The Dynamics of Web-Based Community Safety Groups: Lessons Learned from the Nation of Neighbors

Awalin Sopan[1,2], PJ Rey[3], Jae-wook Ahn[1,2], Catherine Plaisant[2], Ben Shneiderman[1,2]

[1]Dept of Computer Science,
[2]Human-Computer Interaction Lab,
[3]Dept of Sociology
University of Maryland

ABSTRACT
Community safety as a social issue has expanded its reach to web forums, portals and dedicated sites. This paper presents our study of 230 community safety groups whose members communicate through the Nation of Neighbors website. We analyze the patterns of activities within these communities along with their temporal dynamics. We demonstrate both feature-based and temporal analyses of the communities aiming at discovering the characteristics that make such communities successful. We use ManyNets's capability to visualize the overview of multiple networks at once, demonstrating the value of visual analytics for community managers to better understand their communities. Using previously-developed health metrics we distinguish the successful communities, observe the influence of leaders in those communities and establish that larger communities are reporting more crime incidents rather than having discussion on other topics. To our surprise, we did not observe any strong association between the involvement of Law Enforcement personnel and activeness of the communities.

Author Keywords: Technology-Mediated Social Participation, social media, community safety, network analysis, visualization, temporal analysis

ACM Classification Keywords: H.5.m Miscellaneous

INTRODUCTION
General purpose social networking websites like Facebook and micro-blogging sites like Twitter have hundreds of millions of active users. Online social networks targeted towards national priorities, such as disaster planning, crime-watch or food-safety, are also growing gradually. A central research challenge is to understand the determinants of successful growth. The use of a visual analytics tool would guide the community managers to understand the dynamics of these communities. To study this we used data from Nation of Neighbors (NON) [14], which is a platform for online neighborhood crime watch communities. While these communities are geographically distributed, they share the same concern: awareness about the safety of their neighborhood. NON expands on the successful Watch Jefferson County project (Jefferson County, WV, http://watchjeffersoncounty.net), by creating a national portal for residents and law enforcement officers to interact over the shared goals of preventing crime and strengthening communities. The NON website facilitates real-time Neighborhood Watch via citizen reporting and fosters social participation within communities. The members of the communities and Law Enforcement jurisdiction can add their community to NON, report crime or other incidents in their communities, participate in community discussion, share news, photos or documents and manage upcoming events. The system sends real-time email or text message alerts to members.

During this case study we worked closely with the executive director and manager of NON, defined metrics to quantify the success and growth of the communities and refined our visual analytics tool ManyNets [10] to explore and compare communities as well as analyze their growth over time.

ManyNets presents network data in three kinds of tables: network tables, node tables and edge tables, connected to node-link diagrams of the networks. While network visualization tools are primarily focused on visualizing the network structure at a single point in time, ManyNets now allows users to analyze community growth over time. All activities of community members are time stamped and suitable for temporal analysis.

Our team of computer scientists and sociologist worked together to propose new metrics for community analysis and to improve ManyNets' ability to support three types of analysis:

Label	Node Count	Relationship activity_type	Accept	Invites	Post	Reply	Report	TotActivity	InterIntensity	Relationship activity_date	ActiveMonth
Watch-Jefferson-County:4	70		16	68	52	35	207	378	54.00		7
shannondale:17	51		24	67	11	13	59	174	24.86		7
Blue-Ridge-Acres:3	13		7	63	19	12	50	151	25.17		6
charles-town:18	16		7	66	3	5	6	87	12.43		7
river-view-park:32	5		3	40	5	1	9	58	7.25		8
ranson:23	12		5	42	1	0	9	57	8.14		7
breckenridge-north:25	2		1	39	2	0	0	42	5.25		8
summit-point:37	4		1	1	1	0	3	6	1.20		5
avon-bend:38	2		0	0	5	0	1	6	1.00		6
shenandoah-junction:41	2		2	0	1	0	0	3	0.38		8
potomac-hills-subdivision:40	1		1	0	1	0	0	2	0.25		8

Figure 1: ManyNets network table showing 11 communities (one per row) and a selection of the available metrics.

1) **Community level analysis**: We defined and implemented novel metrics to assess community success, compared communities along those metrics and develop hypotheses about factors influencing community growth and member participation.

2) **Member level analysis:** We analyzed the activity of individual community members, defined and implemented metrics to identify leaders and to quantify their impact in the community.

3) **Temporal analysis:** We compared the growth and activity patterns of communities over time.

In this paper we first describe the new metrics and tool features developed for the three types of analysis, then we propose a workflow for community managers to analyze their data and better understand how their community evolves over time. Finally we give examples of insights generated from the analysis of the NON data.

RELATED WORK

By strengthening social control of information and restoring a 'sense of neighborhood,' Neighborhood Watch has been proven effective at lowering crime [5, 16, 30]. Few studies found a significant reduction of crime after crime watch programs were introduced in American communities [4, 21, 31, 32] even though in some cases the positive effects dissipated rapidly [7, 21]. The limited success of traditional neighborhood watch programs indicates that new approaches should be explored in pursuit of improved outcomes. Some such efforts are already under way: notably, Hollaback, SeeClickFix, Crimereports.com and SpotCrime. Nation of Neighbors, however, has the unique and specific goal of reorganizing neighborhood watch with new web-based tools to make it more effective. To function optimally, neighborhood watch programs require involvement, partnership between law enforcement and community members, a common understanding of the problems to be addressed, motivation and organization, and continuation of effort [3]. Having online communities targeted towards neighborhood crime watch gave us the opportunity to analyze this domain more closely and determine how this online social participation can help reduce crime and increase awareness. Many studies (e.g. [6, 28]) describe relevant measures that needed to be analyzed for online community success, but success is defined by the unique mission of the community and its organizers. Moreover, different communities engage in different forms of activity [26]. Close collaboration with the community managers can help analysts identifying these measures, coming up with hypotheses and verifying those using visual analytics.

