The Economics of Mass Surveillance
and the Questionable Value of Anonymous Communications

George Danezis[1] and Bettina Wittneben[2]

[1] K.U. Leuven, ESAT/COSIC,
Kasteelpark Arenberg 10, B-3001 Leuven-Heverlee, Belgium.
george.danezis@esat.kuleuven.be
[2] RSM Erasmus University Rotterdam,
Burg. Oudlaan 50, 3062 PA Rotterdam, The Netherlands
bwittneben@rsm.nl

Abstract. We present a model of surveillance based on social network theory, where observing one participant also leaks some information about third parties. We examine how many nodes an adversary has to observe in order to extract information about the network, but also how the method for choosing these nodes (target selection) greatly influences the resulting intelligence. Our results provide important insights into the actual security of anonymous communication, and their ability to minimise surveillance and disruption in a social network. They also allow us to draw interesting policy conclusions from published interception figures, and get a better estimate of the amount of privacy invasion and the actual volume of surveillance taking place.

1 Introduction

Much contemporary cryptography and computer security focuses on securing communications between Alice and Bob [12, 7]: link confidentiality preserves the secrecy of what they say, authentication ensures that they are talking to each other, anonymity can hide their identities from each other or third parties. At the same time it is futile to deploy such a system, without realising that confidentiality, anonymity and privacy for Alice and Bob, do not exist in a vacuum. In most cases Alice and Bob would be embedded in a social network [16], and their identities and conversations would not only leak information about themselves but also about other actors in the network. Similarly, Alice and Bob's confidential conversations or secrets are likely to be communicated on other links, and may be compromised by observing third parties communicating them.

We start from the standpoint that insecurity has externalities – allowing Eve, an eavesdropper, access to information transiting on the links between Alice and Bob, additionally leaks other people's secrets. At the extreme surveillance on the link may only be there to gain information about third parties. Based on this intuition we define a model of actors participating in spaces, or clubs: if one of the club participants is under surveillance we assume that all information shared

in this club, and the membership list of the club becomes known to Eve. Some questions naturally arise from this model:

- How many people have to be under surveillance to observe all clubs?
- How many people have to be under surveillance to discover most people (their existence but not necessarily their membership to other clubs)?
- How best to choose those to be put under surveillance to maximise returns in terms of the last two questions (effective target selection)?
- How does the lack of information, due to the use of an anonymizing networks, affects the effectiveness of such target selection?

We could attempt to answer these questions by looking at a synthetic family of social networks [9]. Instead, we chose to perform our experiments on harvested social network information, derived from a real world political network. This approach adds authenticity, and credibility to our results since it is easier to argue that the harvested data represents real-world organisational models, than if we used fully synthetic networks. At the same time, the process of answering the questions, from the collection of the data to the analysis, illustrates the methodology that Eve would have to follow in order to put a similar social network under surveillance, given a limited interception budget – very little has been written about this in the past. Even less has been said about how to systematically protect such a network, at all levels, from such an adversary. In particular we show that deployed anonymization solutions protect only partially against target selection that may lead to efficient surveillance.

Both the fresh set of questions that we seek to answer and the realistic nature of the data set used, lead to results with far reaching policy implications. We illustrate this by re-interpreting actual reports of surveillance both legal and illegal, technical and physical, under the light of our models. We show that even low levels of reported surveillance may well be the tip of a massive surveillance operation iceberg.

2 Model & Data Set

We need to define precisely our model before posing or answering our key questions. We consider a social network that has two types of edges or nodes: *people* and *spaces* (sometimes referred for clarity as *clubs*). People may belong to spaces, which we call a *relationship*. We allow relationships to have some strength denoted by a positive integer. This symbolises the relative degree of association of the person to the space.

The social network is in fact a weighted bipartite graph from the set of people to the set of spaces. We do not allow direct links between people, and consider instead that all communication is mediated through a mutually shared club. This allows us to focus on the shared aspects of security, rather than studying traditional pairwise security.

2.1 Extracting the Network from Harvested Data

We were allowed access to some of the archives of the mailing lists used by a large international political network. The earliest posts to some lists were from 2003, and the lists were active at the moment we collected our data, in February 2006. Some other lists used by the same network were private and we did not have access to them – which means that the data set can only be seen as a subset of the 'public' interactions of the people involved. We made no attempt to map direct interactions between participants, but only considered mail sent through the mailing lists.

We mapped each mailing list to a space and each individual email address to a person. We recorded the day each message was sent by each participant to a mailing list for the observed period. We then aggregated the network links over time: if any message was sent between a person and a list, we created a relation, with weight equal to the overall number of messages observed between the parties.

