power Board (for L33tCrew) and vBulletin database schemas [1, 2]. We briefly describe the purpose of each forum.

BlackhatWorld (BH) was founded in approximately 2005 and is primarily English speaking. The main focus of BH is blackhat search engine optimization (SEO), a practice in which users attempt to abusively manipulate search engine algorithms to gain increased page rank. At the time our dataset was obtained, BH did not have a vibrant trading marketplace, as the site was initially oriented towards the discussion of blackhat techniques. Today, however, the site contains over 800 threads in the services-for-sale section and more than 275 threads in the goods-for-sale section.

Carders (CC) is a German-speaking site primarily focused on the monetization of stolen credit card numbers and bank account information. The site is heavily geared towards the exchange of goods and services. The L33tCrew (LC) forum is very similar to CC, both in its content and the types of products exchanged.

Freehack (FH) is another German site, but does not target any one industry. The threads on the forum cover a number of different topics, ranging from crypting (encoding software to make detecting malware more difficult) to video games. Items for sale include Steam (gaming) accounts, automatic account creators, and hacking software. The users typically do not buy or sell stolen credentials. The remaining sites, HackSector (HS) and HackeL1te (HL), are similar although HL is English speaking.

Table 1 summarizes the membership and activity across each of these forums. In total, our analysis covers over 2.5 million posts, 900k private messages, and 100k users. Our dataset also spans a range of time periods, with the FH and HS datasets covering approximately six years, while HL is our shortest at roughly a year. LC is our largest dataset by forum activity, with the largest number of threads (120k), posts (860k), and private messages (500k). The

six forums exhibit different properties with regards to the number of posts and private messages exchanged. The forums geared towards commerce, CC and LC, have a much higher number of private messages, since many business transactions occur over private messaging. For example, we observed 23–26 PMs/user for CC and LC, while the other four forums, whose users primarily swap information, exhibit less than 3 PMs/user. Across all forums a fair number of “lurkers”, or individuals who simply register an account but take no action, exist on each forum. Over 55% of the users on FH, HS and HL are lurkers. Again, the trading forums have a smaller fraction of lurkers with only 35% in CC and 46% in LC.


### 4. UNDERGROUND SOCIAL NETWORKS

In this section, we analyze the structural properties of the six different forums. This task is nontrivial in the context of a forum, since the definition of a link between nodes (users) remains ambiguous. To that end, we consider three types of relationships that exist in the forums: buddy, private message, and thread. Table 2 summarizes the basic social networking statistics for each forum.


A buddy link is the most explicit relationship that exists between users, and is a directed link since buddy requests may be accepted, pending, or denied. Thus, accepted requests result in symmetric links, while pending requests produce unidirectional links. However, the number of explicitly declared buddy relationships across the forums is quite small. Less than 10% of all users in each forum issued a friend request to another user, suggesting that underground forum users do not think of their relationships as persistent, or that the members see no utility in friending other users.

Because buddy links do not fully capture the latent relationships present in the forums, we further analyze the social network by including links that result from private messaging. If user  $u_1$  sends a PM to user  $u_2$ , we establish a directed link from  $u_1$  to  $u_2$ .

Lastly, thread relationships result when two users post in the same sub-forum thread. To establish these links, we order all posts in the same thread by their post times. We then create a link from user  $u_2$  to user  $u_1$  if  $u_2$  posted after  $u_1$ , with the reasoning being that  $u_2$  is interacting with all users in the thread prior to his or her post. Unsurprisingly, Table 2 shows that these one-to-many thread relationships produce the most links and fewest weakly connected components. One may use more advanced techniques (e.g., parsing “[QUOTE]” and “@<username>” expressions) to establish finer-grained thread relationships, but we leave this to future work.



(a) Private Message



(b) Thread

Figure 1: Degree distribution for reciprocated links.




Figure 2: The median social degree for users based on private messages in a 60-day window surrounding their registration dates.

Because so few buddy links exist, we only consider the private message (“PM”) and thread (“thread”) link types in our analyses.