Several projects adopted visualization technologies to support the efforts to reduce crime. Lodha [22] introduced a technique to visualize crime on geo/temporal GIS grids and the idea to visualize the crime statistics on the web-based maps became widely adopted (oakland.crimespotting.org). Social network analysis and network visualizations are actively used for analyzing the networks of criminals or terrorists [24,37, 38]. In particular, we expect these technologies to serve community managers by answering questions such as what forms of invitation are most successful [36], how do new members connected to the existing members [18, 20],what generates high quality reports from the community [35], or how do messages spreading across networks [1, 17, 19]. Vizster provided a visualization method that can analyze online social networks [15]. Hansen and Shneiderman [13] used NodeXL to mine conversation networks.

Dynamic network visualization techniques can help understand how these communities are changing over time. Trier et al. [33] demonstrate the usefulness of using dynamic visualization tools to understand community development. Two common approaches to visualize dynamic networks are 1) plotting summary statistics over time [8] and 2) presenting a separate node-link diagram of the network at each point of time ([27]). Durant et al. [9] presented a snapshot-based network visualization showing different node positions over time. The "movie" approach is also used for dynamic visualization. Moody [23] distinguished 1) Flipbook style movies where node-positions were fixed but connectivity formation was captured and 2) Dynamic movies where node-positions changed over time. The use of sliding time frame to animate the network was introduced in Condor and TeCFlow [11, 12]. Animation approaches might distract or hinder users' attempts to track changes in the network, for example to track new nodes that are possibly responsible for significant change in the network, because nodes and edges keep changing their positions. TempoVis [2] keeps the node positions unchanged assigning colors to new coming nodes and TimeSpring [25] uses a new layout algorithm so the nodes in adjacent timeframes are not too far apart. In contrast, ManyNets uses tabular visualization to compare features of networks. It can display changes in a single network throughout time or it can compare the summary statistics of different networks for a given time period.

DATA, METRICS AND WORKFLOW

Data preparation

We collected the activity log of NON members from January 2005 to December 2011, 6370 activities in total. Activities were classified into 5 categories: report (describes an incident which occurred in the community), post (starting a discussion topic), reply (responding to a previously posted discussion), invitation (soliciting a person to join the community via email) and acceptance (new members joining the community following any email invitation). Two members have an edge connecting them (i.e. relationship) if one of them replies to other's post. Members often replied to a post by simply making a new post. In these cases, the posts were recoded as replies to the original post to accurately reflect conversations. The data also include a unique user ID and the joining date of each member as well as the county name and creation date of each community.

Once the data was recoded, it was imported into ManyNets. Our initial look at the 230 communities showed dramatic spikes in invitation activity in June 2006. Discussions with the Nation of Neighbors manager revealed that it corresponded to a major reorganization of NON and that the older data was not usable. The node (i.e. member) table also showed another anomaly: the member with ID 0 was the most active in most communities. Once again discussions with the community manager provided an explanation (i.e. ID 0 is used whenever a member posts anonymously) and we decided to ignore those anonymous contributions when calculating activity and leadership metrics for community members.

Community Level Analysis

ManyNets automatically creates a network table where each row is a community (Figure 1) and each column a metric. It computes default metrics such as node counts (i.e. number of members), edge counts (total count of activities), connected component count etc. In addition ManyNets allowed us to specify new metrics specific to the NON community, such as number of active months, total number of reports, etc. A distribution column shows the distribution of activity type using small color coded histograms. Here acceptance, invitation, post, replies and reports are red, blue, green, purple and orange respectively. In addition separate columns for each type of activity are provided as well.

ManyNets allows analysts to compare communities based on the metrics by easily sorting, filtering, clustering and selecting communities based on the metrics. Nevertheless it became apparent that more complex metrics were needed to represent the success of some communities [29] so we defined health metrics for each community, added them as new columns in the network table, and used a combination of metrics to identify successful communities for further analysis.

Community level health metrics

Our health metrics are as follows:

Equity: The variance in activity per month per member. Because members tend to make new posts rather than reply, analysts cannot use average number of two-way ties between members. Instead we look at the variance in activity per month per member.

$$\text{Equity} = s(A_C)^2$$

Interaction Intensity: This is the total activity divided by total member-months.

$$I = \text{Interaction Intensity} = \sum (A) / \sum (MM)$$

Average Active Months: This is the average number of months the community members have participated.

$$\overline{M} = \text{Average Active Months} = \sum (MM) / \sum (N)$$

Where,

A_C = Communication Activity = $\sum$ (reports + posts + replies)

A_I = Invitation Activity = $\sum$ (invites sent + invites accepted)

$\sum (A)$ = Total Activity = $\sum (A_{C\,+}\,A_I)$

$\sum (MM)$ = Total Member-Months = $\sum$ (months since each Member registered)

$\sum (N)$ = Number of Community Members

Selecting successful communities for further analysis

To focus our analysis on communities open to the public we first eliminated communities limited to law enforcement agencies (who only post reports), and the community of NON manager which focuses on management issues.

