Organizational Response to the Turmoil of Personnel Turnover[1]

Terrill L. Frantz & Kathleen M. Carley
{terrill, kathleen.carley}@cs.cmu.edu
Carnegie Mellon University

Abstract

We explore the dynamics of a social network within a business organization during momentous personnel turnover events—by applying dynamic network analysis techniques to the Enron email corpus. We investigate how the complex network structure of a real-world organization responds when the appointment of a new CEO is announced, then later resigns; as well as, when the organization is in crisis and announces bankruptcy with massive layoffs. Our findings indicate that the social network became perceptibly more hierarchical when the CEO appointment was made and considerably less hierarchical when the CEO resigned. Further, we found that there is evidence that the organization also became slightly more hierarchical immediately when the bankruptcy and layoff were announced.

Key Words: executive turnover, crisis response, Enron, social network analysis, dynamic social networks, communication networks, email networks

1 Introduction

Organizational crisis and executive turnover is an all too often occurrence in today's society. To the extreme, we have seen the implosion of Enron and Arthur Andersen. On a less destructive scale, there are smaller crises in every day organizational life. Further the organizational ecological system makes for, perhaps, planned crises in organizations. And, being an open system, organizations can be thrust into a crisis from an external event.

It is important to understand the internal dynamics of an organization when it is subjected to a major personnel event—it may be inevitable that one will occur in any organization, large or small. How will the organization respond? What is an atypical response? What is normal? What comes next? How does the social network change, if at all? While countless aspects of the organization may be affected by such events; we ponder the organization-wide response to these events from the perspective of the informal network of relationships within the organization.

The informal network can be a "double edge" sword [FLA1] in an organization. The relationship an employee has with another, if positive, will likely increase their cooperation with one another; or occasionally relationships may indeed, have the opposite effect and decrease their cooperation. Regardless of the outcome, the informal network most certainly has an effect on the organization, and it is therefore, important to understand the dynamics of these relationships.

The broad availability of the Enron email corpus provides researchers with a rich set of data for studying the dynamics of a real-world organization during an organizational crisis. This collection of a half-million emails, provides a history of employee interaction throughout a period in which Enron shifted from a successful organization, to one in turmoil, and, finally, to its dissolve.

During the nearly four-year period (1999-2002) that we will study, Enron internally—and externally—experienced employee turmoil and senior leadership turnover: corruption and wrong-doing inside Enron became public via an employee whistle-blower; from improper financial transactions, the company was forced to declare bankruptcy; it faced public scrutiny in the form of government investigations, and; ultimately, Enron folded. What was once America's most innovative company [FOR1], had tragically imploded and had helped to undermined public confidence in the broader American business community [GEO1].

Three events (out of many that led to the ultimate demise of Enron) that are material enough to possibly affect the whole of the Enron organization are: (a) the promotion of Jeffrey Skilling to CEO on December 13, 2000, (b) Skilling's resignation from the same CEO post eight months later, on August 14, 2001, and (c) Enron filing for bankruptcy, coupled with massive layoffs on December 2-3, 2001. These events are of particular relevance because of their involving highly-public personnel turnover. The Skilling promotion and resignation are atypical executive personnel events, while the

[1]This paper is part of the Dynamics Networks project in CASOS (Center for Computational Analysis of Social and Organizational Systems, http://www.casos.cs.cmu.edu) at Carnegie Mellon University. This work was supported in part by the Office of Naval Research (ONR), United States Navy Grant No. 9620.1.1140071 on Dynamic Network Analysis under the direction of Rebecca Goolsby. Additional support on measures was provided by the DOD and the NSF. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the National Science Foundation or the U.S. government.

near-total layoff and bankruptcy are, fortunately, less frequent events.

Work on executive turnover suggests a flurry of activity as a new executive comes on board which dissipates with time; work on organizational change, in general, suggests that over time organizational interaction stabilizes. Might there be evidence of a flurry of change in the structure of interactions, to be followed by a movement toward stability?

We posit that the structure of Enron's social network changed during the different phases of the shifting organizational climate. We suppose that the structure of the internal interactions, thus the informal network, at times becomes less fluid, less networked, more hierarchical and more rigid; this phenomena of "closing rank" likely occurred at times that correspond to disruptive events, possibly at leadership change, legal announcement, and such.