#### 4.1 Degree Analysis

Figure 1 shows the degree distributions for the six forums. We only consider pairs of nodes with reciprocal links. Figure 1(a) shows that few users exchange private messages on HL, BH, FH and HS, with less than 40% of users involved in more than one private message exchange. Members on CC and LC create more relationships through private messages, with over 50% of their users linked to at least one other forum member, and more than 25% of their users involved in at least 10 conversations. This difference is likely due to users being more actively engaged in commerce, with many transactions taking place over private messages. The users in all forums (except for HL, being relatively new) create a comparable number of links via posting, which makes sense since this is the primary means of communication in forums. CC and LC exhibit more links, with over 50% of their users linked to over 20 other members via public postings (double the amount of other forums).

#### 4.2 Social Network Growth

Figure 2 shows how the median private messaging social degree changes among forum users in the 60 days following their registrations. We only consider reciprocal links when computing the social degree. For each pair of linked users, we looked at the difference between their registration dates and the latest timestamp associated with their first private message exchange. We omit HL and BH due to the noise present in their curves. For the private messaging social graph, there is little change in the median social degree after the first week for users on FH and HS. In contrast, for CC and LC, both sets of users experience growth after the first week, again likely due to the business transactions occurring over PM to a dynamic set of

		German ( <i>DE</i> )			English ( <i>EN</i> )	
		LC	FH	CC	BH	HL
<i>DE</i>	LC	–	13.99	3.16	0.05	0.03
	FH	6.87	–	2.23	0.02	0.08
	CC	7.09	10.19	–	0.02	0.39
<i>EN</i>	BH	0.11	0.10	0.02	–	0.00
	HL	0.25	1.19	1.36	0.00	–

Table 3: Percentage of overlapping email addresses across all pairs of forums. Percentages listed in each row are with respect to the population size of the forum specified in that row. The HS dataset lacked addresses.

		German ( <i>DE</i> )				English ( <i>EN</i> )	
		LC	HS	FH	CC	BH	HL
<i>DE</i>	LC	–	17.34	21.47	6.18	1.57	0.62
	HS	9.61	–	13.13	3.03	1.24	0.44
	FH	10.54	11.63	–	4.43	1.07	0.46
	CC	13.80	12.21	20.19	–	1.86	1.03
<i>EN</i>	BH	3.38	4.82	4.73	1.80	–	0.33
	HL	4.77	6.17	7.20	3.58	1.19	–

Table 4: Percentage of overlapping usernames across all pairs of forums.

partners. After the first month, the median social degree for the users on LC and CC is approximately nine, versus two for FH and HS. The growth rate for the thread social graph (not shown) follows the same trend across all forums: users undergo the largest growth in the first several weeks. In contrast, however, all forums show members continuing to interact with new users via public postings.

#### 4.3 User Overlap

To study the population overlap among forum members, Tables 3 and 4 show the percentage of overlapping email addresses and usernames (respectively) between each pair of forums. Of the six, only HS did not list any email addresses for its registered members. LC and FH share over two thousand email addresses, while LC and CC share 595, roughly 7% of the CC membership. This overlap is unsurprising, since both forums are geared toward the trading of stolen credit cards. However, even the newest English-based forum, HL, shared 6 to 33 email addresses across the German forums. Though email addresses are likely a more reliable metric for establishing population overlap, usernames also provide insight into overlap since users desire unique public identities to maintain their reputations. All forum pairs contained some number of overlapping usernames, with the German forums (CC, FH, HS, LC) sharing an appreciable number of usernames. For example, over 10% of the usernames on CC are present in all three of the other German forums. Likewise, over 17% of the usernames on LC appear in both FH and HS.




Figure 3: Median activity users engaged in prior to transitioning groups for CC.




Figure 4: Median activity users engaged in prior to transitioning groups for FH.




Figure 5: Distribution of users' interactions for PMs on LC.