Using the sorting and filtering capability of ManyNets and the community health metrics we developed a process to filter out inactive or redundant communities:

1) Filter out the communities with no activities at all.

2) Keep only the communities that have at least 5 Invitation activities and at least 5 active members.

3) Sort the table on the Interaction Intensity column; select the communities with the highest Interaction Intensity and remove the others.

Finally we manually reviewed communities that geographically overlapped and kept the largest one. This entire process reduced the dataset to 44 active and independent communities suitable for comparative analysis:

Member-Level Analysis

Member level analysis gives us more insights about the individual members and their activity patterns. Community managers can identify influential members, their role in the community and their connection with other members. In ManyNets each community has a node table showing each member as a row. The columns are members' activity type distribution, degree, number of total activity, joining date and our proposed Leadership metric. Analysts can also select particular members and create their ego networks to visualize the connections of these members with other members.

Member-Level Leadership metric

Using ManyNets to sort communities by total activity and examine the node table of the most active communities, we noticed that the most active communities shared a common trait: each contained one or two members who were far more active than the other community members. This observation guided us to hypothesize about the importance of leadership in these communities and to examine whether these extremely active members influenced the behavior of other members.

The standard betweenness centrality only can be used as a leadership measure for a typical conversational network [34] but the Nation of Neighbors network has reporting and other unidirectional activities that do not resemble a typical conversation network. Absolute activity rates are also problematic to identify influential members, since leadership is relative to the activity of non-leaders. We, thus, decided to examine relative activity rates in each community, looking for outliers who participated an extraordinary amount (two standard deviations above the mean).

$$\text{Leadership} = M_A - (\text{standard deviation}(M_A))\times 2 - M_M$$

Where,

M_A= Total Activity of a Member = $\sum$ all activities by a member

M_M = Mean Activity of All Members

A positive leadership score indicates a member whose activity level is significantly higher than other members in the community.

Temporal Analysis

Temporal changes include growth patterns, changes in activity levels and changes in the type of activity over time. To analyze the data in the temporal dimension, we have added two new features in ManyNets:

Activity distribution over time

We introduced the column Activity date in the network table (Figure 1). Each cell in this column shows the distribution of activity count over time. By observing this column managers can identify different patterns of activity; e.g. a sudden spike in activities in a community, communities where the activity is diminishing, or persistent communities where activity level remains high. It is also useful to detect communities with anomalous activity patterns over time.

For example managers might distinguish new communities (that are less than a few month old) from mature communities. A mature community might be considered successful if its members remain active over time. A new community might be considered a success if the initial members keep recruiting more members, which a manager can see by looking at both activity date and the counts of invitations.

Temporal split of network

We implemented the "Temporal split" feature in ManyNets that splits a network into a series of sub-networks, each one comprising only the activities within a specific time range (a week, month or year). ManyNets opens a new tab showing a network table with one row by week, month or year. This table has all the functionalities of the network table. Once this table is sorted by time it can visualize the changes in activity type over time.

Analysis Workflow using ManyNets

For our case study, we followed a specific workflow that can guide the community managers to explore and analyze online communities. The workflow starts with community level analysis, using the network table which shows all the communities and their metrics. Sorting the communities by community level metrics, (e.g., interaction intensity or total activity) makes it possible to compare the communities. Users can select rows to generate a subset of communities which can be shown in another new tab for further analysis. Another option is to filter out parts of each community by removing unwanted nodes and edges using node/edge filtering; for example, they might want to compare all communities looking only at active members (i.e. removing nodes with no activity).

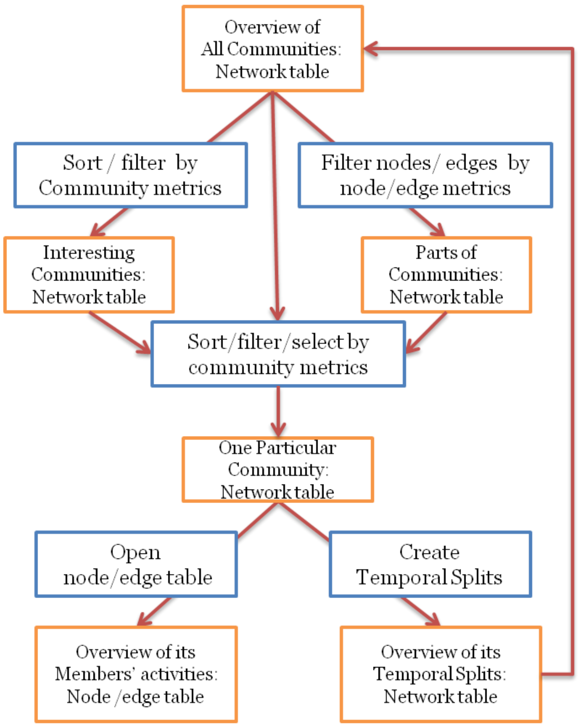

Figure 2: Workflow of analysis.

