

Why I’m not an Entropist

Paul Syverson*

Naval Research Laboratory
syverson@itd.nrl.navy.mil

Abstract. What does it mean to be anonymous in network communications? Our central thesis is that both the theoretical literature and the deployed systems have gotten the answer essentially wrong. The answers have been wrong because they apply the wrong metric to the wrong adversary model. I indicate problems in the established adversary models and metrics for anonymity as well as implications for the design and analysis of anonymous communication systems.

1 Introduction

Anonymous communication is the quintessence of brief encounter. For anyone to whom it is anonymous, such communication cannot be linked to an individual, for example to its source or destination. But it also cannot be linked to any other instances of communication. If it could, it would be part of a communication profile, hence pseudonymous communication rather than anonymous. Anonymity thus guarantees that an encounter is brief rather than leaving this to chance.

Anonymity is also an area of security that is much younger than, for example, confidentiality. Mechanisms and networks for anonymous communication have been designed for only about thirty years [6] and deployed for not quite half that time [21, 35, 19]. We do not have nearly as much experience with deployed systems on which to build theory, definitions, and models as many other areas of security, for example confidentiality or authentication. Though the field has developed and adapted, the basic conception of what it means to be anonymous has persisted throughout its history.

That conception is of anonymity as about indistinguishability within a set. The idea is thus that there is a set of possible senders, and the adversary cannot adequately distinguish amongst the members of that set. (For brevity and convenience I focus discussion on sender anonymity of a message.) I call this *the entropist conception* of anonymous communication. I am motivated in choice of terminology here by general Rényi entropy rather than Shannon entropy specifically (H_2 in the ordering of Rényi entropies). Thus I am including everything from simple cardinality of the set of possible senders (the log of which is Hartley entropy, H_0) to the other extreme of min entropy (H_∞). However, I am more

* A version of this paper was originally presented to the 17th Security Protocols Workshop, Cambridge UK, April 2009, which had as its theme “brief encounters”. This is a substantially revised and expanded version of the originally presented paper.

focused on the conception of anonymity as indistinguishability within a set (possibly according to some probability distribution) than on any particular mathematical characterization of it. I will argue that this basic conception—a conception that underlies virtually all extant work on anonymous communication—is not appropriate for the fundamental role it has been given.

Before getting into details, I set out some general principles for security metrics and system design that I believe should be uncontroversial. It is on the basis of these principles that I will make my case.

- To be useful, security metrics should reflect the difficulty an adversary has in overcoming them. (So, for example, a metric for a security property should not depend on variables whose values do not significantly impact whether the adversary succeeds in attacks on that property.)
- To be meaningful, security metrics should not depend on the values of variables for which we cannot make adequate relevant determinations or predictions. (It might be sufficient to determine or predict bounds, distributions, etc. for those values rather than the exact values themselves, provided they are relevant. For example, to observe that a value will always be finite will usually not reflect a bound that is relevant to a practical security metric.)
- Though security through obscurity is rarely a good system design strategy, if your design provides your adversary with an explicit target that he has available resources to overcome, you are doomed to fail.

It is crucial to understand that our fundamental complaint against entropism is *not* that entropy fails to ever properly say anything about the uncertainty we may have about senders or receivers of messages. The problem is that entropism fails to tell us much about the adversary’s amount of knowledge (or lack of knowledge) that is useful in saying how secure our practical anonymity systems are, or how secure a given use of one of them is. Starting from an entropist conception of anonymity has led to a focus on system properties that are not the most important to system security, has led to system assumptions that are not reasonable in practice, has led to adversary models that are not reasonable in practice, and has led to both theory and system design for anonymous communication that fail in significant ways. The remainder of the paper investigates these points in light of the above principles. We concentrate herein on problems with the existing approach, but we will at least briefly consider potential alternatives to replace it with.

2 Wait. I thought you said you’re *not* an entropist.

There are many circumstances where the entropist conception of anonymity makes perfect sense, for example, in an election where there are registered voters casting ballots. An entropist approach can protect the anonymity of voters by confounding of associating voters with ballots and provides a good metric for that protection. (This assumes that appropriate other protections ensure that each registration represents exactly one voter, that each voter casts at most one

ballot per election, and assumes that enough eligible voters vote, that enough voters cast ballots for each of the options—if that is of concern, etc.)

As a more exotic example, in the 1970s the United States and the Soviet Union sought to limit the number of nuclear missiles they both had through cooperation and inspection while maintaining security guarantees. A large part of one of the major initiatives involved shuttling Minuteman missiles about in a field of silos to preclude successful first-strike attacks directly on all available missiles. The plan also included techniques for communicating to the adversary an authenticated, integrity-protected report from each silo. The report indicated whether a missile was in the silo or not, but without the adversary being able to determine which silo the report was from (except by silo-unique but random identifier). In a field of 1000 missile silos, the adversary could be sure that exactly 100 would be occupied but would not know which ones were occupied. Note that this anonymity system used “dummy packets” in the form of huge “transportainer” trucks continually shuttling either actual missiles or dummy loads between silos. Talk about system overhead! See [40] for more details about the plan.

As a less exotic example, suppose an intimidating colleague or perhaps the boss is repeatedly but obviously doing something that is inappropriate and bothers several coworkers: leaving the food area or bathroom a mess, telling jokes or making comments that some find offensive, etc. Nobody wants to be the one to complain to the offender. Also they don’t want to embarrass or perhaps encourage ire by giving the impression that this was a topic of full group discussion and decision making, whether or not that is true. They just want it to stop. Thus they wouldn’t want to use a collaborative anonymity mechanism to send this message unless it was also being used to send others. Put differently, they want the recipient to know that one of them sent the message, rather than that all of them sent it, even though only one copy was delivered.