1.1 Research Questions

Using Enron's internal email communications history as empirical evidence and applying social network analysis techniques to the data, we explore the dynamics of the internal social network and correlate the structural characteristics with the actual event history to answer the following research questions:

(a) Is there evidence that Enron's structure of interactions became less fluid, less networked, more hierarchical and more rigid during the crisis period—as crisis theory suggests?

(b) Did the interactions within Enron follow a pattern of outburst followed by movement towards stability as a response to executive turnover events?

To answer these questions, we need to undertake the challenging process of studying the network as a whole, which can be daunting [SHR1].

2 Methodology

We obtained the version of the Enron email corpus dataset, which is publicly available, from the University of Southern California's Information Sciences Institute (ISI) [ISI1]. This version originates from the much-studied, Carnegie Mellon University dataset supplied by William Cohen [COH1]; however, the ISI version we used has been post-processed by removing blank, duplicate, and other unusable emails. The ISI pruning process also removed computer-generated messages that typically are sent to large number of recipients automatically—these are not of importance to this study because they do not reflect the social, human-to-human interaction we are investigating. Our use of the ISI dataset simply saved duplicate, manual processing effort on our part; the content of the ISI dataset is certainly of a more ready-to-use quality.

From the original 500 thousand emails in the CMU dataset, the ISI dataset is pruned into a dataset of about 250 thousand high-quality emails; the number of unique email addresses in this subset, number over 20,000. From this dataset, we further filtered the data so that only Enron-to-Enron emails are analyzed. We removed any emails that were not sent from an Enron email address to at least one other Enron email address. Any non-Enron email addresses were discarded from the dataset, thus not affecting the network analysis. To determine an Enron address, we performed a simple text compare on the address using a case-insensitive regular expression looking for the word token "enron". This process captures addresses from both the main Enron organization as well as most of the subsidiaries of the company.

The raw data is separated into weekly bins, using the ISO 8601 calendar standard, with slight modification; this is consistent with the scheme Microsoft uses in the Excel spreadsheet program. ISO 8601 corresponds with the Gregorian calendar; it uses the same year number—in the form, *ccyy*—, but it defines a methodology for enumerating the weeks within the calendar year. ISO 8601 defines a week-interval of seven days, starting on Monday, and assigns an ordinal number for each week from 1 to 52, or 53. For processing simplicity, we follow the Microsoft Excel method of always making January 1, week number 1 of its year, rather than sometimes placing January 1 in the last week of the prior year. We also fold week number 53, into week number 52; leaving 52 weeks as a fixed number of weeks in any year. The straddling process we choose to implement, with its systemic pros and cons, will not affect our analysis in any meaningful way for the analysis we are doing in this study. This process resulted in 165 weekly bins— numbering from 1999-20 to 2002-28, May 10, 1999 to July 14, 2002 respectively.

For each of the weekly bins, an adjacency matrix of the email activity for each was constructed. Each email transaction—a unique email send from one email address to one or more recipients—was parsed to identify the single sender, and the, often multiple, receiver(s). These email addresses were added to the rows and columns of the adjacency matrix and a unity value was placed in the intersection of the send address and each of the receiver addresses. Essentially, we ultimately have constructed a square matrix that maintains the directionality of each transaction. We do not maintain the number of occurrences for the sender-receiver pairs, simply a binary (0/1) indication of a directed interaction between the addresses.

As a result of this procedure, we have 165 distinct adjacency matrices that liken to 165 one-mode, un-weighted digraphs. Each digraph correspond to a network that represents the history of email transactions, specifically an indication of at least one email being sent between two individual email addresses These digraph networks are the basis for calculating several network measures.

2.1 Computed Measures

We compute two distinct types of descriptive measures from each network; they are basic exploratory data analysis indicators and measures that quantify the structural characteristics of the network. We form a temporal array of these values for our analysis.

The basic exploratory data analysis measure that are calculated for each time-based instance of the network include: the count of the emails sent, the number of distinct email addresses in the network (without any distinction of being a sender or a receiver), and the number of unique ties (between an ordered pair). We do not present the density of these networks because all the networks are extremely sparse, with density rarely approaching 1%.