Category	Threads		Users		Top Subcategory
	B	S	B	S	
payments	5,294	5,074	1,354	1,281	paysafecard
game-related	935	951	449	459	steam
credit cards	597	798	339	421	unspecified cc
accounts	761	566	382	356	ebay
merchandise	390	518	246	334	iphone
software/keys	355	485	214	296	key/serial
services	155	562	119	384	carder
victim logs	380	334	237	232	viclog
mail/drop srvs	347	292	248	203	packstation
fraud tools	203	343	132	239	socks

Table 5: Top 10 most commonly traded merchandise categories on CC

Category	Threads		Users		Top Subcategory
	B	S	B	S	
payments	8,507	8,092	1,539	1,409	paysafecard
game-related	2,379	2,584	924	987	steam
accounts	2,119	2,067	850	974	rapidshare
credit cards	996	1160	467	566	unspecified cc
software/keys	729	1410	422	740	key/serial
fraud tools	652	1155	363	601	socks
tutorials/guides	950	537	562	393	tutorials
mail/drop srvs	751	681	407	364	packstation
merchandise	493	721	264	404	ipod
services	266	916	176	555	carder

Table 6: Top 10 most commonly traded merchandise categories on LC.

## 4.4 Group Elevation

Users that join a forum are assigned a group, which roughly corresponds to their social status on the site. Generally, users start in the pending authorization group, meaning they must perform some action (e.g., respond to email confirmation) or undergo some type of scrutiny before being given access to the forum. Once the user has jumped through the necessary hoops, they begin in the “newbie” group. After some activity, users are generally elevated to a non-newbie group and advance from there. Figures 3 and 4 show the median amount of activity that users engaged in prior to transitioning to higher group levels for CC and FH (BH and HL were similar to FH). All the forums place a large emphasis on public postings versus private messaging, indicating that reputation comes from being publicly active on the forums. Users with greater standing in the CC forum have the most balanced amount of activity, posting and private messaging in roughly equal amounts.

## 4.5 User Interaction Analysis

Figure 5 shows how private message interactions are distributed among users’ “associates” (i.e., fellow members they are linked with) on LC, which has the greatest number of PMs. We looked at these distributions to determine the extent to which users interact with different individuals. For each user, we compute a distribution of private messaging events over the user’s associates. We then looked at the 60%, 70%, and 80% points in that distribution. Figure 5 suggests that users on LC exchange private messages with a diverse set of individuals, versus users on traditional OSNs, who interact with few of their friends. Wilson *et al.* [11] found that, for users on Facebook, 20% of their friends account for 70% of their interactions. In contrast, for users on LC, approximately 70% of their associates are responsible for 70% of their private messages. The corresponding graph for users linked via threads is similar.

## 5. MARKETPLACE

In this section we look at the types of goods and services exchanged on LC and CC, the two forums with the most well devel-

oped and active trading marketplaces. We first look at what types of goods are traded among these two underground communities, and then analyze how social degree and reputation affect trading.

### 5.1 Merchandise

To determine what types of items are available on the forums, we extracted thread titles containing the markers “[B]” or “[S]”, denoting items that are being traded for and sought after, respectively. We then wrote over 500 regular expressions to bin the items into 18 categories; these hand-defined categories include merchandise, banking information, drugs, mailing and dropping services, and a number of other commonly observed wares/services. We created the categories based on domain knowledge of illicit goods and by randomly sampling trading thread titles. Using our regular expressions, we categorized 87% of the 14,430 CC threads and 77% of 31,923 LC threads. Because users typically list several items for trade in a single thread, a thread may be counted in multiple categories. There is a long tail of merchandise types that we did not cover with our regular expressions; for example, on LC, threads mention such items as “Internet hack N95” or “Proteine - Inko XTREME Muscle Gainer”, while on CC, threads offer up such goods as “Conrad.de Kundenlogins” or “Pall Mall umsonst”.

Tables 5 and 6 show the top 10 most commonly traded items on CC and LC (respectively), ordered by the number of total binned threads in the designated category. The thread column shows the number of thread titles containing terms associated with the category, while the user column shows the number of distinct users who created those threads. The “B” and “S” columns denote threads where items were being traded for or sought after, respectively.