The methodology used only created relations between senders and the mailing lists – it is possible that many members of the mailing lists acted only as passive observers during that period and have not been recorded in our network graph. Those passive members may in fact be valuable when reasoning about the resilience of the network [14], or the speed by which information can be disseminated. Unfortunately we did not have the means to uncover them, and this again means that we only have a partial picture of the network, which limits our ability to draw fuller conclusions.

The mailing list software performed some spam filtering that was not perfect. This means that some unsolicited mail has been archived. This is a problem since spammers have a very large degree, as they usually send their emails to many lists at the time. Yet filtering them on the basis of degree only, may distort our data, and may remove the most important genuine nodes from our data set. Instead we note that spammers have a very low weight per link – they usually only send one email using the same identity. To filter those out, to keep only genuine people we chose to remove all links of strength less than 5. This means that we consider a person associated with a list if they have sent to that list, using the same identity, no less than 5 emails. This threshold was chosen and only keep meaningful relations, and avoid the inclusion in our data set of very occasional relations (such as one email in the three year period processed.)

2.2 Characteristics of the Studied Network

The studied network contains 2338 people in 373 spaces forming 3879 total relationships. Figure 1 plots the degree distribution[1] of people and lists. It is clear that, as we expect, they both follow a power law distribution [1], the parameters of which are also plotted.

The graph contains 8 disconnected components, the largest of which contains most nodes (the others containing less then 4 nodes each.) Figure 5, in the

[1] The degree of a node is the number of other nodes it is connected to.

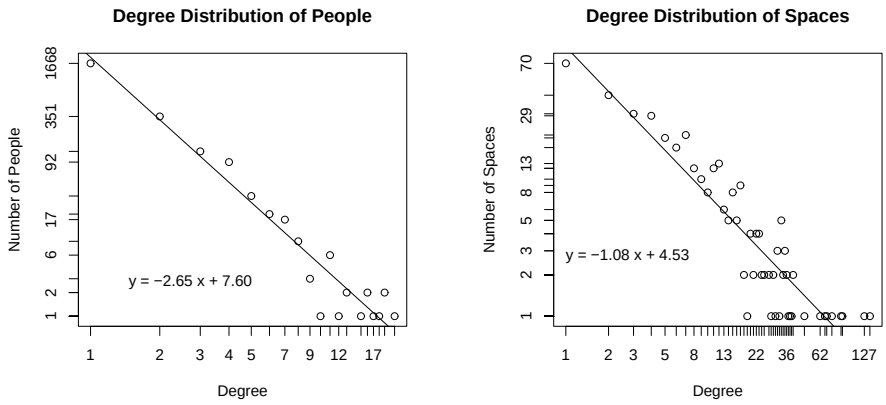

Fig. 1. Degree distribution of Spaces and People, plotted with a log-log scale. The good fit of the straight line indicates that the degree distributions follow a power law. The parameters of the regression are also plotted.

appendix, represents a stylised illustration of the largest component. Red nodes denote spaces and blue nodes denote people. The intensity of the links represents the strength of a persons link towards a space relative to the other spaces he or she belongs to. Many weak ties were removed for clarity, breaking this main component further down, and only people acting as bridges between spaces are plotted.

3 Effectiveness of Partial Surveillance

Our first experiment attempts to measure the effectiveness of partial surveillance. In our model we assume that if one member of a space is under surveillance all relationships associated with that space become known to an observer. Monitoring one participant in a newsgroups, or private mailing lists, or bugging one participant into a public or private meeting room, has such a property. Using the phone records of one individual for investigation, also leaks information about others, which also has a similar property. The adversary's goal is to observe a small set of people, according to its surveillance budget, that maximises the information the adversary gained in terms of uncovering people, clubs and relations.

The initial target selection strategy is extremely simple: the adversary chooses to put under surveillance those people with highest degree, excluding all spaces already under surveillance. The algorithm to do this requires ranking candidate people according to the number of spaces they participate in that are not under surveillance, and picking the first one. The process is repeated as many times as the surveillance budget allows.