After selecting one particular community, the next step is to analyze its members and its temporal evolution. For member level analysis users can open the member table in another new tab to compare members and find members who are more active, who are good at reporting, who are good at inviting other members to the community, who are consistently active, who are from law enforcement, etc. If the activity distribution over time is interesting, the may want to split the data by week, month or year.

EXAMPLE OF INSIGHTS

Following the analysis workflow we produced insights about the activity patterns, growth patterns, and the leadership in the communities.

Activity patterns of all communities

To explore the activity patterns in communities, we looked at the "activity type" column. By simply scrolling we could see that some communities had a lot of reporting while others mostly had invitations. To study the prevalence of those 2 types of distributions we used the heatmap column overview feature of ManyNets [39]. This heatmap column overview rendered each community's activity type distribution as heatmaps stacked one after another. This heatmap view of distribution presents color as a function of histogram bar-height, from white to blue, blue being the color for the maximum value. The communities were clustered by similarity (Figure 3), clearly separating the communities with lots of reports at one end and the ones with lots of invitations at the other end.

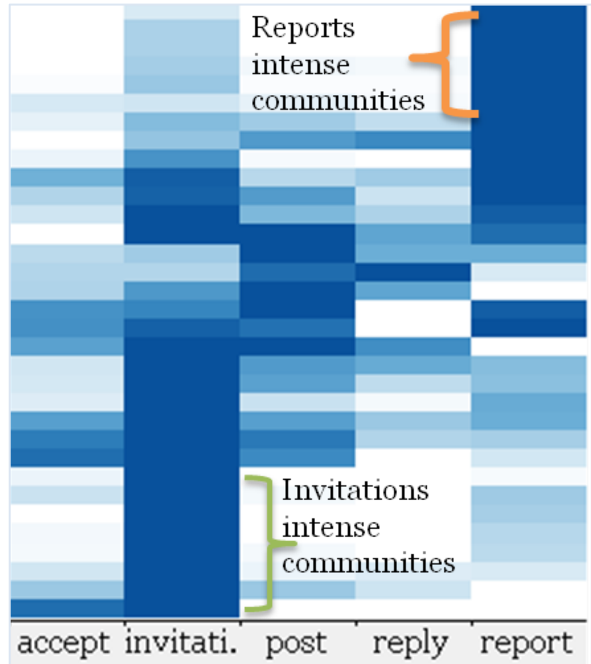

Figure 3: Heatmap overview of the activity type column for the 44 most active communities (one per row). Rows are grouped by similarity. Two clusters of Invitations intense and Reports intense communities are marked.

The histogram of the activity type for all the 230 communities showed that invitation was the most common activity (Figure 4, left), but after filtering down to the 44 larger and active communities (as described earlier), a different pattern emerged where we observed more reports than invitations (Figure 4, right).

To see if there was a correlation between the size of the communities and the activity patterns in the 44 communities, we generated a side by side overview of the activity type distribution column and the total member count column (Figure 5) sorted the rows according to the total member count. This showed that the larger communities have more reports than any other activity.

Growth Patterns of Communities

Using the "activity date" distribution column, analysts observed the rise and fall in activity over time. This column showed the total count of activity per day for each community, starting from July 2009 to December 2011. It showed different patterns in the activity in the communities. Some communities had persistent activities throughout the whole time. In contrast some communities started with high activity but gradually, their activity diminished.

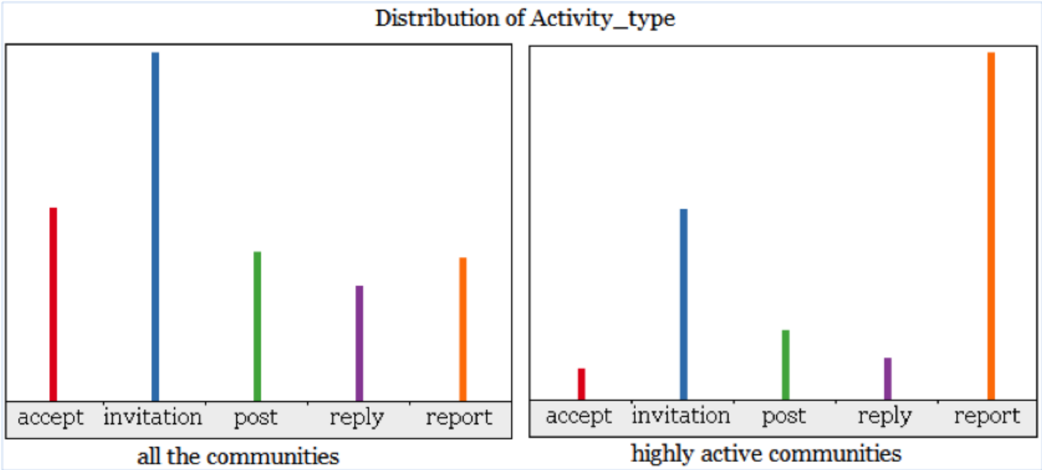

Figure 4: Relative distribution of activity type in communities. Left) all communities and right) active successful communities.