To measure and quantify the underlying structure, we evaluate each network according to four network hierarchy measures developed specifically for this purpose [KRA1]: connectedness, graph hierarchy, efficiency, and least upper bound (LUB). These four measures quantify the four conditions Krackhardt considered necessary for a graph to be considered a hierarchy—specifically an *outtree*. All four measures are computed as a real value from 0 to 1.

Connectedness is defined by Krackhardt as:

$$\text{Connectedness} = 1 - [V / N(N-1)/2)] \qquad (1)$$

where V is the number of pairs of points that cannot reach one another, which is the numerator divided by the total number of possible ties: N is the number of nodes and total possible ties in a undirected graph is N(N-1)/2. This is a measure of the undirected underlying graph that indicates the connectedness of a node with any other node in the network, either directly adjacent or via a walk. This is a measure of how separated nodes are from one another.

Graph hierarchy is defined as:

$$\text{Graph Hierarchy} = 1 - [V/\text{MaxV}] \qquad (2)$$

where V is the number of pairs of points that cannot reach one another, which is the numerator, divided by MaxV, which is the total number of pairs of points where there is a tie present between them. This measure indicates the reachability of nodes with regard to the directionality of the ties, that is, as per the digraph. The closer this value is to 1, the more strictly ordered are the relationships in the communications network. That is, an email transaction from address A does not also have a paired transaction from B back to A (in the same network snapshot.)

Graph efficiency is defined as:

$$\text{Graph efficiency} = 1 - [\ V / \text{MaxV}\ ] \qquad (3)$$

where V is the number of links greater than Nn-1, and MaxV is the maximum number of ties possible. This is an indicator of the number of redundant ties in the underlying network. A high value, close to 1, implies that the network is fully connected and that the redundant ties make any single node, or tie, not critical to the communicative success in the graph.

Least Upper Boundedness (LUB), the most complex measure of the four to calculate, is defined as:

$$\text{LUB} = 1 - [V / \text{MaxV}] \qquad (4)$$

where V is the number of nodes that LUB and MaxV is the maximum possible nodes that could possibly have no LUB. The least upper bound is a node that is common to a pair of nodes, through a geodesic for each of the pair to the common LUB node. Krackhardt offers that the LUB is the common boss for two employees, either directly or up though the organization chart.

$$\text{MaxV} = (Nn-1)(Nn-2) / 2 \qquad (5)$$

These measures, for this study, are calculated using the Organization Risk Analyzer (ORA) [CAR1] software. We, first, formatted each of the 165 networks into DYNETML format, which is an XML-based standard for representing complex social networks; then, we ran each of the networks in ORA, resulting in separate result-sets. These result-sets where then combined into one for comparative analysis.

2.2 Critical-Event Timeline

The timeline of pertinent events is as follows:

(a) 2000 Dec. 20 (week 2000-49)
- Skilling named CEO

(b) 2001 Aug. 14 (week 2001-33)
- Skilling resigns as CEO

(c) 2001 Dec. 2-3 (weeks 2001-49 & -50)
- Bankruptcy filing and mass lay-off

3 Results

In this section, we provide a statistical profile of the sample network data over time by providing basic count measures, along with measures specific to the network structure, as described in the prior section.

To start,—for the exploratory data analysis step—three basic count measures are calculated as a foundation for providing a context for evaluating the more complicated measures. The number of emails sent in each week is determined, as are the number of unique email addresses and unique pair-wise ties among the email addresses.

Each figure in this section has three vertical bars positioned on the graph to identify the specific weeks that pertain to the three events. The left-most bar is the Skilling to-CEO announcement, the middle bar is the Skilling resignation announcement, and the right-most vertical bar is the bankruptcy and layoff announcement.

Figure 1 shows the volume of emails sent each week, over time. The average number of emails per week is 1,110 (*n*=165, std. dev. =1,134.8), with the maximum being 7,474 emails, which were sent around the last week of October 2001. There appears to be systematic reduction in the number of emails sent each year, that correspond with the year-end holiday period (weeks 50, 51, and 52); this would be expected for a commercial business in the United States.

