Defining "Anonymity" in Networked Communication, version 1

Joan Feigenbaum[1]
Technical Report YALEU/DCS/TR-1448
December 2011

Support for anonymous communication in hostile environments is the main goal of DARPA's "Safer Warfighter Communications" program (DARPA-BAA-10-69: SAFER). Despite the fact that the word is regularly encountered in common parlance, "anonymity" is actually a subtle concept – one that the computer-science community has worked hard (and not yet completely successfully) to define precisely. This briefing paper explains the concept of anonymity in the context of the SAFER program; it is intended for a general audience rather than for one with professional expertise in computer science and data networking.

Approved for public release: distribution unlimited
Keywords: Anonymity, blocking, disruption, observability, pseudonymity

---

Defining "Anonymity" in Networked Communication

Joan Feigenbaum
http://www.cs.yale.edu/homes/jf/

Version 1, December 2011

Support for anonymous communication in hostile environments is the main goal of DARPA's "Safer Warfighter Communications" program (DARPA-BAA-10-69: SAFER). Despite the fact that the word is regularly encountered in common parlance, "anonymity" is actually a subtle concept – one that the computer-science community has worked hard (and not yet completely successfully) to define precisely. This briefing paper explains the concept of anonymity in the context of the SAFER program; it is intended for a general audience rather than for one with professional expertise in computer science and data networking.

1. Introduction

In 2010, DARPA started the "Safer Warfighter Communications" program, referred to by program performers and in the rest of this document as "SAFER." As stated in the BAA (DARPA-BAA-10-69), "The goal of the … SAFER program is to develop technology that will enable safe, resilient communications over the Internet, particularly in situations in which a third party is attempting to discover the identity or location of the end users or to block the communication." In other words, the goal of the SAFER program is to enable anonymous communication in hostile environments.

Anonymous communication is a complex research area, characterized by multiple, subtle goals and a great deal of confusing terminology. The purpose of this paper is to explain some (but not all) of the key concepts in the area in a manner that is accessible by a general audience and to relate these concepts to the SAFER program. In particular, we will clarify what it means for third parties to identify or locate the senders or receivers of a given communication stream or to disable the communication stream altogether; as stated in the above quote from the SAFER BAA, the mission of the program is to develop technology that prevents third parties from accomplishing these goals. We present the first goal in Section 3 below and the second in Section 4 after introducing some basic building blocks in Section 2. Three SAFER use cases put forth in the SAFER CONOPS document [4] are explored in Section 5.

Note that the SAFER BAA speaks of "identification and location of end users." Throughout this paper, we assume that "locating" is about determining the network locations at which SAFER technology is being used (without necessarily learning anything about the human beings who are using it) and that "identifying" is about determining the real-world identities of SAFER users. In the computer-security literature, these activities are described in technical terms such as "observable," "anonymous," and "pseudonymous"; we will use these established terms in our discussion below.

2. Domain of Discourse

We refer to users of SAFER technologies as *communicators*. The *network locations* (*e.g.*, personal computers, mobile phones, web servers, Internet domains – basically anything with an IP address) at which the communicators create, send, receive, use, and store SAFER traffic exist in a larger networked *environment* or *surrounding* in which non-SAFER network technologies are also in use, and some (perhaps most) of the users are not SAFER communicators.

Communicators use SAFER technologies in order to evade *adversaries*, *i.e.*, those third parties referred to in the BAA who seek to "discover the identity or location of the end users or to block the communication." Adversaries who take explicit steps to disrupt the activities of the communicators, *e.g.*, by blocking or corrupting their traffic streams or by shutting down their local networks or websites, are often called *attackers*. Often, however, the adversary's goal is not to attack the communicators but rather to exploit them. An adversary who identifies a SAFER traffic stream that he leaves in place and listens in on is called an *eavesdropper*. One who identifies a group of SAFER users that he leaves in place and manages to join as an active participant is called an *infiltrator*.

Following Pfitzmann and Hansen [6], we use the acronym IOI (for *items of interest*) in our definitions. IOIs can be of any type; for example, they can be humans (*e.g.*, senders and receivers of SAFER traffic are both IOI sets), locations (*e.g.*, websites or Internet domains), or data items (*e.g.*, email messages or streams of video traffic). In a nutshell, the adversary's goal is to discover relationships among IOIs and to act on these discoveries.

One essential notion in our discussion of anonymity is *linkability*. An adversary has succeeded in *linking two sets A and B of IOIs* if he has obtained significant information about whether pairs *(a, b)* are related, where *a* is in *A* and *b* is in *B*, that he did not have before he deployed his resources in the relevant environment. Here, the meaning of "related" should be clear in context; for example, if *A* is a set of email senders, and *B* is a set of email messages, then *a* and *b* are related if *a* sent *b*. The meaning of "obtaining significant information" is highly technical and beyond the scope of this paper, but it is based on the information-theoretic notion of change in probability. For example, if the adversary's *a priori* belief (*i.e.*, what he thought the state of things was before he deployed his resources) was that each of SAFER communicators a_1 and a_2 was equally likely to have sent each of messages b_1 and b_2, and his *a posteriori* knowledge (*i.e.*, what he knows after having deployed his resources) is that the probability that a_1 sent b_1 is .9, and the probability that a_2 sent b_2 is .75, then he has succeeded in linking $\{a_1, a_2\}$ and $\{b_1, b_2\}$.

A second essential notion is *distinguishability*. Let *A* and *B* be disjoint sets of IOIs, and let *S* be the union of *A* and *B*. An adversary has succeeded in *distinguishing A and B* if he has obtained significant information about whether an element *x* of *S* is in *A* or *B*. For example, an adversary may be interested in distinguishing email messages transmitted by an anonymous-email service developed under the SAFER program and email messages transmitted by a standard (non-anonymous) email service such as GMAIL.