The items most commonly traded for are offline/online payments, including PayPal, cash, Ukash, and PaySafeCards (PSC). Over 5% of all threads involve trading for offline/online payments on both forums. Traders in the underground market prefer PSC, a type of prepaid online currency that is widely used in Europe. Gaming accounts, in particular Steam, are the second most commonly traded item; credit cards and accounts make up the next two traded for




Figure 6: Top percentile of traders vs percentage of goods traded.




Figure 7: Number of LC trading threads created vs. median social degree.




Figure 8: PM social graph degree vs median number of trading thread response PMs for LC.

items. While not shown in the table, drugs also made up a notable number of traded items on CC: over 100 threads listed weed, 25 mentioned Viagra, and 15 specified speed.

Next, we look at the number of trading threads the top merchants are responsible for. Figure 6 shows the relationship between the top percentile of traders on both forums and the percentage of goods traded. For example, the top 10% of the traders on both sites (measured by the number of trading threads created by the user) account for 40–50% of the goods traded. One implication of this trading distribution is that law enforcement can pursue the top tier traders to shut down much of the marketplace activity. The top traders can create multiple accounts to mask their activity levels, but accruing reputation for possibly numerous accounts is nontrivial.

## 5.2 How Social Degree Affects Trading

Now, we analyze the effect of the social network on trading. Figure 7 shows how the median social degree measured in private message links increases for users after posting trading threads on LC. We order each user’s trading threads by time. Subsequently, we compute the user’s PM social degree just before he or she creates each thread. We then bin that social degree with respect to the thread’s posting order, taking the median over all users. The results suggest that, as users trade more, they continue to interact with increasingly larger numbers of individuals and hence, potential customers. For LC, the median out and in degrees for the PM social graphs increases at a roughly constant rate of four for every trading thread created. The CC graph looks similar; the PM out/in degree increases at a rate between one and two for each thread.

We now investigate how PM in/out degree affects the response to a trading thread. Figure 8 plots the in/out PM degrees for LC users before they post a trading thread against the number of PMs they receive in the following week. In comparison, LC and CC members receive on average 0.07 (standard deviation  $\sigma = 0.68$ ) and 0.11 ( $\sigma = 1.19$ ) private messages per week when issuing no posts prior to that week. The graph suggests that traders with higher PM social graph degrees receive more PMs in the week after they post a trading thread. For CC, the numbers look similar, but become increasingly noisier after an in/out degree of 30.

Lastly, we look at what fraction of private messages are sent by new users to traders in the week following a trading thread post. Figure 9 shows that for approximately 30% of all trading threads, the posters receive only PMs from individuals they have interacted with before. For 50% of the trading threads on LC and CC, 60–75% of the PMs came from prior acquaintances. Finally, for approximately 20–25% of all trading threads, the trader interacted with only new people.

## 5.3 Effect of Group Status

Figure 10 shows how a user’s group affects the responsiveness to a trading thread; we focus on CC because LC does not contain any data regarding user group transitions. We looked at the trading

threads posted by users at a certain group level. We then determined how many PMs the users received in the week after posting trading threads. Figure 10 suggests that a user’s group status does influence how many PMs the user receives upon posting a trading thread. Verified vendors and VIP members receive between 2–3 times more response PMs than “newbies”.

## 5.4 Effect of Ratings

The only forum with an explicit rating system for trading transactions is CC. The rating system is ternary: traders receive either a positive, negative, or neutral feedback. Surprisingly, most of the reviews are positive: of the 3,157 reviews (20% of discernible trading threads), only 67 were negative and 2 were neutral. Of the 67 negatively rated users, 43 were banned, and the rest were generally rated negatively due to “poor” or “unfriendly” service. We suspect that bad traders are outed publicly in separate thread posts (see Section 6.2) and banned before being rated. Because so few users were rated poorly, we do not differentiate between the rating types.