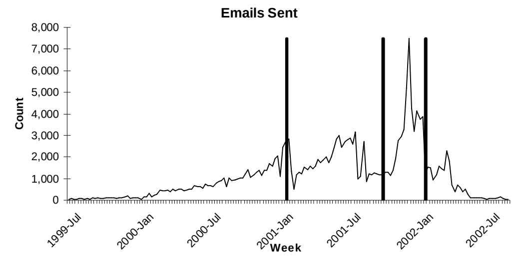

Figure 1. Number of emails sent each week.

Figure 2 shows the number of distinct email addresses tied to any email (as sender, or receiver) in a given week—either as a sender or a receiver. The average number of addresses each week is 1,503 (*n*=165, std. dev. =1,202.7), with the maximum being 4,759 distinct email addresses, which occurred during the middle of October 2001. These are *unique* email addresses, thus, for example, if two emails are sent from the same email address in the same week, there would only be one address recorded in this count.

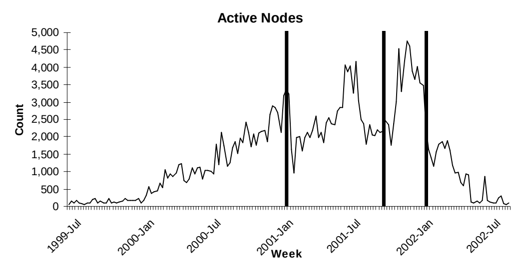

Figure 2. Number of active email addresses per week.

Figure 3 shows the number of unique ties between a pair of email addresses, according to an email connection—a sender tied to a receiver of any type—, in a given week. The average number of weekly ties is 3,365 (*n*=165, std. dev. =3,260.9), with the maximum being 14,966 ties during the end of October 2001. These are unique ties, thus for example, if one address sends two separate emails to the same receiver, there is only one tie recorded in this count.

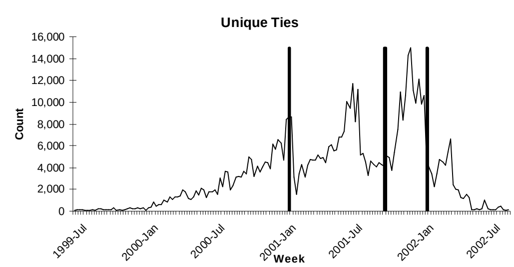

Figure 3. Number of unique sender-to-receiver ties per week.

Figure 4 shows the connectedness measure over the 165 weeks. The average value is 0.882 (n=165, std. dev. = 0.168). This measure is an indication that nearly all email addresses are connected to all others, albeit, via a walk of the underlying graph. It appears that the connectedness fluxuated and varied little during the weeks of the CEO and bankruptcy/lay-off events. It is readily apparent that there is a large range difference and variability in the tails of the dataset.

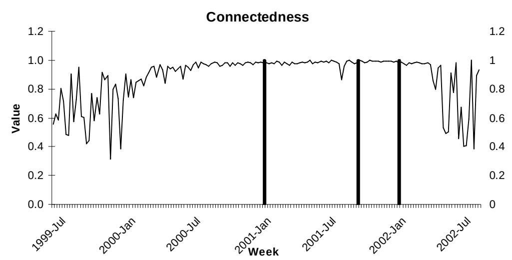

Figure 4. Connectedness measure over time.

Figure 5 shows the hierarchy measure over the 165 weeks. The average value is 0.943 (n=165, std. dev. = 0.025). It can be readily seen that the hierarchy measure rose suddenly and rapidly when Skilling as appointed CEO and that when he resigned, the hierarchy value began a multi-week slide to significantly less heirachy value. It also appears that there may have been a distinct long-term shift in the value post CEO resignation. It also appears that the bankruptcy and layoff announcement resulted in a short term increase in hierarchy, albeit not as great as the CEO-related events.