Finally, a similar example that comes closer to our intended topic below. Suppose it is important to release some information to the press or the public without attribution. The information is posted to a public site. It is known that only a relatively small number of people had prior access to the information, so it was one of them, but we don’t know which one. (Let us assume and ignore that the relevant set of people is known, and that there is no doubt that they are the only ones who had the information. Also assume that the authenticity of the information is evident or confirmed once revealed.) It could be government information (good or bad) of important public interest where overt indication of which knowledgeable person released the information would be a distraction from the public good of the release. Alternatively, perhaps the release is intentional and important but cannot be overtly approved for some reason, and if the “leaker” were known, s/he would be punished, officially or otherwise.

Entropism is about using entropy as the meaning, the criterion for anonymity or how anonymous something is. That has drawbacks because, as an average, it necessarily does not capture everything important. To the extent that entropy is the measure of anonymity it drives system design to maximize entropy, which

may not be the same thing as maximizing anonymity in ways that matter. It is just one number, or a few numbers, such as the Shannon entropy and the min entropy. Any of these might be fine as a measure of anonymity in an appropriate context. None is fine as *the* measure of anonymity. As *the* measure it effectively becomes the definition for a relatively complex concept which is no better served by limiting to any single definition than is the overall field security by limiting to a single definition of ‘security’.

In all of the just given cases, there is a set of possible sources (or locations) and the relevant security issue is the uncertainty within that set. For some of the cases, the exact set may not be known, but reasonable bounds can be given for both set size and uncertainty. Uncertainty on a known or approximately known set is clearly the significant concern in these cases, but what are examples where entropy is not appropriate? We will argue that entropism is the wrong approach to design and evaluation of the most well known and widely used systems for protecting the anonymity of communication.

3 Protecting Anonymity

Tor is a very widely used and deployed onion-routing network, which means that it gets its protection from creating cryptographic circuits along routes that an adversary is unlikely to observe or control. As in previous onion-routing systems, in the currently deployed Tor design this is achieved by choosing unpredictable routes through the network. This could, however, also derive in part from the inherent difficulty an adversary has attacking a particular part of the network [25, 26] or from other sources. We will return to this below. Also like other onion-routing systems, Tor passes traffic bidirectionally along those circuits with minimal latency [44]. It is common in the literature to characterize the anonymity protection Tor provides in entropist terms. For example, “[t]he anonymity provided by Tor relies on the size of the anonymity set. Currently, there are around 1500 ORs, and an estimated quarter million Tor users.” [43].

Assume I am a user of the network who picks his entry node the same way the vast majority of users do and then visits a hostile website. On the entropist conception, my anonymity is the same as any other user of the network. With about a quarter million active users that gives me a (Shannon) entropy of about 18 with respect to being identified as going to this website. But the number of users is only an estimate, and the system is designed so that nobody should be able to actually see connections coming from all of them. Further, talking about the number of users ignores how attacks actually happen.

An adversary observing an entry node and an exit node of the Tor network through which I am, e.g., browsing the web can trivially link the two ends of the connection and correlate source to destination. This has been an acknowledged feature of the design since its inception [14]. By the same timing correlation, an adversary that controls or observes a destination website and the entry node will be able to identify the source address of the user accessing the website. Against these attacks the number of other users is all but irrelevant.

In this case, entropy fails our second criterion for a security metric. We cannot determine the number of users on the network in general or the number using circuits with the same entry and exit nodes at the time of the attack. We can estimate these numbers based on network measurements such as given above. That would, however, only tell us the average number of users during various periods in the past. What should matter for my (entropist) anonymity with respect to the adversary is my distinguishability from other users during the attack. If we cannot know when the attack occurs, could we give at least a lower bound on the anonymity set size during any reasonable candidate attack period? The answer is no. This means that we cannot use set size or indistinguishability within a set as a metric of my security when using the system. Could it be because nobody knows the anonymity set size so that my protection is absolute?

No. Entropy also fails our first criterion for a useful security metric: it should reflect the difficulty an adversary has in overcoming security and should not depend on variables whose values do not significantly impact the adversary's success. We could imagine rare cases where two Tor circuits through the same entry and exit nodes might be hard to distinguish by at least a completely passive adversary, but it is well known and accepted that whether you are the only user of those nodes during the attack or there are a hundred others, linking your source and destination is trivial. So the anonymity set size cannot be determined, and the difficulty of the most salient attacks are not affected by what it might be.

Perhaps we are discussing uncertainty on the wrong set. Perhaps the number of users has only an indirect effect on the security of Tor traffic. We have repeatedly observed that the primary threat we are discussing is the end-to-end timing correlation. The relevant set is thus the number of locations at which traffic might enter and exit the network.

In the scenario in which I go to a hostile website, what will matter is the nodes through which I can enter the system. If I make that choice the same way as everybody else, then I currently have on the order of three hundred choices.¹ So an entropy of a little over eight. But that assumes that the probability of entry node choice is uniform. This is not actually true. It is weighted by node bandwidth and other factors that we will not go into. We could still calculate the entropy of a typical connection entering the network. But even if we ignore that and assume the choice of entry nodes was uniform we must ignore something else as well—how likely a node is to be compromised. Suppose that, instead of choosing uniformly (or possibly weighted to help maintain performance), I always choose my entry nodes to be just one of a handful of nodes. If the nodes in that handful are all run by highly trusted and technically competent friends who would never betray me, my anonymity is clearly better than picking uniformly from the much larger set: the hostile website will never identify me through the

¹ For those who know about entry guards, we are describing the choice of entry guards not the subsequent choice of route once guards are chosen. For those who do not know about entry guards, please do not be distracted by this footnote.

cooperation of a hostile entry node.² My anonymity protection has far more to do with the security of that first node I choose in my path than with the number of other users or even other possible first nodes. Again, this is a variable that does not affect the difficulty of attacking my anonymity.