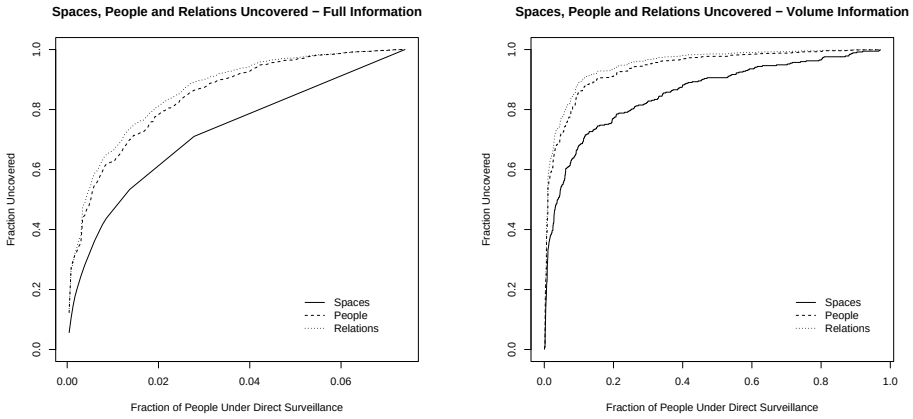

Fig. 2. Fraction of spaces, nodes, and relations uncovered by the eavesdropper by observing a certain number of nodes. On the left the targets for surveillance are chosen using full information (upper bound), on the right the nodes with highest traffic volumes are chosen first. (Note the different scales on the x axis.)

The results of the first experiment are plotted in figure 2. The left hand side of figure 2 illustrates the results of this strategy. The horizontal axis denotes the fraction of people under surveillance, chosen according to the degree strategy described above. The vertical axis denotes the fraction of people, spaces and relations that were uncovered. Note that after a very small fraction (0.08 or 8%) of the network under targeted surveillance the full network information is already available to the eavesdropper.

Our strategy is particularly well tuned for uncovering spaces, and therefore we do no expect other strategies to provide significantly better surveillance results than those depicted by the solid black line (i.e. other strategies would have a graph that is bound upwards by our graph's strategy.) Minor modifications could lead to similarly good strategies for uncovering people and relations, but in practice their graphs in the left hand Figure 2 are a good approximation. Note that the graph denoting the worse surveillance strategy, or the best defensive counter-surveillance measure, would be a straight line from $(0, 0)$ to $(1, 1)$. No information is leaked when there is no surveillance – all information is leaked only when everything is under surveillance. Any amount of surveillance only provides a proportionate amount of intelligence.

3.1 Target Selection with Partial Information

The target selection strategy based on the degree of nodes is optimal, yet difficult to use in practice, hence only useful as an upper bound. Employing it would

require the adversary to already know all the relationships between people and clubs. If that was the case there would be no need for surveillance in the first place.

We present how the fruits of partial surveillance lessen if the target selection has to be done with less information. We consider that the adversary can only observe the aggregated mass of relations of each user. In terms of our model, an adversary can only observe the sum of the weights of all the relations of a node to decide if it is to be put under surveillance. This is a good model of the degree of protection offered by current anonymous communication infrastructures [3, 8, 4, 2]: they hide the correspondents, but not the communication acts. As a result an adversary can estimate the volume of messages sent by Alice, but not her degree or correspondents.

Degrading the adversary's ability to perform target selection lowers the intelligence outputs for the same surveillance budget. The right hand side figure 2 plots the fraction of information uncovered by the adversary against the fraction of nodes under surveillance, chosen according to their relation weights. The full network information becomes available only after nearly all nodes are put under surveillance (note the horizontal axis scale difference between the two plots.) The degradation is sharp: to get 80% of relations an adversary with full information needs to observe about 3% of nodes versus 30% of nodes when only weights are known. Yet the two graphs are comparable early on, since 50% of all relations become known extremely quickly, with less than 5% of nodes being put under surveillance in both cases.

4 Discussion

The two models of target selection examined are of quite some practical importance. The first, where full network information is known to the adversary, can be seen to represent target selection in networks that do not provide any anonymity protection. In those, the adversary can easily obtain the traffic data necessary to chose appropriate nodes, and minimize the number of those to be put under surveillance to obtain most information in the network.

The second strategy, which relies only on aggregate volume information being available, represents the information available to an adversary in an anonymous communication network, without cover traffic. In such a system the exact relationships cannot be observed, but the overall volume of data sent by each participant can be inferred. An adversary in such an anonymous network has to chose a small number of participants and put them under more intensive surveillance (i.e. using listening devices, back doors on computers, ...)

Our results demonstrate that the use of anonymized communications makes target selection harder, and reduces the value of the intelligence extracted from surveillance. At the same time, the protection they offer is far from perfect: 50% of the nodes will have some of their relationships uncovered after minimal closer surveillance of about 5% of the nodes. This result falls short of the established perception of security such networks should provide.

On the other hand, the perception that anonymous communications are good at hiding small groupings of people, that take care not to 'stick out' is true. In the case the adversary has only volume data, small spaces keep getting discovered as more people are put under surveillance. Some remain hidden until the surveillance operation includes close to every participant.