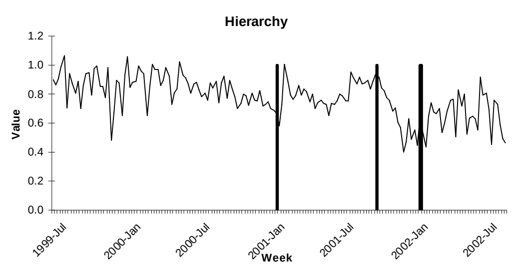

Figure 5. Hierarchy measure over time.

Figure 6 shows the efficiency measure over the 165 weeks. The average value is 0.996 (n=165, std. dev. = 0.007). This value is consitently within a very small range and fluxated very little during the 165 weeks in the dataset. Like the connectedness measure, it appears, however, that there is some change in variability and the range at both tails. There are many more ties in the network than are necessary y y y y to keep the network connected. There is little concern for possible fractions occurring in the communications network.

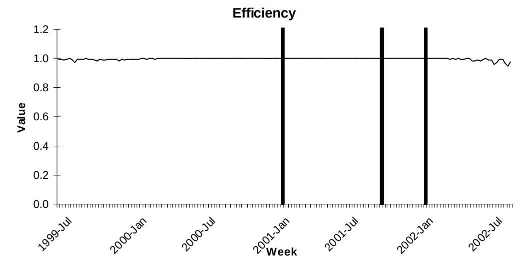

Figure 6. Efficiency measure over time.

Figure 7 shows the upper bound measure over the 165 weeks. The average value is 0.760 (n=165, std. dev. = 0.200). While it can be easily seen that immediately after each of the three events, that the upper bound decreased, the variance of the measure throughout the entire dataset held this same pattern of fluxation.

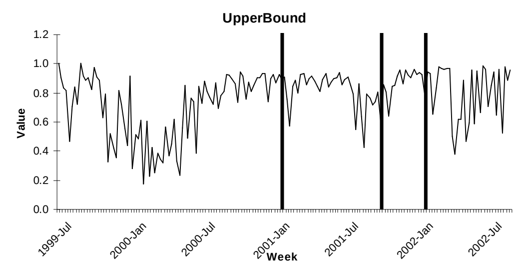

Figure 7. Upper Bound measure over time.

These seven measures are merely a handful of the entire set of social network analysis measure available to analysts. For our analysis, however, they provide several useful perspectives leading to a quantitative profile for the network and its dynamics over time, particularing during business-as-usual and during significant turnover events. We discuss these results, from a substantive organization-dynamics perspective, in the next section.

4 Discussion

The results indicate that the social network had changed subsequent to the major personnel-events. While three of the structural measures (connectedness, efficiency, and upper bound) provided no evidence of any noteworthy change, the hierarchy measure did indicate that, indeed, there were momentous changes to the structure following these events.

Our first research question: “Is there evidence that Enron’s structure of interactions became less fluid, less networked, more hierarchical and more rigid during the crisis period—as crisis theory suggests?"; can be answered in the affirmative. Indeed, the structure changed, as the hierarchy measure unmistakably indicates.

The evidence is clear that the organization become *more* hierarchical immediately following the *CEO-appointment* event. While not as noticeable, it seems the organization also became somewhat more hierarchical following the bankruptcy/layoff event.

There is also clear evidence that the organization became *less* hierarchical following the *CEO-resignation* event. In the months following this event, it appears that structure may have even shifted to a new equilibrium level of hierarchy.

We reason that our findings are consistent with the more generalized theories that during times of uncertainly, organizations become more decentralized as individual search for information from a broader group of colleagues in their ego network.

5 Limitations

The primary limitation of this study, along with any study using the Enron email corpus, is the original data collection process. The data was collected in a manner not fully documented, thus has some likely bias in the process. The data is available from a public disclosure order on Enron. The true process for what data was submitted to the court, and what was not, is unknown.

Another limitation or this study is the fact that only one sample organization has been explored. While our findings are useful, this study alone cannot claim much generalize-ability of the findings with this single sample.

We recognize that social network structures can be materially different at different levels of analysis. In this study we did not investigate other frames of the data. For example, constructing the weekly networks using only the email receiver fields (to, cc, bcc) equally, versus weighting the relationship according to the addresses' address field may provide a different structure. We also recognize that using a different binning rule when separating the data, for example, studying the daily or monthly views, may possibly lead to different results.