We now consider the effect of the rating system on the amount of interest a trader receives (either in the form of response posts or private messages) after posting a trading thread. Figure 11 shows the effects of the first 10 ratings on the median number of PMs users receive in the week following a trading thread post. With zero ratings, the median response PM count is one, but with a single rating, the median count rises to 15. The median PM count continues to increase from there, though not in a well-defined manner. While rated traders receive more private messages, they do not experience an increase in the number of response posts; the median response post count remains constant at one. We speculate that the rating system lends more credibility to a trader’s threads, and people are not so quick to question the trader’s reputation in the public space.

## 5.5 Activity To First Rating

We now analyze how much activity users must participate in before they earn enough trust such that they engage in a business transaction with another forum member. We measure the number of actions users take (postings, private messaging) before they receive their first ratings. Relying on the rating system is subject to error, since users can conduct business outside of the forum, but it provides us with some idea about how trust is earned. Figure 12 shows that, before approximately 50% of users received their first ratings, they posted around 60 times in 50 different threads, received about 35 private messages from 13 users, and sent around 33 private messages to 13 users.

## 6. BANNED ANALYSIS

Some users in underground forums behave maliciously towards other members. In this section, we look at the top reasons why users are banned from the forums. We also investigate different properties associated with accusations of fraud in the marketplace.




Figure 9: Fraction of PMs received from new users after posting a trading thread.




Figure 10: User group vs median number of trading thread response PMs for CC.




Figure 11: Effect of ratings on response to trading threads.




Figure 12: CDF showing the activities a user participated in prior to receiving a first rating.

Stat Name	Banned		Unbanned	
	Accuser	Accused	Accuser	Accused
Num PMs Sent	243.0	34.5	22.0	73.5
Num PMs Recv	271.0	31.0	24.5	57.5
PM InDegree	101.0	12.0	10.0	21.5
PM OutDegree	104.0	12.0	7.0	22.0
Num Posts	299.5	50.0	72.5	134.0
Thread InDegree	500.5	182.0	180.0	333.0
Thread OutDegree	527.0	183.5	198.0	332.0
Total Pairs	314		62	

Table 8: Statistics for ripper accusers and accusees. The numbers represent the medians across all accuser/accused pairs.

comprise over 10% of the bans. Rippers are individuals who rip off other members, and threads are created to identify these users.

Forum	Rank 1		Rank 2		Rank 3	
	Reason	%	Reason	%	Reason	%
BW (22)	spammer	40.9	dup. acc	22.7	infl. posts	22.7
CC (1,587)	dup. acc	60.7	ripper	12.1	spammer	7.3
FH (447)	dup. acc	30.6	malware	16.6	spammer	14.1
HS (317)	spammer	24.6	malware	10.7	dup. acc	9.8
HL (75)	infl. posts	52.0	trade-rel.	34.7	spammer	13.3
LC (247)	misuse	37.7	spammer	17.4	ripper	17.0

Table 7: The top three reasons why users are banned on each of the forums.

## 6.1 Why Users are Banned

The fraction of users that are banned on the forums is non-trivial. For example, the moderators on CC have banned over 20% of the users that appear in their members table. Because the individuals who participate in underground forums generally lack some scruples, this is not surprising. Table 7 shows the top three reasons why users are banned in the forums. To generate the data, we analyzed a specific table in five of the forums (BW, CC, FH, HS, HL) that explicitly holds information regarding user bans. LC did not have such a table, so we analyzed the warnings issued to banned users and assumed that the final warning received by the user resulted in the ban. Next, we created 13 categories for the most commonly appearing banned reasons and wrote 35 regular expressions to bin them. Again, users can be banned for multiple reasons. Also, not every ban is explained (the reason is sometimes left blank), so the percentages shown are with respect to the explained bans (represented by the numbers in parenthesis). These regular expressions covered over 70% of the banned reasons across all five forums.