We can transpose this observation into the real world by saying that a large surveillance budget will be necessary to eliminate every single cell of a decentralized terrorist network, ultimately with unsure results. Cells of a few people will be active and undetected until nearly everyone is put under surveillance.

4.1 The Diminishing Returns of Surveillance

Our results also show that surveillance as an activity has diminishing returns. A great amount of intelligence is produced for the first few units of intelligence budget invested, but less and less intelligence is produced as this budget is increased. In case imperfect information is available for target selection, getting close to 100% intelligence, would require observing nearly everyone, with extremely high marginal costs after the first 20% of people are already under surveillance.

This has to be judged in terms of the opportunity cost of such activity, and the marginal value of the intelligence collected at each stage. In case a social network is using surveillance countermeasures, such as anonymous communications, or modifying its structure to keep the degree or volume of important nodes low, it might be extremely expensive until Eve can observe anything useful. That can be best illustrated by the example of a cell structured terrorist network that embeds itself into a larger community. A large amount of the surveillance budget would be 'wasted' by putting under surveillance high degree nodes, that are not associated with a terrorist cell, and only after a large part of the community is under surveillance and intelligence concerning the network would be collected.

This has important policy implications, in that it illustrates that the degree of collateral privacy violation, for the amount of actual useful intelligence will be quite substantial. Furthermore, the useless, from the point of view of intelligence collection, privacy violation will happen early and cheaply (lets be reminded that 80% of participants are observed when only 20% are under surveillance). On the other hand, the hardened cell structures will not be observable until much later, and at a much greater cost per unit of intelligence return.

4.2 Reinterpreting Interception Figures

We will attempt to use some results from our model and experiments to reinterpret some levels of reported surveillance and assess the true levels of privacy invasion taking place in each case.

The United Kingdom Interception of Communications Commissioner, is under statutory duty to report the aggregate numbers of surveillance warrants issued under the Regulation of Investigatory Powers Act 2000 [11]. These figures [13] include the number of warrants for phone tapping and opening letters. In his report he states that 1849 such warrants were issued during the period

from 1 January 2004 to 31 December 2004 by the Home secretary (124 further warrants were issued by the Scottish Executive). According to RIPA those warrants do not have to be attached to a phone number or other network address, but can be attached to a person allowing law enforcement surveillance of any medium he or she uses.

Our models suggest that the figure of 1849 warrants gives a lower bound on the actual number of people under some kind of surveillance in the UK during that period. This would assume that these 1849 people form a disconnected clique, that contains all those that are of interest to law enforcement. This is counter intuitive and seems a gross underestimate when compared with the crime rates, the estimated criminal population of the United Kingdom, and the prison population. Yet, it may have been a realistic assumption if there were assurances that interception was only used to clarify homicides, whose level is very low in the UK (the 1998 crime statistics [10] suggest there were 1227 homicides in the UK in that year.) We know though that surveillance can take place to clarify other crimes, and therefore we can safely say that this lower bound would be much lower than the actual number of those that surveillance affects. Furthermore homicide suspects definitely do not form a clique (with about 37% of homicide prime suspects being in the family or lovers of the victims [10].)

Assuming that the graphs in figure 2 hold true for all social networks (a thesis that requires much further data collection and research to prove or disprove beyond doubt) we can calculate an estimate of the actual number of those under surveillance. In this calculation we shall use the graphs resulting from the unprotected surveillance: the current communication technologies, such as POTS, GSM and most Internet based technologies do allow for complex target selection based on traffic data [17]. We approximate the region of the graph between $x = 0$ and $x = 0.01$ by the line $y = \frac{0.5}{0.01}x$, and observe that the fraction of those under surveillance against the total population of England and Wales (to which the interception figures correspond) is well into this range ($1849/45000000 \approx 0.00004$). We can therefore estimate the number of those likely to be under surveillance as:

$$\text{Estimate of those concerned} = \frac{0.5}{0.01} \cdot 1849 \approx 92000 \text{ people} \tag{1}$$

This makes the unrealistic assumption that all those targeted are distributed equally amongst the UK population. That in turns leads to an error and therefore this number must be seen as an order of magnitude, of those that are in fact affected by the surveillance of the, relatively few 1849 people named in the warrants. For each warrant issued we could say that about 50 people had information collected about them, which is consistent with our intuition about the size of the close social circle of the person under surveillance.