Using entropy as a security metric for anonymity conflates the adversary’s uncertainty with the effort needed to reduce that uncertainty. This conflation is fine as long as the effort needed to reduce uncertainty is more or less uniform in the degree of uncertainty (as would be the case for, e.g., key length for a well designed block cipher). When it is not, failing to focus on anonymity protection distorts both security analysis and secure system design.³

Even within the entropist conception something akin to the distinction between entropistic anonymity and anonymity protection is recognized. In [37], Pfitzmann and Hansen distinguish between anonymity and strength of anonymity: “Robustness of anonymity characterizes how stable the quantity of anonymity is against changes in the particular setting, e.g., a stronger attacker or different probability distributions. We might use quality of anonymity as a term comprising both quantity and robustness of anonymity. To keep this text as simple as possible, we will mainly discuss the quantity of anonymity in the following, using the wording ‘strength of anonymity’.” But they still make entropy the primary measure of how well anonymity is protected. “All other things being equal, global anonymity is the stronger, the larger the respective anonymity set is and the more evenly distributed the sending or receiving, respectively, of the subjects within that set is.” They do note how individuals may have weak anonymity regardless of the average uncertainty across the anonymity set. “Even if global anonymity is strong, one (or a few) individual subjects might be quite likely, so their anonymity is weak. W.r.t. these ‘likely suspects’, nothing is changed if the anonymity set is made larger and sending and receiving of the other subjects are, e.g., distributed evenly. That way, arbitrarily strong global anonymity can be achieved without doing anything for the ‘likely suspects’ [7].” This distinction in [7], however, is intended to cope with the influence of statistical outliers on average uncertainty within a set versus uncertainty concerning an individual and thus to suggest using quantiles as a metric for those cases. It does not at all question the basic idea that anonymity is based on a level of uncertainty from within a known set.

To reiterate, the problem with the entropist conception is not that entropy entirely fails to reflect uncertainty (once the situation is properly described). The problem is that focusing on the size of a known set of users and their distinguishability within it skews system design as well as what gets analyzed in

² Whether or not these nodes are known to be affiliated with me now becomes relevant, which depends on the probability that the second node in my path is compromised (and ignoring links as possible places of compromise) [34]. We will touch briefly on such issues below.

³ For a cryptographic analogue see the discussion of locally unpredictable delaying functions in [20] wherein most of the bits of a function might be easy to compute, but a few of them require much more effort.

a way that obscures rather than enhances our understanding of the anonymity protection provided by our systems. Wanting to communicate anonymously and so making entropy the focus of your anonymity design is like wanting to be married and so making expected-number-of-spouses the focus of your matrimony design. If that is indeed your primary focus, you are likely to end up with a strategy of, e.g., proposing indiscriminately to as many people as possible.⁴

If probability of node compromise could be assumed to be roughly uniform (and ignoring links), then network size would be a primary determinant of anonymity protection. Similarly, if we could produce bigger, more uniform looking (to the adversary) sets of senders and recipients, that might actually be useful. But that is not a realistic view of how large, widely used anonymity networks work. They are comprised of fairly dynamic and diverse collections of users communicating over nodes that are diversely trusted by diverse parties and that are diversely configured versions of diverse platforms. And these nodes are connected over diversely trusted links (based on ASes, ISPs, geography, etc.). Unlike designing a closed secure system, there is no point in even discussing trying to make the degree of security of all of the different parts of the network roughly comparable. We will touch on more centrally operated systems below. We will also see that the number of nodes in an open system does play a role in the protection a system offers, but only in combination with the resistance those various nodes have against adversaries.

In security analysis, most adversaries are worst-case, possibly subject to some constraints on ability. But they can attack any message anywhere. ‘Attack’ might just mean observe, but the adversary can be in the worst possible place. If there is a subset of nodes, or messages that he is able to attack, it is assumed that these can be the optimally most effective nodes or messages. We will next consider what an adversary can do and where.

4 What is an adversary?

“Keep your friends close, and your enemies closer.”

Sun-tzu — Chinese general and military strategist (c. 400 BCE)

Computer security, including anonymity, is fundamentally tied up with the idea of an adversary. Indeed the primary difference between security and other types of computer assurance is the assumption in security that there is a potential for intentional misuse of some sort as opposed to simply accidental errors or defects. An old security maxim puts it this way: a system without an adversary model cannot be insecure, just surprising.

Research into design and analysis of anonymous communication began in the cryptologic community where vetting and deploying cryptosystems and protocols takes a long time and where a practical break in a deployed system can

⁴ Thanks to Nick Mathewson for this analogy.

have serious and widespread consequences. A nice example of this is the evolution from initially discovered weaknesses to practical attacks on MD5. Some vulnerabilities are only of theoretical importance, however. Though they may motivate further research, they will never themselves be implemented in a significant practical attack even if systems are never modified in response to them. They will never be the lowest hanging fruit. Contrast this with the Dolev-Yao adversary model [17], an attack against which typically implies dangerous real-world attacks against any implementation of a protocol that is vulnerable to it. Anonymous communications research started with the crypto approach of looking at how hard it is to break the design of an isolated component that might be crucial. The community has still not fully recognized the very limited usefulness such an approach would have for designing and analyzing practical widescale systems for anonymous communication.

Mix networks get their security from the mixing done by their component mixes, and may or may not use route unpredictability to enhance security. Onion routing networks primarily get their security from choosing routes that are difficult for the adversary to observe, which for designs deployed to date has meant choosing unpredictable routes through a network. And onion routers typically employ no mixing at all. This gets at the essence of the two even if it is a bit too quick on both sides.⁵ Mixes are also usually intended to resist an adversary that can observe all traffic everywhere and, in some threat models, to actively change traffic. Onion routing assumes that an adversary who observes both ends of a communication path will completely break the anonymity of its traffic. Thus, onion routing networks are designed to resist a local adversary, one that can only see a subset of the network and the traffic on it.