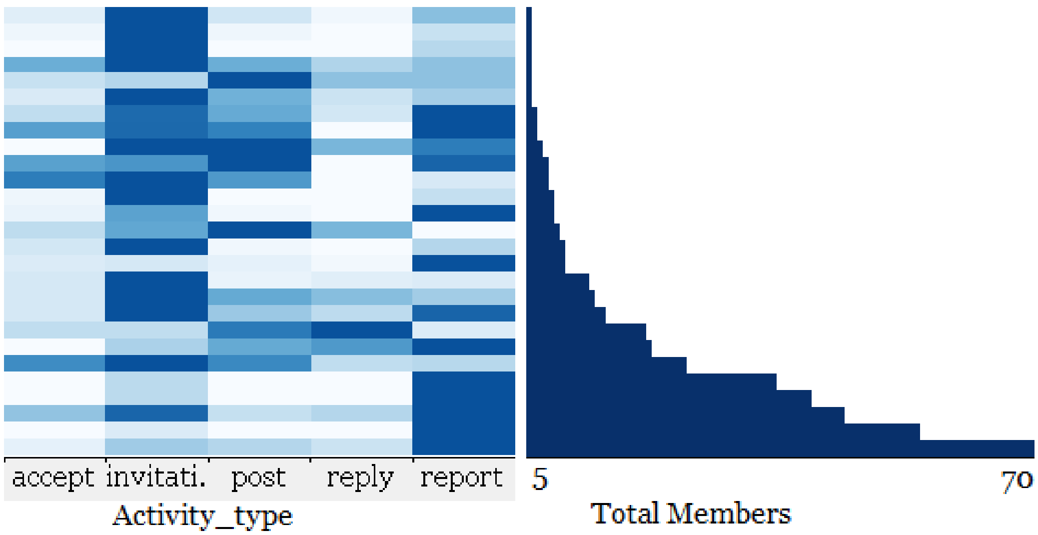

Figure 5: Side by side overview of activity type distribution (on the left) and Total members (on the right). Large communities have comparatively more reports than any other activities, smaller communities have more invitations seeking further growth.

We also identified some outliers using this distribution column. For example one community had a spike of activity and then there was no activity for a long duration and finally a small spike. To analyze what happened, we opened the edge table for this community that showed the date and type for each activity. There we saw that since the community started, member 1188 invited other users and posted in the community, while other users accepted an invitation from 1188, and it all happened in March and April 2010, then nothing happened until January 2011, when member 1275 made a report that happened to be the only report in this community.

After sorting the activity date distribution column in the community table according to their first activity, we noticed that all the top ones were from Jefferson County; they not only started earlier than most other communities, they also continued to be active. To see which communities became stale, we sorted the column by the last date of activity and identified which communities were not active anymore.

Leadership

Our member level analysis aimed at finding out leaders and the influence of law enforcement people. After selecting the most active communities we added a column showing the number of law enforcement people in each community and none of the successful communities had one, indicating the involvement of law enforcement people was not influential enough for the members. After that we filtered out the members whose Leadership value was below zero thus creating networks with activities only by the leaders. Now the community table contained the 16 networks generated by the leaders and their activity. From the distributions of activity type, we could also see that among the activities performed by the leaders, invitations were most prevalent (Figure 6). This indicates that leaders are busier recruiting new members in the community than replying to

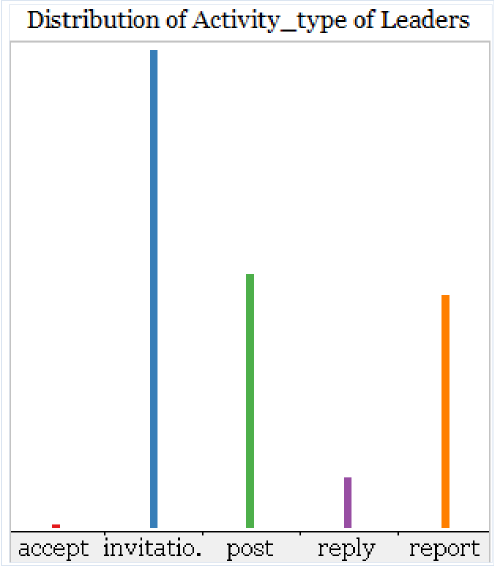

Figure 6: Activity type distribution of all the leaders.

discussions or submitting reports. Also most of the invitations in any community were sent by the leaders, and invitations sent by other members were sparse. From its Node count column, users could see only 16 communities had at least one leader, 3 of them had 2 leaders, all others having 1 (Figure 7 shows part of this table), and none of them had any member from law and enforcement. So, it is not the law enforcement people who are leading the communities.

Again, in the activity date column, we could see that the temporal patterns in the bottom 4 communities (marked with red in Figure 7) were very similar and they were all from Jefferson County. Although the leaders of these communities were initially very active their activities decreased over time. After filtering out these communities we could see a different scenario: for the remaining communities, the leaders were more involved in making posts and replies than the filtered out communities.

This observation made us interested in comparing the leaders' activity networks among different communities using node-link diagrams. In the node-link view the nodes were ranked by Leadership value of the members (red indicting the node with maximum Leadership value, blue being the lowest) (Figure 8). In the "Watch-Jefferson-County" community, there were many disconnected nodes, meaning the members were creating reports and invitations rather than having conversation with each other. The leader himself was also connected with a few other members as the connected component had only 10 nodes even though this was the largest community. In contrast, the "Duncuns" community had a lot of replies, so we opened its node-link diagram. This community's members had more conversations (posts and replies) as seen by its connected component and the s leader was also directly more connected with other members meaning the leader made posts and replied to others' posts.