A second example we shall present makes use of the data presented by the Greek government concerning an illegal wiretapping operation that was uncovered in February 2006. At the time the operation was uncovered, there were 103 numbers under surveillance, mostly belonging to government politicians and ministers, including the Prime Minister himself, the heads of the police, army and secret service, lawyers, journalists and a handful of left wing political ac-

tivists. Since the list has become public, we know that those targeted would indeed have very high degree, and do not just form a clique. Applying the same calculation as above we can say that in this case the actual number of those affected by the surveillance (i.e. who had private information concerning them being discussed in intercepted calls) would be closer to:

$$\text{Estimate of those concerned} = \frac{0.5}{0.01} \cdot 103 \approx 5150 \text{ people} \tag{2}$$

This is a gross underestimate, since the network we studied was based on a decentralized political network, while most of those under surveillance in Greece were very much part of a hierarchical government structure. As a result we expect their degree to be much higher than the degrees we observe in our data, leading to more people being affected by the surveillance. Furthermore, we expect those third parties affected by the surveillance to also have a special status within the command and control structure (since they were talking to ministers after all).

4.3 Reinterpreting Physical Surveillance Figures and Tactics

Our models of surveillance are sensitive to the actual topology of people and clubs, as well as surveillance uncovering all people in a space, but are not necessarily bound to technology. The next example we shall use is from physical surveillance of political events in London, UK.

The London Metropolitan Police makes extensive use of surveillance at political events in London. This includes photography, video, CCTV and the use of spotters (ofter referred to as 'the forward intelligence team' or FIT), used in a blanket fashion at events, such as public gathering, meetings and demonstrations.

Mark Thomas, a UK based journalist and comedian, found a discarded sheet of a table of faces, used by the police to remember those to be kept under surveillance during a large demonstration in London. He writes [15]:

> This mugshot from the CO11 Public Order Intelligence Unit is part of a "spotter card", designed to help police identify troublesome protesters at the biennial arms-fest before they can commit any nefarious acts. Each spotter card has pictures of 24 individuals, lettered A to X. I am, as you see, Suspect H.

Other reports also indicate that the surveillance surrounding this particular event was not confined to the demonstrations themselves, but also the organizing meetings [5]:

> Following the decision to have regional meetings, London met last week. This meeting was attended by four cops and two photographers who stood and took photos inside the Festival Hall. We are not therefore going to publish further details of meetings but if you want to attend please email [...]

Further reports show that police often follows those under directed surveillance and records their interactions with others, as well as performing blanked stop-and-searches before allowing anyone to leave a demonstration.

These reports show that our model would be appropriate to understand the extent of surveillance actually taking place. A small set of 24 people are set as the "targets" of surveillance. Any space where a target is present becomes a legitimate space for surveillance, and teams of police and photographers record everyone present there.

We conjecture again that the ratio of those under surveillance to those who have actually their privacy violated, and intelligence gathered on them, follows the data we have collected. We can estimate some bounds on the number of people that are effectively under surveillance. There were thousands of people at the demonstrations where the spotter card was uncovered, yet there were about 24 faces on the card. This is less than 1% and we shall use our previous calculations to estimate the number of people who had information and intelligence recorded about them:

$$\text{Estimate of those concerned} = \frac{0.5}{0.01} \cdot 24 \approx 1200 \text{ people} \tag{3}$$

The calculation above assumes that the target selection uses a near optimal algorithm, choosing targets according to their degree distribution. This is confirmed by looking at whose face is on the spotter card. It contained mostly people involved in groups directly organizing the infrastructure of the demonstration (outreach, coordination, etc.).

4.4 The Failure of Adaptive Target Selection

We have presented so far two types of target selection. The first, where the adversary can use full information, is unrealistic and can be used as an upper bound. Similarly, one would expect to see the target selection strategy that does not take into account the degree but only the volume of communications to act as a lower bound for the effectiveness of selecting nodes for intelligence gathering. One might think that an adaptive strategy, that selects nodes in accordance with intelligence gathered thus far, would easily outperform target selection based on traffic volume alone.

This intuition turns out not to be correct, since we have shown that two simple adaptive strategies *are inferior* to performing target selection based on mere volume, for the purposes of designating nodes to put under surveillance. Our results are summarized in figure 3.

The first adaptive strategy, prioritizes nodes with *high known degree*. The intuition behind this strategy is that those nodes are likely to also have links with spaces that have not yet been observed. Observing them would therefore uncover those spaces and their participants.