Given the fundamental differences in the mechanisms they employ, the adversaries they are intended to resist, and their basic designs (not to mention typical applications) it might seem impossible or at least astonishing that anyone who works in this area would ever confuse the two. Yet for years it has been common for publications by even top researchers in anonymous communication to refer to onion routing networks as mixnets or vice versa. All deployed onion routing networks do use some form of layered encryption on traffic they carry, encryption that is gradually removed as it passes through the network. And this is also true of decryption mixnets (re-encryption mixes work differently). Thus there is a clear similarity between the two in at least this respect. Still, given the differences, it would be surprising if this were enough to confuse an expert. But, if you start from an entropist conception of anonymity the confusion becomes less surprising: if you start from an entropist conception of anonymity, all anonymity designs are trying to make a given set of users (or user communications) less dis-

⁵ Other typical and highly salient distinctions include that all existing onion routing network designs are for carrying bidirectional low-latency traffic over cryptographic circuits while public mixnets are designed for carrying unidirectional high-latency traffic in connectionless messages. (An exception is the Web MIXes design [3] as deployed in JonDonym [27], which creates bidirectional circuits through a mix cascade to carry public web traffic.)

tinguishable. If one motivates design by starting with such a set and seeing how well the system obscures identification of its elements, the security contributions of an onion routing approach are harder to see. Any distinction between onion routing networks and mixnets, if recognized at all, is then likely to be couched only in terms of differences in intended application or engineering tradeoffs of security (in the entropist model) versus performance. Even we designers of onion routing systems have been occasionally guilty of falling into this idiom.

But does using the entropist conception reveal potential low-hanging fruit for current or future real systems? To answer this we should consider adversaries that can conduct practical attacks. For brevity's sake, I describe only the one that matters most in examining if entropism is the best approach for widely-used systems like Tor [44], Mixmaster [30], and the Anonymizer [2], viz: **The Man**.⁶

The Man owns big chunks of the anonymity infrastructure, either because he simply set them up himself, or because they are not hardened against takeover. He can also get access to ISPs, backbones, and websites, will know ancillary things, and, if targeting you, will have you specifically under physical surveillance. Think organized crime, state level actors (intelligence, secret police), etc. The Man subsumes the other adversaries we might consider.

In the literature, a standard anonymity adversary is the global-passive adversary (GPA), who controls no nodes in the anonymity network but can observe all traffic on all links to and from that network as well as between nodes of the network. This adversary can observe all sending and receiving behavior by all principals interacting with the network. It thus fits nicely in the entropist model of anonymity and facilitates formal analysis therein. While nice for producing both positive and negative provable results, the GPA is too strong to be realistic because it can observe absolutely every link everywhere regardless of network size. Good security counsels a conservative view, assuming an attacker that is stronger than anything one would encounter in practice. The GPA might thus be considered an overstatement of The Man—except that the GPA is also much too weak; it cannot even briefly delay a packet passing over a single link under its observation. As has long been recognized, the GPA is both too strong and too weak for low-latency distributed systems like onion routing [38, 42, 14].

All low-latency systems as currently designed and deployed are essentially broken against The Man, but often much weaker adversaries are adequate. Single proxies are vulnerable to timing attacks by much weaker adversaries, for example, a single well-placed network observer, and of course communication is vulnerable to the proxy itself if it is corrupt or compromised. Realtime web cascades, such as JonDonym [27] are generally to be comprised of nodes run by recognized but not mutually trusting authorities and thus to distribute trust possibly without the complications of a large, diversely trusted network. A single observer, as would

⁶ We ignore herein any attacks other than by compromising or observing network elements, and by altering and observing the timing and/or volumes of traffic sent over network elements and connecting links. For example, any attacks by injecting identifying information into an anonymity circuit or anonymous message are beyond our scope.

threaten a single proxy, will not be effective. Two such well-placed observers or corrupt insiders, however, would be. If one could enforce uniform sending and receiving behavior of a persistent set of users, a stronger adversary might be needed. But there is to date no accepted scheme for practically doing this against even a passive adversary. And, if there were, that would not defeat an adversary that selectively blocks or manipulates traffic at sources or destinations. Versions of onion routing that do not use entry guards [34] are statistically broken against a small number of well-placed attackers for longterm repeated connections between the same sources and destinations.

As has already been noted, an adversary observing an entry node and an exit node of the existing Tor network can trivially link the two ends of the connection and correlate source to destination regardless of the number of users of these or any other nodes. Anonymity is broken if the endpoints are owned and not if they are not. There are also potential attacks such as website fingerprinting [22, 28] or network latency [23] that require only the entry node to be observed by the adversary. But, in contrast to the entropist view, this again is not significantly affected by the number of other simultaneous circuits initiating at that node and not affected at all by the numbers of circuits elsewhere in the network.

The problem Mixmaster faces against The Man is not the strength of its component protection mechanisms per se but its usability and the implications thereof. (Similarly for Mixminion [9]). With enough compromised nodes it becomes more likely that messages will traverse an entirely compromised path, but many paths might be at most partially compromised. And in mix systems, unlike onion-routing systems, it is not trivial to recognize the same traffic when it is observed in two places. Thus, the number of messages entering and leaving uncompromised mixes does affect uncertainty, and that is indeed the focus of much analysis of mix systems. But there are important factors that largely obviate the protection this might afford.