The most common problem across the forums involves duplicate accounts (dup. acc), which appears in the top three reasons in four of the forums. Users often create duplicate accounts to circumvent a prior ban. Another problem in the forums is, ironically, spamming and malware attacks, with BW and HS particularly inundated with spammers. In the commerce oriented forums (LC and CC), rippers

## 6.2 Accusers vs. Accused

We next investigate the repercussions of accusing one member of being a ripper on the CC forum. To do this, we extracted all the threads where a user accused another member of being a ripper. The titles of these threads often take the form “Ripper <username>”. Once we identified both parties, we compared the amount of activity the users engaged in prior to the accusation time. Table 8 shows the median values for several statistics about the accusers and the accused. We see that, in the cases where the accused person was ultimately banned, the accusers were much more active on the CC forum. For example, accusers had more than eight times the number of links in their PM graphs than the banned accused users, and the accusers had roughly twice as many links in their thread graphs. In the cases where the accusation did not result in a ban, the accused exhibited more activity than the accusers. Also, the unbanned accused users had a larger number of links in their PM (e.g., 10 vs. 21 for PM indegree) and thread (e.g., 180 vs. 330 for thread indegree) graphs than their banned counterparts.

## 7. CONCLUSION

This paper has characterized the social network makeup for six underground forums, how users interact, how baseline reputation is established and how it changes over time. This work is a first step in a larger research agenda to understand the social dynamics of the underground and how they impact e-crime market efficiencies.

## Acknowledgments

We thank the reviewers for their valuable suggestions, and Vyas Sekar for his assistance with the accompanied public reviews and comments. This work was supported by NSF grants NSF-0433668 and NSF-0831138, ONR MURI grant N000140911081, and by generous support from Google, Microsoft, Yahoo, Cisco, and the UCSD Center for Networked Systems (CNS). McCoy was supported by a CCC-CRA-NSF Computing Innovation Fellowship.

## 8. REFERENCES

- [1] Ivision Power Board. <http://www.invisionpower.com>.
- [2] vBulletin. <http://www.vbulletin.com>.
- [3] H. Fallmann, G. Wondracek, and C. Platzer. Covertly Probing Underground Economy Marketplaces. In *Proceedings of DIMVA*, July 2010.
- [4] J. Franklin, V. Paxson, A. Perrig, and S. Savage. An Inquiry into the Nature and Causes of the Wealth of Internet Miscreants. In *Proceedings of CCS*, October 2007.
- [5] T. J. Holt and E. Lampke. Exploring Stolen Data Markets Online: Products and Market Forces. *Criminal Justice Studies*, 23(1):33–50, 2010.
- [6] A. Mislove, M. Marcon, K. P. Gummadi, P. Druschel, and B. Bhattacharjee. Measurement and Analysis of Online Social Networks. In *Proceedings of IMC*, October 2007.
- [7] J. Radianti. A Study of a Social Behavior inside the Online Black Markets. In *Proceedings of SECURWARE*, July 2010.
- [8] A. Shilman. The Underground Credentials Market. *Computer Fraud and Security*, 2010(3):5–8, 2010.
- [9] B. Stone-Gross, T. Holz, G. Stringhini, and G. Vigna. The Underground Economy of Spam: A Botmaster’s Perspective of Coordinating Large-Scale Spam Campaigns. In *Proceedings of LEET*, 2011.
- [10] R. Thomas and J. Martin. The Underground Economy: Priceless. *login:*, 31(6):7–16, Dec. 2006.
- [11] C. Wilson, B. Boe, A. Sala, K. P. N. Puttaswamy, and B. Y. Zhao. User Interactions in Social Networks and their Implications. In *Proceedings of EuroSys*, April 2009.
- [12] J. Zhuge, T. Holz, C. Song, J. Guo, X. Han, and W. Zou. Studying Malicious Websites and the Underground Economy on the Chinese Web. In *Proceedings of WEIS*, June 2008.

# Summary Review Documentation for “An Analysis of Underground Forums”

Authors: M. Motoyama, D. McCoy, K. Levchenko, S. Savage, G. Voelker

## Reviewer #1

**Strengths:** This is a new, interesting direction for research. The dataset is unique and this paper will generate a lot of excitement/discussion.