The second adaptive strategy, uses *structural equivalence and uniqueness* [6] to select targets. A node is selected if its known degree is high, but also if there are fewer other known nodes that share its position in the network. This allows for

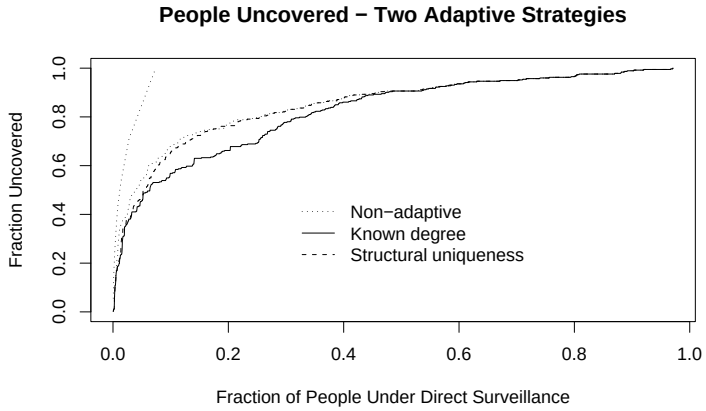

Fig. 3. Comparing adaptive target selection, with selection by volume or full information. The thinnest two lines reproduce the results of figure 2, i.e the percentage of nodes that need to be under surveillance to uncover people in the social network. The other two lines represent the simple minded adaptive strategies.

nodes that act as unique bridges between disjoint spaces to be selected with high priority, based on the assumption that they may also be acting as bridges with unknown spaces. This structural uniqueness strategy outperforms the simple minded known degree strategy, but they are both worse in selecting targets than the simple minded volume based target selection.

This preliminary result has profound repercussions on anonymous communication system design. Not only the information leaked by such systems, as they are deployed today [3, 8, 4, 2], does not hide the volume information necessary to do target selection, but also target selection using the volume information leaked is better than using simple minded approaches based on partial structural information. This means that anonymous communication systems should prioritize protecting volume information, rather than communication relationships, which would require a radical shift in architecture.

4.5 Target Selection for Disruption

We complete our assessment of the use of anonymizing networks to protect users against target selection, by considering an adversary that intends to disrupt the network (so far the objective was just to observe). Such an adversary selects people (not spaces) according to a target selection strategy, and removes them, with the aim to split the network into smaller components. A more sophisticated adversary may aim to block flows of information, and lessen routing efficiency, but we shall only consider here the simple case.

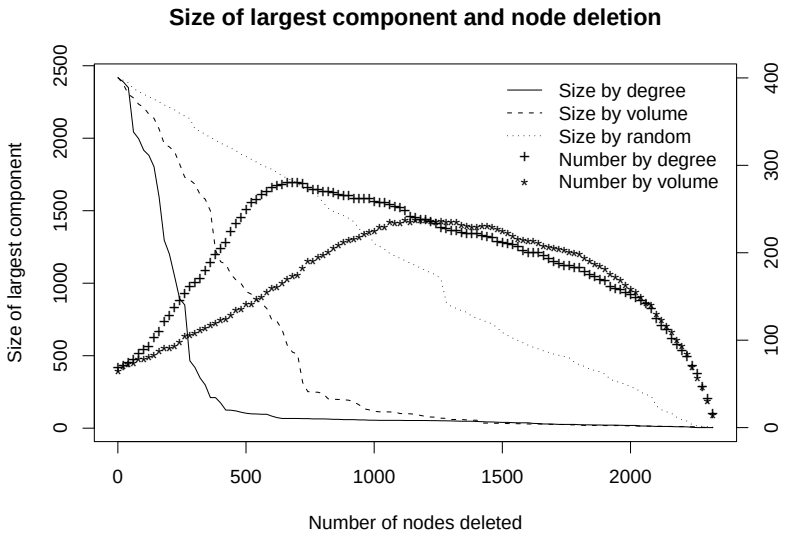

Fig. 4. The size of the largest graph component (lines – left axis), and the number of components (points – right axis), as a function of nodes killed according to the two strategies.

We apply the same target selection strategies to this problem as to the previous ones. The first strategy assumes that the adversary has full information about the network, and removes the people with highest degree (the *by degree* strategy). On the second case, the adversary only has volume information (possibly because an anonymizing network is used) and has to make choices based only on this. Such an adversary chooses people to remove according to the volume they transmit (we call this strategy *by volume*).

Figure 4 summarizes the results of our simulations. The lines plot the size of the largest connected component in the network, for both strategies (left axis) and when nodes are removed at random. The dots represent the number of disjoint components (right axis). We observe that for both strategies there is a sharp drop in the size of the largest connected component as the adversary removes people (we do not allow the adversary to remove spaces). The effect of partial information, in the form of volume information only, on the target selection is not as dramatic as one may expect. The adversary is required to remove twice as many nodes to reduce the largest component to the same size as in the case the target selection was performed with full information.