User configurability and interface play a role [41, 13, 15] in the amount of use an anonymity system gets, but just as important is the latency inherent in the design: the less the latency the more the system becomes vulnerable to correlation attacks. But, keeping the latency high as in remailer mix networks means that only the users with the most sensitive requirements will use the system [1]. And the most common interactive bidirectional applications such as web browsing, chat, or remote login will not work at all, which is probably an even larger factor limiting adoption. Tor has an estimated quarter million or more concurrent users routing over thousands of network nodes. Though Mixmaster has been around longer, it has had perhaps a few hundred users on any given day, and the network has never been more than a few dozen nodes. Mixmaster's numbers have not grown significantly for years, and the incentive issues we have been discussing imply this to be an inherent limitation. If your goal is to have legal deniability for activity this may be adequate, at least initially before The Man has decided to focus on you. But if your adversary *is* The Man, then avoiding suspicion at all is likely an important goal. As such a few hundred initially possible sources or destinations hidden by a few dozen nodes is just

inadequate. Simple externalities such as geographic locations associated with IP addresses, previously known associations, etc. are likely to narrow the initial set even further very quickly [12]. And The Man can own many of the nodes in the network if he so desires so that the already small number of initially unexposed source-destination pairs shrinks commensurately.

The important thing is that the number of possible targets and number of initially needed network observation/compromise points is small enough to be within the capabilities of The Man to bring other resources to bear irrespective of the ability of the anonymization network to render its users indistinguishable. The issue is not simply a question of uncertainty in identifying a large set of senders and/or recipients, it is an issue of available resources and their expected effectiveness. Yes, The Man strives to reduce his uncertainty about linking senders and recipients, but the measure of his ability to succeed is not the size of those sets or probability distributions on them. Rather it is the resources it will take him to place himself in a position to learn those linking relations. And uncertainty about which elements to attack to improve his position may also play a role, but its role is not as important as the expected cost and benefit of attacking each of them, especially if he has the capability to effectively attack all of them—as is the case for the deployed Mixmaster network.

It is true that entropist approaches might help narrow the initial set of a few hundred senders of a given message. For example, if communication pairs are maintained over hundreds of communication rounds, then statistical disclosure [10, 29] may be useful. And The Man should be able to further reduce anonymity within network traffic by DoS attacks on uncompromised nodes [16, 4]. If that would be too easily detected, The Man can combine DoS with bridging those nodes via trickling and flooding from his own clients [39] and owned nodes or links. Note that while these active attacks can be analyzed in the entropist model for effectiveness, no entropist view is necessary to simply deploy them in order to slightly improve the efficiency of an already effective compromise. They are just not the attack threats that matter most when going against The Man.

In sum, all current low-latency anonymity systems are broken against The Man. Onion-routing systems by nature of their potential for wide distribution are relatively resistant to other kinds of adversaries. Mix networks, such as Mixmaster and Mixminion, do have measures against The Man. But they are overkill against other expected attackers. This, coupled with the overhead of trying to use them, limits the numbers of both users and infrastructure providers to a point that they are vulnerable to direct inspection and attacks by The Man for which the anonymizing network becomes irrelevant. While entropist techniques can play some role here, they are not the central issue.

5 So what is anonymous communication?

What the entropist conception gets right is that, like other security properties, anonymity protection is determined by the amount of work an adversary needs to do to overcome it. What it gets wrong is the kinds of work the adversary

needs to do. We have observed so far that an anonymity breaking adversary generally does not do best by focusing directly on reducing an initially known set of senders (receivers, etc.) to the point that the sender is uniquely determined or has low enough entropy to be considered busted. The entropist approach will yield theoretically interesting analyses. It will yield system designs that might be useful in controlled settings like elections, where we can control or predict the number and nonidentity of participants and where anonymity within expected-size sets is useful. But the entropist approach is not appropriate for general communication on large diversely shared networks like the internet. I do not set out here a definitive answer of what to replace it with: this is a breaking paper not a design paper. I do, however, offer some hints and suggestions.

As we have noted, putatively more secure designs, such as DC-nets or mix networks, are actually not secure against the intended adversary unless there is a significant increase in both network size and persistent users. And usability and incentive issues make that unlikely. A network might scale up to a point that, even if The Man were able to compromise much communication passing through the network, any individual communication would not be likely to be de-anonymized. P2P anonymity designs have this potential, although they require further scrutiny because they are complex, and attacks on, e.g., network discovery or versions of sybil attacks seem hard to eradicate. Decentralized designs with semicentralized management of network information, such as Tor's, also have this potential. But in either case one is forced to make some assumptions about the likelihood that a given percentage of nodes is under hostile control, and this can be impossible to gauge. Any design that assumes a roughly uniform distribution of resistance to attack across all nodes is likely doomed if The Man can own botnets within the network (and why shouldn't he?). A useful metric should therefore evaluate rather than assume how hard it is for the adversary to own a portion of the network. This much is true of mix networks as well as onion-routing networks and true of peer-to-peer as well as more centralized versions of either of these.

How can we gauge an adversary's difficulty in owning chunks of the network and the traffic flowing over them? We need not analyze how hard it is to contribute hostile servers, break into nodes, steal administrative passwords, control or observe inputs and outputs to nodes etc. For our metrics we just need a representation of the possible outcomes of that analysis. Our observations suggest that it will be useful if those outcomes are not considered uniform across all nodes. So it is not just a question of implementing the strongest component design and deploying it widely. We will also need to consider how much the adversary learns from the nonuniformity itself. We should also consider initial awareness of the network, whether the adversary knows the same network elements and structure as do the communicants he is attacking. Such considerations are a focus of our current ongoing research.

6 Rotor and other immodest proposals

In this paper we have argued that entropist metrics for anonymity are not appropriate to widely shared and deployed anonymity networks. We have observed that they rely on variables for which it is unrealistic to attach values and that would not significantly affect how well anonymity is protected if they could be determined. For less widely deployed networks, whether centralized or distributed, they provide the most relevant adversary with a set of communicating pairs that he has the resources to deanonymize.