After using ManyNets to identify the leaders we then exported the data to conduct a regression analysis in STATA. The regression analysis supported the hypothesis developed with the aid of ManyNets' visualization, i.e. the presence of super-active members strongly correlates with the growth of a community [29].

Growth Pattern of a Single Community

After sampling the communities, "Watch-Jefferson-County" community appeared to be the most active one, so we selected it for temporal analysis. We split it to observe its activity over time. In Figure 9, each row represents the network for a month, sorted by time from July 2009 to February 2011. Users could see that initially there were different types of activities but gradually the proportion of reports grew larger while no more invitations and acceptance occurred lately. This community accumulated members first and only after having enough members did they start posting crime reports and having discussion about community safety. As more people became involved in the community, the number of reports increased.

Label	Node Count	Relationship activity_type	Accept	Invites	Post	Reply	Report	TotActivity	Relationship activity_date
laurelbrooke:215	1		0	32	24	2	5	63	
kendallplace:204	1		0	29	15	2	7	53	
sylvangrove:136	1		0	6	1	0	21	28	
killianpines:185	1		0	12	16	5	8	41	
monrovia-hollow:51	1		1	21	28	3	0	53	
mtpleasant:127	1		0	4	6	1	16	27	
lickingknoxcounty:139	1		0	4	11	0	10	25	
shannondale:17	2		0	63	4	4	16	87	
mountain-watch:31	1		0	39	0	0	26	65	
charles-town:18	1		0	39	0	0	1	40	
ranson:23	1		0	39	0	0	1	40	

Figure 7: Network table comparing the activity of leaders in 11 communities.

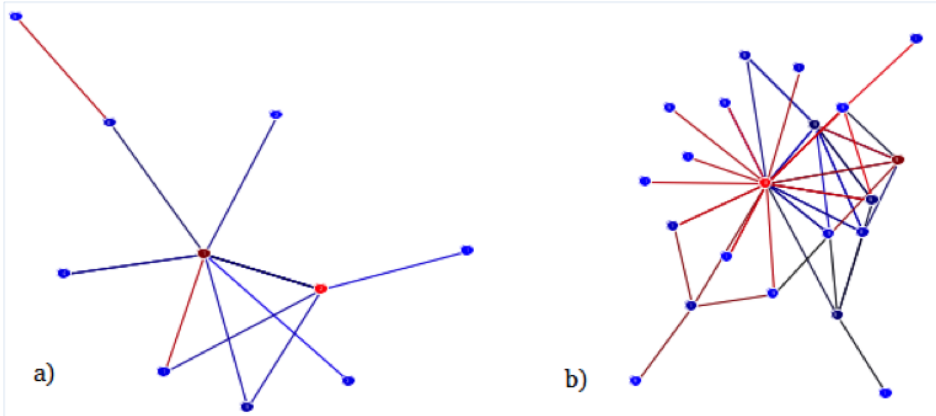

Figure 8: Node-link view of the connected components of a)Watch-Jefferson-County community and b)Duncans community. Red nodes are the leaders.

DISCUSSION

In successful and persistent communities, we observed the presence of leaders who are active most of the time and sending invitations to people. Community managers might want to encourage such leadership and support their activities online or offline (e.g. encourage them to arrange community safety activities).

We observed that as the number of members grows, there are fewer invitations sent and more reports posted. Highly active communities appear to have more reports than any other activity. We expected law enforcement involvement to heavily influence activity levels, but we found no evidence to support this hypothesis in this case study; however, the number of law enforcement officers is relatively small, so there may just be too few cases for adequate analysis. Also by clustering communities together we found out different trends of activities in communities.

Improvement Suggestion for ManyNets

We received several improvement suggestions while doing the case study. One suggestion was to expand the distribution columns on demand, so now the activity type column shows the distribution of activities as a histogram inside a table cell; the expansion option generates new columns for each type of activity-per-user request. The distribution column compares the relative proportion of each type of activity inside a community, whereas the expanded columns can perform cross-community comparison of each type of activity. Simple spreadsheet based tool cannot show and manipulate distribution columns. We also implemented on demand global and local scaling for the distributions within the table cells. If users want to compare the values inside the histograms with other rows, then the will choose global scaling; if the comparison is to be made with other activities in the same community, then they will choose local scaling of the histogram bar height. Also to make the comparison easier separate colors were assigned for each type of activity.

Improvement Suggestions for NON interface

Only a small percentage of community members are involved in crime reporting. If community managers can integrate their tool into popular social networking sites and have a way to export reports and discussion from those, then it might get more member participation.

We also observed that mostly the leaders are inviting people to the communities and other members are not so active in inviting more people. The interface should have more obvious options to encourage members to send invitations. Another observation was that even though members intended to reply to a particular post, they created a new post for that instead of replying to the initial post, so the interface had more room for improvement and relevant discussion can be performed within the same post and other members can follow up properly. As 10% of the replies were recoded, indicating that the interface was not sufficiently intuitive may need some additional design work.

The community members using the website may have modest technical skills, but they are concerned about their neighborhoods. So improving the website interface can increase their participation level.