There are other proposed methods [TAN1] for measuring hierarchy of an organization that may, or may not, prove to be more valid in assessing the structure. For example, Mackenzie [MAC1] proposes an alternative to the Krachhardt [KRAx] proposal that we applied in this study. To date, no cross-comparison has been performed on these different methods. It is possible that they may yield different findings when applied to the same data; we have not confirmed or rejected this possibility.

6 Conclusions and Future Work

In this paper, we describe how the Enron email corpus was studied with a focus on the social network structure following three public events, namely, events that involve very-public personnel turnover. We constructed social networks from the email interactions over a nearly four-year period and took weekly measurements that quantify the structure of the network.

The results lead us to conclude that the network did indeed change its structure, immediately and over time, following such events. In this case, when a CEO was announced, immediately the social network became much more hierarchical in the near term, then returned to its prior equilibrium state. When the CEO announced his resignation from the post, the social network began a long period of changing its structure from this more hierarchical state to much less rigid form. It seems that this change may have permanently shifted the social network into a new equilibrium of a much less hierarchical structure. We also conclude from this study that, in this case, the bankruptcy and layoff announcement preceded a less sudden and less dramatic change in the structure of the social network to becoming only slightly more rigid, or hierarchical.

Our findings and conclusions are important pieces of evidence that can be considered along with other research to reach a better understanding of broadly defined organization dynamics and more narrowly, the organizational response to personnel turnover. We cannot formulate generalized concludes from this study—this is only one sample in an area with much variation and complexity—but these findings do introduce evidence for the ultimate formulation of a generalized theory.

From this study, we see several areas for future related work. Mainly, as more email corpus datasets become available from other real-world organizations, this same analysis should be conducted in order to develop a comparison between the organizations — it is likely that all organizations will not respond to similar events in precisely the same way. We also envision a deeper analysis into the social network following these events using text analysis techniques. At this stage we can only ponder what the content of the communications were following these events; is the communication directive, inquisitive, or gossipy?

Regardless, this study, along with many others, touches only the surface of organization dynamics; there is much more research necessary.

7 References

[CAR1] Carley, K. M., & Reminga, J. (2004). ORA: Organization risk analyzer. Technical Report, Carnegie Mellon University, Institute for Software Research International. CMU-ISRI-04-106

[COH1] Cohen, W. W. (n.d.) Enron Email Dataset. See http://www.cs.cmu.edu/~enron/

[FLA1] Flache, A. (1996). *The double edge of networks.* Amsterdam: Thesis Publishers.

[FOR1] Fortune Magazine

[GEO1] George, A. M., & Evuleocha, S. (2003). Denials, excuses, justifications, and aopologies: Restoring tarnished reputations after the years of corporate malfeasance. What worked and what didn't. In Proceedings of the 2003 Association for Business Communication Annual Convention.

[ISI1] Shetty, J., & Adibi, J. (n.d.). The Enron Database schema and brief statistical report. Retrieved November 4, 2004, from www.isi.edu/~adibi/Enron/Enron_Dataset_Report.pdf

[KRA2] Krackhardt, D., & Stern, R. N. (1988). Informal networks and organizational crises: An experimental simulation. *Social Psychology Quarterly*, 51(2), 123-140.

[KRA1] Krackhardt, D. (1994). Graph theoretical dimensions of informal organizations. In Kathleen M. Carley and Michael J. Prietula (Eds.), *Computational Organization Theory*, pp. 89-111. Hillsdale, NJ: Lawrence Erlbaum Associates.

[MAC1] Mackenzie, K. D. (1978). *Organizational structures*. Arlington Heights, IL: AHM Publishing.

[SHR1] Shrader, C. B., Lincoln, J. R., & Hoffman, A. N. (1989). The network structures of organizations: Effects of task contingencies and distributional form. *Human Relations, 42*, 43-66.

[TAN1] Tangmaunarunkit, H., Govindan, R., Jamin, S., Shenker, S., & Willinger, W. (2003). *On characterizing network topology*. University of Southern California, Computer Science Department, Technical report, 03-782. CiteSeer: http://citeseer.ist.psu.edu/605282.html