We end with some implications for anonymity research and system design. In the mid nineties (before decent SSL encryption was available in every browser) encrypted traffic was relatively rare on the internet; the mere fact that a message was encrypted made it interesting. Thus, we used to say that if you encrypt your traffic at all, you should make sure you encrypt it well because it is likely to be scrutinized. The same point applies to anonymizing your communication. Unlike choosing good algorithms and increasing keylength for encryption, however, usability and incentive limitations make it unlikely that users and nodes for mix networks will ever increase adequately to exceed the resources of The Man. And even if they did, there is no way to have any meaningful assurance about the distribution of noncollaborating users. More importantly, there is no way to have meaningful assurance about the distribution of noncollaborating network nodes. Our first modest suggestion is thus that existing mix networks for general internet use should simply be abandoned for other than research purposes. They should continue to be studied for their inherent interest. And, they should be used for applications where it is possible to manage and measure the sets of distinct users and anonymity providers and the probability distributions on their behaviors, voting being the clearest example. But for general internet use, they are overkill against almost every adversary except unrealistic ones like the GPA or incredibly strong ones like The Man. And, because of usability and incentive limitations, in practice they do not scale enough to protect against The Man anyway. On the other hand a widely distributed network like Tor may already offer better, though still inadequate, protection. If we modify the design of onion-routing networks like Tor so that trust is incorporated into routing (as we discuss briefly below), then it should be harder for The Man to attack where he must to be effective.

Also, we should not worry (too much) about partitioning attacks. In distributing network directories, Tor has conservatively opted for distributing awareness of the entire directory to clients. This has presented scaling challenges, but the ability to differentiate users by partitioning what they know and don't know is not fully understood. And there have been published epistemic attacks on both mix and onion-routing network designs [8, 11]. But unless the adversary is approaching The Man in both capability and intent, this is not a significant concern. And if the adversary is on the order of The Man, then this is not effective. A caveat is that the design should not allow simple actions on the part of a small player, one that can affect at most a small part of the infrastructure, to partition to a point that it can effectively isolate a small number of clients with

little effort. Entropy is not useless for characterizing anonymity, it is just not the primary measure of a practical anonymity network’s security and as such should not be the primary driver of design.

If the primary measure of anonymity is how hard it is for an adversary to observe or own parts of the network that allow him to conduct attacks, then we can represent the amount of trust we have that a node will not be compromised and then route accordingly. In addition to anonymity system nodes, the same basic trust notion can be applied to links and destinations. We are currently researching such modeling and analysis of trust and have several initial results for optimal route selections assuming simply that network nodes are trusted at different levels [25]. Our results are based on a partial adversary that can attempt to compromise a fixed fraction of network nodes and succeeds for any given node n_i with probability p_{c_i} . This is a variant of the roving adversary introduced by Ostrovsky and Yung [33] and applied to anonymous routing in [42]. What we add is the idea of trust in a node n_i , represented as $1 - p_{c_i}$ and applied as just described.

We expect the adversary model that we ultimately develop fully will add mobility to the above, much as Ostrovsky and Yung’s original work did. This will allow for an iterative attack in which an adversary can discover parts of the network, trust relationships, communication patterns, etc. and then apply what is learned at each stage to conduct further attacks. Anonymity then becomes a question of what an adversary can uncover given a budget of resources expended over an attack of a given duration.

Once better understood, one could imagine a trust-based onion routing network: roTor. This would be a variant of Tor built to counter such an adversary that gets its protection through trust; hence the name.⁷

A rotor network may have a place for mixing too. Routing based on trust could mean that the expected chance that a given communication is compromised is quite low. But, for very sensitive communications it still may pay to have some defense in depth so that observing in a few opportune spots will not yield easy timing correlations. The goal is likely to be hiding of particular linkings rather than avoiding suspicion that the communicants talk at all. But that

⁷ Tor began as *the* onion routing, designed as one of the original NRL onion routing projects, in contrast to the many other onion routing designs that followed the NRL original. It is also a recursive acronym for *Tor’s onion routing*. (‘Tor’ has never been an acronym for *The Onion Router*—as generally misstated in the press, nor has it ever been proper to write it ‘TOR’.) Though that is perhaps enough of a hint to make the following implicitly obvious to the reader, we note that ‘rotor’ is a sesquipalendromic sesqurecursive acronym for *rotor onion T³or’s onion rotor*, which itself expands into *rotor onion [Tor’s trusted through] onion rotor*. As also implied by the name, the T-words rotate so that it also expands to *rotor onion [trusted through Tor’s] onion rotor* and to *rotor onion [through Tor’s trusted] onion rotor*, all three of which should be viewed as in the same equivalence class. Defining the algebraic structure for the equivalence relation is left as an exercise for the reader. Note also that the nodes of a rotor network are the network rotors; thus a rotor network *in sensu composito* is also a rotor network *in sensu diviso*.

is likely to yield rather different designs from existing mixes and mix networks, for example, embedding messages to be mixed by an intermediate node inside of web traffic on a low-latency circuit.

This last section has been fairly speculative so as to attempt to give at least one vision of what definitions and system designs might replace entropism. Whether or not these ultimately prevail, however, is not the central point of this paper. The central point is that the model on which all existing work on anonymity, the entropist model, is broken for open or widely shared communications. That result remains, even if we do not know definitively what to put in its place.

None of the usability, incentive, or usage observations made herein are particularly novel. However, engendered by its fundamental entropism, the community still operates on a premise that mix networks are more secure. Those of us who have made such observations have even sometimes moderated our assumptions so that mix networks are described as more secure in principle if not in current deployment. But taking the next steps implied by these observations has not even been overtly broached. All of us have still clung to an entropist view of anonymity and thus to the systems and analysis it engenders. Hopefully, this paper will give us the push we need to let go.