CONCLUSIONS AND FUTURE WORK

This paper shows how we can start with an overview of the attributes of hundreds of communities and then filter down to successful communities, analyze their member activities, and identify the leaders using visual analytics. It also presents an analysis workflow along with case study of an online community safety platform. Having the capability to generate both statistical and visual insights integrated in the same tool along with its filtering features provided the leverage of rapid reiteration within one tool without going back and forth among several tools. This way ManyNets can help managers of web-based communities to analyze their community, learn how members are using the web-based platform and identify the successful cases. One remark from the NON community manager, Art Hanson

Label	Node Count	Relationship activity_type	Accept	Invites	Post	Reply	Report	TotActivity
2009-07Jul	16		7	17	12	11	11	58
2009-08Aug	4		2	3	5	2	9	21
2009-09Sep	7		1	11	7	3	10	32
2009-10Oct	6		0	18	8	5	11	42
2009-11Nov	12		1	5	2	1	16	25
2009-12Dec	7		0	2	3	2	6	13
2010-01Jan	8		0	10	3	2	10	25
2010-02Feb	6		0	0	1	0	6	7
2010-03Mar	4		0	1	0	0	9	10
2010-04Apr	9		0	0	0	0	10	10
2010-05May	7		0	0	3	0	8	11
2010-06Jun	9		0	0	0	0	19	19
2010-07Jul	7		0	0	1	2	11	14
2010-08Aug	10		0	1	2	1	16	20
2010-09Sep	9		0	0	2	1	19	22
2010-10Oct	7		0	0	0	2	10	12
2010-11Nov	6		0	0	1	2	5	8
2010-12Dec	6		0	0	1	1	4	6
2011-01Jan	6		0	0	0	0	11	11
2011-02Feb	6		0	0	1	0	6	7

Figure 9: Temporal splits of the Watch-Jefferson-County community. After January 2010, more reports are posted whereas number of invitations dropped.

was: *"Your observations and analysis of what contributes to a 'successful' community will be very helpful going forward - I am hoping to implement some of your measures as built-in tools to help our community managers".*

In the future, we want to analyze the passive activities, too. Members who read the posts and visit the websites but do not participate actively are also very important for the success of an online community. Another direction of work is to perform content analysis of the conversations, visualization of the topic distribution inside the tool and observation of their temporal changes. Finally possible direction is to suggest possible interventions to the manager that are likely to increase participation.

ACKNOWLEDGMENTS

We cordially thank Professor Alan Neustadtl, Manuel Freire, graduate student Jee-hye Kang and Robert Gove and NON community manager Art Hanson for their valuable feedback and development suggestions. This work is supported in part by the National Science Foundation, grant # IIS - 0968521 "Supporting A Nation of Neighbors with Community Analysis Visualization Environment".

REFERENCES

1. Adar, E. and Adamic, L. A. Tracking information epidemics in blogspace. In *Proc. Web Intelligence, 2005. Proceedings. The 2005 IEEE/WIC/ACM International Conference on* (2005), 207 - 214.
2. Ahn, J., Taieb-Maimon, M., Sopan, A., Plaisant, C. and Shneiderman, B. Temporal Visualization of Social Network Dynamics: Prototypes for Nation of Neighbors. *Social Computing, Behavioral-Cultural Modeling and Prediction* (2011), 309-316.
3. Baker, T., Baker, J. and Zezza, R. Neighborhood watch: A leadership challenge. *FBI Law Enforcement Bulletin*, *68*, 2 (1999), 12-18.
4. Bennett, T. *Evaluating neighbourhood watch*. Gower Pub Co., (1990).
5. Bennett, T., Holloway, K., Farrington, D. and Bennett, T. The effectiveness of neighborhood watch. *Campbell Systematic Reviews*, *18* (2008).
6. Butler, B. S. Membership size, communication activity, and sustainability: A resource-based model of online social structures. *Information Systems Research*, *12*, 4 (2001), 346-362.
7. Cirel, P., of Law Enforcement, N. I. and Justice, C. *Community Crime Prevention Program: Seattle, Washington: An Exemplary Project*. US Government Printing Office, (1977).
8. Doreian, P., Kapuscinski, R., Krackhardt, D. and Szczypula, J. A brief history of balance through time. *The Journal of Mathematical Sociology*, *21*, 1 (1996), 113-131.
9. Durant, K., McCray, A. and Safran, C. Modeling the Temporal Evolution of an Online Cancer Forum. In *Proc. 1st ACM International Health Informatics Symposium (IHI 2010)* (2010), 356-365.
10. Freire, M., Plaisant, C., Shneiderman, B. and Golbeck, J. ManyNets: an interface for multiple network analysis and visualization. In *Proc. Proceeding of the 28th annual SIGCHI conference on Human factors in computing systems* (2010), 213-222.
11. Gloor, P., Laubacher, R., Zhao, Y. and Dynes, S. *Temporal visualization and analysis of social networks*. In *NAACSOS Conference*, (2004), 27-29