**Weaknesses:** The connection of the actual analysis presented to the motivation is slightly weak (e.g., its different from conventional OSNs, reputations, dynamics of graph).

**Comments to Authors:** I would have liked to see a better analysis of the following: (1) how are these graphs structurally different from regular OSNs? (2) is the temporal dynamics of these graphs different from OSNs -- e.g., are there frequent interactions between a pair of entities or are people constantly finding new friends etc?

You say the dataset is “public” -- Is it public to your group or more broadly??

Why are German forums surprisingly a high fraction of your dataset-- are underground groups more active there?

Section 4: I like how you break down relationships based on buddy/PM/thread. I wonder if the thread category can be more finely divided; e.g., don't the forum posts have some sort of “Reply-To” semantics where user 1 responds to user 2's comments etc?

Section 5.1: you claim that most trading can be shut down -- wouldnt the obvious response then be to create a lot of sybil-like proxies or pseudonyms?

Section 5.4 -- the majority of positive ratings makes me worry about collusion/sybil effects here. Is it possible for you to rule these out?

Section 5.5 -- what fraction of transactions have ratings? Isnt the number of posts to first transaction or the time to first PM/thread response to trading post a more accurate measure of what you are trying to show here?

## Reviewer #2

**Strengths:** This is the first study of the kind i have encountered. Moreover, the 6 datasets have sufficient differences to make the study interesting.

**Weaknesses:** None of the results actually surprised me or stayed in my mind after I finished reviewing the paper (this is typically a bad sign in my opinion). Second, I do not really see any strong methodological contribution that could be the lesson learnt here. More importantly, the metrics studied as rarely precisely defined or discussed - the paper appears as a series of plots of different metrics.

**Comments to Authors:** I wish I were more enthusiastic about this work, but I am really not. I think the authors deserve credit for having looked for those data sets and for having analyzed them under the same lens. However, I do not really see anything that exciting coming out of the study itself. Moreover, I do not really see a strong methodological contribution either.

In particular, for a measurement paper, the authors have done a poor job describing the studied metrics. As an example, section 5.2 discusses “the average social degree for a user after posting trading threads”. I would like to see the precise definition of this metric.

Second, and given that this is a measurement paper as well, all metrics that present average values should also present confidence intervals. I think extracting conclusions based on the average values alone can be very misleading.

In summary, I do not really feel like I learnt something new from this paper. I would love if the authors had managed to frame the discussion in a more impactful way...

More detailed comments: you talk about using regular expressions to identify specific categories. Those matches tend to lead to non 100% matches. It would be nice to know what the remaining threads contained that did not match the regular expression. Have you looked through them?

## Reviewer #3

**Strengths:** The first strength is the novelty of the paper. The second strength is its impact to the community. It brings knowledge about background economy to the network research community, hopefully will trigger more research in this direction. The background section (although brief -- but this is a short paper) provides pointers to background reading in the area, that I believe majority of network researchers still do not have a comprehensive understanding.

**Weaknesses:** I wish the paper could go one step deeper to discuss how the results from this paper can be applied to mitigate the underground problems.

**Comments to Authors:** None.

## Reviewer #4

**Strengths:** Interesting study of underground forums based on complete datasets. Some interesting observations regarding banned accounts and the effect of rating.

**Weaknesses:** The authors claim that underground forums are quite different than regular ones, but never attempt to quantify or examine this difference. The classification of topics is based on a manual process.



**Comments to Authors:** This is a well-written paper examining underground forums. Overall, I enjoyed reading the paper, although the analysis could go in more depth describing in more detail some of the findings observed.

What I felt is definitely missing from the paper, is a comparison with regular forums. It is unclear to me whether characteristics are significantly different as the authors argue.

In particular, I would argue that, intuitively, observations in section 4 do hold for most of the online forums at least qualitatively. Quantifying the differences would be interesting here.