This is further evidence that current anonymizing networks are doing a poor job at protecting against target selection. They leak volume information that gives the adversary a considerable advantage in comparison to removing nodes at random (also plotted in figure 4). If anonymous communication networks were

perfect we would expect the adversary to do no better than using a random strategy.

5 Conclusions

In this work we present to the world of computer security a novel approach to studying surveillance. Instead of looking at the confidentiality between Alice and Bob, we take an approach based on social networks, consider that information leaks through third parties, and assess security on this basis. This is an important way of re-framing the problem, that should be considered seriously by engineers who want to build secure *systems* rather than just *channels*.

Our data set is comparatively large and represents a realistic set of interaction in a decentralized worldwide network between people and spaces over a long period of time. This data set has been used for the core of our analysis, but is far from being exhausted. In particular, we only considered aggregate traffic volumes and never took into account the exact days that actions were taking place between people and spaces. Mining such data may lead the adversary to even better traffic selection strategies. This is an open research problem, and our (anonymized) data and software tools, used for its analysis, are available free for all to use.

We also attained some concrete results, with both technological and policy consequences. First, we show that there exists a very small set of people, who, if put under surveillance, will leak all relationships in the network. Such a set is easy to find if topological information is available, but as expected, becomes more difficult to find if only volume data is available. The later case corresponds to the use of an anonymizing network, which still leaks enough volume information to allow for good target selection for surveillance (more than 50% of relations to be uncovered after 10% under surveillance). The volume information leaked is so informational to perform target selection that trivial adaptive strategies, based on highest degree and structural uniqueness, of those under surveillance do not beat it. When looking at target selection for disruption, by an adversary that tries to fragment the network, the results are similar. An adversary that bases its target selection on volume data only, performs worse (by a factor of 2 between effort and results), but is still much more effective than an adversary that removes nodes at random.

The fact that volume information can be used so effectively to perform target selection should be a warning call to all those designing anonymous communication systems. Our results show that mere *unlinkability* is not sufficient to foil target selection, and full *unobservability* is necessary to hide all information, including the volume of interactions. Such system would require a lot of cover traffic and delays, and are likely to be very expensive.

Finally, we used our results to conjecture that the volume of actual surveillance taking place exceeds by far the published and publicized numbers. We have used examples form lawful and unlawful telecommunication surveillance as well as an example from physical surveillance to show that behind low figures can

hide a large surveillance operation. From a policy point of view it might be easy to justify the close surveillance of a hard core of 24 political activists in London. It is much more difficult to justify the privacy invasion of about 1200 people with whom they interact, that also find themselves being under surveillance by proxy.

Similarly, the debate about traffic data retention can be illuminated by our findings. The availability of such data makes the job of target selection trivial allowing the selection of a handful of targets that are exactly within the 8% of the people that would provide most information, allowing the optimal invasion of everyone elses privacy.

While such invasion of privacy of most people is easy and cheap, uncovering small degree hardened disconnected sub-networks remains hard, and takes substantial surveillance effort. As it is often the case the innocent, who do not protect themselves, are the first victims of surveillance, while hardened targets remain elusive for a long time.

Acknowledgments

Ross J. Anderson motivated this study by suggesting in many talks and discussions the idea of building secure, infiltration resistant, peer-to-peer systems based on 'clubs'. George Danezis is supported by the FWO (Fund for Scientific Research – Flanders). This research was done in the context of the E.U. integrated project PRIME.

References

1. Lada A. Adamic. Zipf, power-laws, and pareto – a ranking tutorial. *Glottometrics*, (3):143–150, 2002.
2. Oliver Berthold, Hannes Federrath, and Stefan Köpsell. Web MIXes: A system for anonymous and unobservable Internet access. In H. Federrath, editor, *Designing Privacy Enhancing Technologies*, volume 2009 of *LNCS*, pages 115–129. Springer-Verlag, July 2000.
3. George Danezis, Roger Dingledine, and Nick Mathewson. Mixminion: Design of a Type III Anonymous Remailer Protocol. In *IEEE Symposium on Security and Privacy*, Berkeley, CA, 11-14 May 2003.
4. Roger Dingledine, Nick Mathewson, and Paul Syverson. Tor: The second-generation onion router. In *Proceedings of the 13th USENIX Security Symposium*, August 2004.
5. London DISARM DSEi. Police harassment of dsei meetings and future meetings. UK Independent Media Centre, February 28 2004.
6. Peter Klerks. The network paradigm applied to criminal organisations. In *Connections 24(3)*, 2001.
7. Alfred J. Menezes, Paul C. Van Oorschot, and Scott A. Vanstone. *Handbook of Applied Cryptography*. CRC Press, 1996. ISBN: 0-8493-8523-7.
8. U. Moeller, L. Cottrell, P. Palfrader, and L. Sassaman. Mixmaster protocol version 2. Technical report, Network Working Group, May 25 2004. Internet-Draft.