12. Gloor, P. A. and Zhao, Y. *Analyzing actors and their discussion topics by semantic social network analysis*. In *Tenth International Conference on Information Visualization (IV 2006)*, (2006), 130–135.
13. Hansen, D., Shneiderman, B. and Smith, M. Visualizing Threaded Conversation Networks: Mining Message Boards and Email Lists for Actionable Insights. In A. An, P. Lingras, S. Petty and R. Huang (eds.) *Active Media Technology*, *6335*, (2010), 47-62.
14. Hanson, A. *Nation of Neighbors*. (2012).
15. Heer, J. and Boyd, D. Vizster: visualizing online social networks. In *Proc. Information Visualization, 2005. INFOVIS 2005. IEEE Symposium on* (2005), 32 -39.
16. Holloway, K., Bennett, T. and Farrington, D. P. Crime Prevention Research Review No. 3: Does Neighborhood Watch Reduce Crime? *Washington, DC: US Department of Justice Office of Community Oriented Policing Services* (2008).
17. Katz, N., Lazer, D., Arrow, H. and Contractor, N. Network Theory and Small Groups. *Small Group Research*, *35*, 3 (2004), 307-332.
18. Kumar, R., Novak, J. and Tomkins, A. Structure and Evolution of Online Social Networks. In P. S. S. Yu, J. Han and C. Faloutsos (eds.) *Link Mining: Models, Algorithms, and Applications*, (2010), 337-357.
19. Leskovec, J., Backstrom, L., Kumar, R. and Tomkins, A. Microscopic evolution of social networks. In *Proc. KDD '08: Proceeding of the 14th ACM SIGKDD international conference on Knowledge discovery and data mining*, ACM (2008), 462-470.
20. Leskovec, J., Kleinberg, J. and Faloutsos, C. Graph evolution: Densification and shrinking diameters. *ACM Trans. Knowl. Discov. Data*, *1* (2007).
21. Lindsay, B. and McGillis, D. Citywide community crime prevention: An assessment of the Seattle program. *Community crime prevention: Does it work* (1986), 46-67.
22. Lodha, S. K. and Verma, A. K. Spatio-temporal visualization of urban crimes on a GIS grid. In *Proc. Proceedings of the 8th ACM international symposium on Advances in geographic information systems*, ACM (2000), 174-179.
23. Moody, J., McFarland, D. and Bender-deMoll, S. Dynamic network visualization. *American Journal of Sociology*, *110*, 4 (2005), 1206-1241.
24. Perer, A. and Shneiderman, B. Balancing Systematic and Flexible Exploration of Social Networks. *IEEE Trans. Vis. Comput. Graph*, *12*, 5 (2006), 693-700.
25. Peterson, E. Time spring layout for visualization of dynamic social networks. In *Proc. Network Science Workshop (NSW), 2011 IEEE* (2011), 98 -104.
26. Plaisant, C., Komlodi, A. and Lindsay, F. Evaluation challenges for a Federation of heterogeneous information providers: the case of NASA's Earth Science Information Partnerships. In *Proc. Enabling Technologies: Infrastructure for Collaborative Enterprises, 2000. (WET ICE 2000). Proeedings. IEEE 9th International Workshops on* (2000), 130 -135.
27. Powell, W. W., White, D. R., Koput, K. W. and Owen-Smith, J. Network Dynamics and Field Evolution: The Growth of Interorganizational Collaboration in the Life Sciences. *American Journal of Sociology*, *110*, 4 (2005), 1132-1205.
28. Preece, J. Sociability and usability in online communities: determining and measuring success. *Behaviour and Information Technology*, *20*, 5 (2001), 347-356.
29. Rey, P. J., Sopan, A., Ahn, J., Kang, J. H., Plaisant, C. and Shneiderman, B. The influence of Leadership on Activity Levels on an Online Crime Watch Community. Technical report, submitted to *Proc. American Sociology Association Annual Meeting* (2012).
30. Rosenbaum, D. P. Community crime prevention: A review and synthesis of the literature. *Justice Quarterly*, *5*, 3 (1988), 323-395.
31. Rosenbaum, D. P., Lewis, D. A. and Grant, J. A. Neighborhood-based crime prevention: Assessing the efficacy of community organizing in Chicago. *Sage Criminal Justice Annuals*, *22* (1986).
32. Skogan, W. G. *Disorder and decline: Crime and the spiral of decay in American neighborhoods*. Univ of California Press, (1992).
33. Trier, M. Towards Dynamic Visualization for Understanding Evolution of Digital Communication Networks. *Information Systems Research*, *19*, 3 (2008), 335-350.
34. Tyler, J. R., Wilkinson, D. M. and Huberman, B. A. E-mail as spectroscopy: Automated discovery of community structure within organizations. *The Information Society*, *21*, 2 (2005), 143-153.
35. Viegas, F. B. and Smith, M. Newsgroup Crowds and AuthorLines: visualizing the activity of individuals in conversational cyberspaces. In *Proc. System Sciences, 2004. Proceedings of the 37th Annual Hawaii International Conference on* (2004), 10 pp.
36. Violi, N., Shneiderman, B., Hanson, A. and Rey, P. J. Motivation for Participation in Online Neighborhood Watch Communities: An Empirical Study Involving Invitation Letters. In *Proc. SocialCom/PASSAT* (2011), 760-765.
37. Wiil, U. K., Memon, N. and Gniadek, J. CrimeFighter: A Toolbox for Counterterrorism. *Knowledge Discovery, Knowlege Engineering and Knowledge Management* (2011), 337-350.
38. Xu, J. and Chen, H. Criminal network analysis and visualization. *Commun. ACM*, *48* (June 2005), 100-107.
39. Sopan, A., Freire, M., Taieb-Maimon, M., Golbeck, J., Shneiderman, B., Shneiderman, B.: Exploring Distributions: Design and Evaluation. *HCIL tech report*(2010).