In section 5, I found the analysis a bit incomplete. For example, the authors describe the number of PMs after thread response versus the user degree (figure 6). However, it is hard to get any context from this figure without a similar one that is not conditioned on thread response. What is the PMs/week without thread participation? Also the figure presents averages that might be misleading or provide an incomplete picture. Some percentiles would be helpful. This is true for other figures in the same section.

I did not understand at all Section 5.3. Figure 7 shows that there is some correlation between user status and PMs received but the text claims otherwise. Am I missing something here?

I found the analysis in sections 5.5 and 6 interesting and perhaps the authors could devote more space to these observations. Section 5.5 shows that one needs significant activity before a rating is received. Again comparison with regular forums would be interesting here.

Section 6 implies that banning decisions do depend on the activity and social characteristics of users. Providing similar numbers as in Table 5 conditioning on the reason for ban would perhaps be interesting here.

## Reviewer #5

**Strengths:** Interesting topic, good data sets, reasonably good analysis.

**Weaknesses:** Analysis has limited depth and merely presented basic points, implications of findings are not discussed.

**Comments to Authors:** I found the topic of this paper interesting and overall like the paper. I found it rather ironic that these forums are called underground and then such a detailed data sets of their transaction is available even through wikileaks. The obtained data sets appear to be database (or log) on the original server. It is hard to get such information (e.g. level of seniority of users in a forum, number of private messages, content of messages, ...) even for forums such as twitter and fb that are not underground. I found it intriguing that such a data is available for several underground forums. It is helpful if authors spell out what information are available in these data sets in section 3. For example, availability of users' level of seniority and destination of private messages (for later analysis in the paper) were surprise to me.

For some of the analysis, authors come up with set of expressions to classify messages. Is there a methodology or known list for such an expression or authors have other ways to define these expressions?

It is useful to have some explanations on how the presented findings can be used to disrupt these venues or defend against their attacks (or other purposes), and some minor explanation on how the observed behavior is different from other forums that are not underground. Of course such a discussion should be limited given the length of this paper.

It is useful if authors provide a link to their data sets on wikileaks so the rest of the research community can examine them.

## Response from the Authors

The reviewers inquired about many interesting ways in which the analyses could be extended. We share these reactions. In the spirit (and constraints) of short papers, we view this paper as exploratory work seeding the analysis of the social networking properties of underground forums. We plan to expand on this topic in future, more comprehensive work.

In preparing the camera-ready version, we have further emphasized the differences between the commerce-oriented forums (Carders and L33tCrew) and those focused primarily on the exchange of information. We placed all the reciprocal degree distributions for the underground forums on the same graphs to emphasize the structural differences between the forum types. We point out how the trading forums are much more interconnected among the different link types. Also, we have omitted statements about the differences between underground and regular forums, since we do not make direct comparisons between the two types in this work.

Several reviewers inquired about making the data public; unfortunately, the data is not ours to share. As mentioned in the paper, however, others have leaked the data.

One reviewer asked how we might disrupt the forums. We discuss how 10% of the traders are responsible for 40-50% of the goods traded. Targeting the forum members with the most activity could disrupt, or at least impede, the marketplace. Another reviewer pointed out that traders can mask their activity levels by creating multiple accounts. However, reputation is hard to accrue, and we do not believe creating multiple accounts is a viable option for high volume traders.

We made clarifications where possible, especially regarding the metrics we define. We included a brief description of the information available in the back-end databases, along with pointers about the various packages powering each forum.

Another reviewer commented on the lack of negative ratings. Many users with negative ratings are banned, and the rest are generally accused of providing "poor" service. Also, we discuss in Section 6 how ripper accusations are levied against forum members outside the rating system.

Some reviewers commented on our methodology for classifying the merchandise being traded for. We expanded on the methodology in the paper. Briefly, we developed our list of regular expressions using domain knowledge, and by randomly sampling threads to capture the most frequently traded items. Also, we looked at threads we did not classify under our regular expressions; we found a long tail of traded items.