Wait, what about interference attacks that allow remote observation without compromising or directly observing nodes and links [32, 31, 18]? Sorry. We've got an answer, but there's no room in the margins. See the next paper.

Acknowledgments

I thank Claudia Diaz, George Danezis, Roger Dingledine, Aaron Johnson, and Nick Mathewson for helpful discussions that led to better articulations of the ideas in this paper. (Nonetheless any remaining incoherence is my own.) This work supported by ONR.

References

1. Alessandro Acquisti, Roger Dingledine, and Paul Syverson. On the economics of anonymity. In Rebecca N. Wright, editor, *Financial Cryptography, 7th International Conference, FC 2003*, pages 84–102. Springer-Verlag, LNCS 2742, 2003.
2. The Anonymizer. <http://www.anonymizer.com/>, 2009. Homepage of the company that offers the Anonymizer Proxy Service. Original Anonymizer first described in [5].
3. Oliver Berthold, Hannes Federrath, and Stefan Köpsell. Web MIXes: A system for anonymous and unobservable Internet access. In Hannes Federrath, editor, *Designing Privacy Enhancing Technologies: International Workshop on Design Issues in Anonymity and Unobservability*, pages 115–129. Springer-Verlag, LNCS 2009, July 2000.
4. Nikita Borisov, George Danezis, Prateek Mittal, and Parisa Tabriz. Denial of service or denial of security? How attacks on reliability can compromise anonymity. In Sabrina De Capitani di Vimercati, Paul Syverson, and David Evans, editors,

- CCS'07: Proceedings of the 14th ACM Conference on Computer and Communications Security*, pages 92–102. ACM Press, 2007.
5. Justin Boyan. The Anonymizer: Protecting user privacy on the web. *CMC Magazine*, September 1997.
 6. David Chaum. Untraceable electronic mail, return addresses, and digital pseudonyms. *Communications of the ACM*, 4(2):84–88, February 1981.
 7. Sebastian Clauß and Stefan Schiffner. Structuring anonymity networks. In Atsuhiko Goto, editor, *DIM'06: Proceedings of the 2006 ACM Workshop on Digital Identity Management*, pages 55–62, Alexandria, VA, USA, 2006. ACM Press.
 8. George Danezis and Richard Clayton. Route fingerprinting in anonymous communications. In *Sixth IEEE International Conference on Peer-to-Peer Computing, P2P 2006*, pages 69–72. IEEE Computer Society Press, 2006.
 9. George Danezis, Roger Dingledine, and Nick Mathewson. Mixminion: Design of a type III anonymous remailer protocol. In *Proceedings, 2003 IEEE Symposium on Security and Privacy*, pages 2–15, Berkeley, CA, May 2003. IEEE Computer Society.
 10. George Danezis and Andrei Serjantov. Statistical disclosure or intersection attacks on anonymity systems. In Jessica Fridrich, editor, *Information Hiding: 6th International Workshop, IH 2004*, pages 293–308. Springer-Verlag, LNCS 3200, May 2004.
 11. George Danezis and Paul Syverson. Bridging and fingerprinting: Epistemic attacks on route selection. In Nikita Borisov and Ian Goldberg, editors, *Privacy Enhancing Technologies: Eighth International Symposium, PETS 2008*, pages 151–166. Springer-Verlag, LNCS 5134, July 2008.
 12. George Danezis and Bettina Wittneben. The economics of mass surveillance and the questionable value of anonymous communications. In Ross Anderson, editor, *Fifth Workshop on the Economics of Information Security (WEIS 2006)*, June 2006.
 13. Roger Dingledine and Nick Mathewson. Anonymity loves company: Usability and the network effect. In Ross Anderson, editor, *Fifth Workshop on the Economics of Information Security (WEIS 2006)*, June 2006.
 14. Roger Dingledine, Nick Mathewson, and Paul Syverson. Tor: The second-generation onion router. In *Proceedings of the 13th USENIX Security Symposium*, pages 303–319. USENIX Association, August 2004.
 15. Roger Dingledine, Nick Mathewson, and Paul Syverson. Deploying low-latency anonymity: Design challenges and social factors. *IEEE Security & Privacy*, 5(5):83–87, September/October 2007.
 16. Roger Dingledine, Vitaly Shmatikov, and Paul Syverson. Synchronous batching: From cascades to free routes. In David Martin and Andrei Serjantov, editors, *Privacy Enhancing Technologies: 4th International Workshop, PET 2004*, pages 186–206. Springer-Verlag, LNCS 3424, May 2005.
 17. Danny Dolev and Andrew C. Yao. On the security of public-key protocols. *IEEE Transactions on Information Theory*, 2(29):198–208, March 1983.
 18. Nathan S. Evans, Roger Dingledine, and Christian Grothoff. A practical congestion attack on Tor using long paths. In *Proceedings of the 18th USENIX Security Symposium*, pages 33–50, Montreal, Canada, August 2009. USENIX Association.
 19. David M. Goldschlag, Michael G. Reed, and Paul F. Syverson. Hiding routing information. In Ross Anderson, editor, *Information Hiding: First International Workshop*, pages 137–150. Springer-Verlag, LNCS 1174, 1996.