9. Shishir Nagaraja and Ross J. Anderson. The topology of covert conflict. Technical report, University of Cambridge, Computer laboratory, 2005.
10. Home Office. Statistics on race and the criminal justice system. Home Office Press, 1998.
11. UK Parliament. Regulation of investigatory powers act 2000. TSO (The Stationary office), ISBN 0 10 542300 9, 2000.
12. Ron L. Rivest, A. Shamir, and Leonard M. Adleman. A method for obtaining digital signatures and public-key cryptosystems. *Communications of the ACM*, 21(2):120–126, February 1978.
13. The Rt Hon sir Swindon Thomas. Report of the interception of communications commissioner for 2004. TSO (The Stationary office), November 2005.
14. sKathleen M. Carley, Ju-Sung Lee, and David Krackhardt. Destabilizing networks. In *Connections 24(3)*, 2001.
15. Mark Thomas. Selling torture in london's docklands. The New Statesman, September 26 2005.
16. Stanley Wasserman, Katherine Faust, Dawn Iacobucci, and Mark Granovetter. *Social Network Analysis : Methods and Applications (Structural Analysis in the Social Sciences)*. Cambridge University Press, 1994.
17. John Young and Erich M. On obtaining "lawful interception" documents. `http://www.quintessenz.org/etsi`.

A Additional Figures

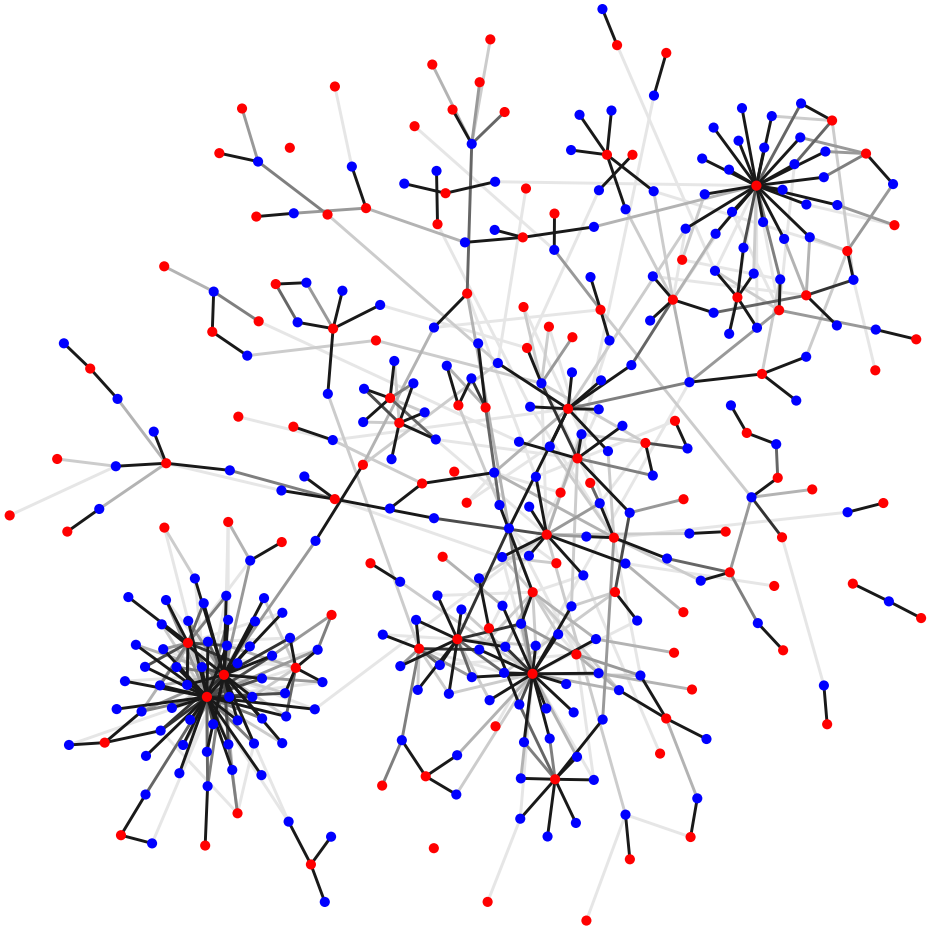

Fig. 5. The largest component of the network studied. Red nodes denote spaces, and blue nodes those people that act as bridges between spaces. The intensity of the link indicates the strength of the association between spaces and people.