20. David M. Goldschlag, Stuart G. Stubblebine, and Paul F. Syverson. Temporarily hidden bit commitment and lottery applications. *International Journal of Information Security*, 9(1):33–50, February 2010.
21. Sabine Helmers. A brief history of anon.penet.fi - the legendary anonymous remailer. *CMC Magazine*, September 1997.
22. Andrew Hintz. Fingerprinting websites using traffic analysis. In Roger Dingle-dine and Paul Syverson, editors, *Privacy Enhancing Technologies: Second International Workshop, PET 2002*, pages 171–178, San Francisco, CA, USA, April 2002. Springer-Verlag, LNCS 2482.
23. Nicholas Hopper, Eugene Y. Vasserman, and Eric Chan-Tin. How much anonymity does network latency leak? In Sabrina De Capitani di Vimercati, Paul Syverson, and David Evans, editors, *CCS'07: Proceedings of the 14th ACM Conference on Computer and Communications Security*, pages 82–91. ACM Press, 2007. Expanded and revised version in [24].
24. Nicholas Hopper, Eugene Y. Vasserman, and Eric Chan-Tin. How much anonymity does network latency leak? *ACM Transactions on Information and System Security*, 13(2):13–28, 2010. February.
25. Aaron Johnson and Paul Syverson. More anonymous onion routing through trust. In *22nd IEEE Computer Security Foundations Symposium, CSF 2009*, pages 3–12, Port Jefferson, New York, USA, July 2009. IEEE Computer Society.
26. Aaron Johnson, Paul Syverson, Roger Dingledine, and Nick Mathewson. Trust-based anonymous communication: Adversary models and routing algorithms. In *CCS'11: Proceedings of the 18th ACM Conference on Computer and Communications Security*. ACM Press, October 2011.
27. JonDonym – the internet anonymisation service. <https://www.jondos.de/en/>, 2008. Commercial version of the Java Anon Proxy (JAP). Initially published description in [3].
28. Marc Liberatore, , and Brian Neil Levine. Inferring the source of encrypted HTTP connections. In Rebecca N. Wright, Sabrina De Capitani di Vimercati, and Vitaly Shmatikov, editors, *CCS'06: Proceedings of the 13th ACM Conference on Computer and Communications Security*, pages 255–263. ACM Press, 2006.
29. Nick Mathewson and Roger Dingledine. Practical traffic analysis: Extending and resisting statistical disclosure. In David Martin and Andrei Serjantov, editors, *Privacy Enhancing Technologies: 4th International Workshop, PET 2004*. Springer-Verlag, LNCS 3424, 2005.
30. Ulf Möller, Lance Cottrell, Peter Palfrader, and Len Sassaman. Mixmaster protocol - version 3. IETF Internet Draft, 2003.
31. Steven J. Murdoch. Hot or not: Revealing hidden services by their clock skew. In Rebecca N. Wright, Sabrina De Capitani di Vimercati, and Vitaly Shmatikov, editors, *CCS'06: Proceedings of the 13th ACM Conference on Computer and Communications Security*, pages 27–36. ACM Press, 2006.
32. Steven J. Murdoch and George Danezis. Low-cost traffic analysis of Tor. In *2005 IEEE Symposium on Security and Privacy, (IEEE S&P 2005) Proceedings*, pages 183–195. IEEE CS, May 2005.
33. Rafail Ostrovsky and Moti Yung. How to withstand mobile virus attacks. In *Proceedings of the Tenth ACM Symposium on Principles of Distributed Computing (PODC '91)*, pages 51–59. ACM Press, 1991.
34. Lasse Øverlier and Paul Syverson. Locating hidden servers. In *2006 IEEE Symposium on Security and Privacy (S&P 2006), Proceedings*, pages 100–114. IEEE CS, May 2006.

35. Sameer Parekh. Prospects for remailers: where is anonymity heading on the internet? *First Monday*, 1(2), August 5 1996. On-line journal <http://www.firstmonday.dk/issues/issue2/remailers/>.
36. Andreas Pfitzmann and Marit Köhntopp. Anonymity, unobservability, and pseudonymity – a proposal for terminology. In Hannes Federrath, editor, *Designing Privacy Enhancing Technologies: International Workshop on Design Issues in Anonymity and Unobservability*, pages 1–9. Springer-Verlag, LNCS 2009, July 2000. A regularly updated version of this paper is maintained at http://dud.inf.tu-dresden.de/Anon_Terminology.shtml.
37. Andreas Pfitzmann and Marit Köhntopp. A terminology for talking about privacy by data minimization: Anonymity, unlinkability, undetectability, unobservability, pseudonymity, and identity management, version v0.32. http://dud.inf.tu-dresden.de/Anon_Terminology.shtml, December 2009. Regularly revised and updated version of [36].
38. Michael G. Reed, Paul F. Syverson, and David M. Goldschlag. Anonymous connections and onion routing. *IEEE Journal on Selected Areas in Communications*, 16(4):482–494, May 1998.
39. Andrei Serjantov, Roger Dingledine, and Paul Syverson. From a trickle to a flood: Active attacks on several mix types. In Fabien A.P. Petitcolas, editor, *Information Hiding: 5th International Workshop, IH 2002*, pages 36–52. Springer-Verlag, LNCS 2578, 2002.
40. Gustavus J. Simmons. The history of subliminal channels. *IEEE Journal on Selected Areas in Communications*, 16(4):452–462, May 1998.
41. Paul Syverson, Michael Reed, and David Goldschlag. Onion Routing access configurations. In *Proceedings DARPA Information Survivability Conference & Exposition, DISCEX'00*, volume 1, pages 34–40. IEEE CS Press, 1999.
42. Paul Syverson, Gene Tsudik, Michael Reed, and Carl Landwehr. Towards an analysis of onion routing security. In Hannes Federrath, editor, *Designing Privacy Enhancing Technologies: International Workshop on Design Issues in Anonymity and Unobservability*, pages 96–114. Springer-Verlag, LNCS 2009, July 2000.
43. Can Tang and Ian Goldberg. An improved algorithm for Tor circuit scheduling. Technical Report CACR 2010-06, University of Waterloo, Center for Applied Cryptography Research, 2010. <http://www.cacr.math.uwaterloo.ca/techreports/2010/cacr2010-06.pdf>.
44. The Tor Project. <https://www.torproject.org/>, 2009. Homepage of the non-profit organization that maintains and develops the Tor network. Original Tor design first published in [14].