3. Anonymity and Pseudonymity

In common parlance, network traffic is usually considered to be "anonymous" if the adversary cannot correctly "identify" the communicators who create and exchange it. For example, if sender *S* sends message *M* to receiver *R* using an anonymous-email service, then an adversary who observes *M* as it travels through the network will not be able to identify *S* as the

sender of M or R as the receiver of M, even if this adversary knows that S and R are users of this service. The word "attribute" is often used to describe the same anonymity property: An anonymous-email service is one that prevents the adversary from correctly attributing a particular email message to a particular sender or receiver.

To describe the desired properties of SAFER technologies, we need to be more precise about what we mean by anonymity and to consider several different types of anonymity. We do this by considering linkability of various types of IOIs.

The IOIs in our discussion of identification include communicators (*i.e.,* users of SAFER technology in the network environment of interest), network elements (such as web servers) in that environment, and subsets thereof. Which subsets are of interest depends on the goal of the SAFER technology under study; for example, if the goal of this technology is to enable anonymous email, then senders and receivers of email are both sets of IOIs. Note that, in this section, we assume that the adversary already knows that SAFER technology is being used and who is using it. We wish to prevent him from correctly associating particular streams or pieces of SAFER traffic with the particular communicators who create and use them or perhaps from correctly associating particular pairs of SAFER communicators with each other.

The following types of anonymity are studied in the computer-security literature and are highly relevant to the SAFER program.

- o Sender anonymity: Let T be a SAFER communication technology, S_T be the set of communicators that send traffic using T, and U_T be the units of SAFER traffic sent by members of S_T, *e.g.,* the email messages, video streams, or documents. Note that T may be used for one-to-one communication, as in the case of an anonymous-email service, or for one-to-many communication, as in the case of anonymous publication of controversial documents. We say that T *is sender-anonymous* if the IOI sets S_T and U_T are unlinkable.
- o Receiver anonymity: This property is completely analogous to sender anonymity; in this case, the IOI sets that must be unlinkable are R_T and U_T, where R_T is the set of communicators that receive traffic using T. Once again, T may be one-to-one or one-to-many.
- o Client anonymity: Let B be an anonymous-web-browsing tool, W a website to which the adversary would like to restrict (or at least monitor) access, C_B the set of communicators who access W using B, and M_{CBW} the set of messages that B transmits (in either direction) between C_B and W. We say that B *is client-anonymous* if the IOI sets C_B and M_{CBW} are unlinkable.
- o Server anonymity: As in the case of client anonymity, B is an anonymous-web-browsing tool. Each communicator C can use B to access a set W_B of websites; note that W_B may in fact contain most or even all of the servers on the World Wide Web and thus that, if B is a powerful and popular tool, the adversary will face resistance if he simply bans its use. M_{CBW} is the set of messages that B transmits between C and W_B. We say that B *is server-anonymous* if the IOI sets W_B and M_{CBW} are unlinkable.
- o Relationship anonymity: Communication technology T is relationship-anonymous if the IOI sets S_T and R_T are unlinkable. Similarly web-browsing tool B is relationship-anonymous if C_B and W_B are unlinkable. Note that, if T (respectively, B) is relationship-anonymous, it must be sender-anonymous or receiver-anonymous or both (respectively, client-anonymous or server-anonymous or both).

o <u>Pseudonymity</u>: Communication technology *T* (respectively, web-browsing tool *B*) is *pseudonymous* if the IOI set U_T (respectively, M_{CBW}) of units of traffic that it transmits is linkable to itself but unlinkable to any IOI set of communicators. For example, it may be feasible for the adversary to determine that a set of email messages were sent by the same user of a pseudonymous-email tool, but it must be infeasible for him to determine who that user is.

4. Observation and Disruption

In this section, the adversary's goal is to determine who is using SAFER technology, to determine where it is being used, and/or to stop it from being used effectively. Note that this is a switch from Section 3, in which we assumed that the set of users was known and that the adversary's goal was to determine who was saying what or who was talking to whom. Although the SAFER BAA speaks of "block[ing] the communication," we will consider the more general concept of "disrupting the communication." "Blocking" is a crude form of disruption in which the adversary completely stops the SAFER traffic from reaching its destination. Less crude forms include "corrupting," in which the adversary alters the traffic stream so that it is unusable when it reaches its destination, and "malling," in which the adversary alters the traffic stream so that it is usable but does not have the effect that its creator intended. The term "malling" is standard in the cryptography literature but not in the anonymity literature; a "malleable" encryption system is one in which an adversary can transform a ciphertext *C* into a related text *C'* in such a way that the receiver will succeed in decrypting *C'* but will obtain a plaintext that differs (in a manner controlled by the adversary) from the plaintext that the sender encrypted.

The IOIs of interest in this section include network traffic, network users, and network locations. The adversary's goal is to determine which of the items are associated with SAFER technologies. To describe this goal more precisely, we use the notion of indistinguishability.

Let *T* be a SAFER technology and **T** be a set of technologies of which *T* is a member. For example, *T* may be an anonymous-email service, **T** may be the set of all email services, and the adversary may be interested in isolating email traffic that's generated using *T* within the set of all email traffic. Alternatively, *T* may be an anonymous-email service, and **T** may be a broader set of communication technologies that produce traffic in which email could be "camouflaged." The adversary is still trying to isolate the traffic produced and camouflaged by *T* in the set of all traffic produced by **T**, but now his job may be harder, because the traffic produced by *T* does not present itself to the world (*e.g.*, in packet headers) as email. Let **U** be the set of all users of technologies in **T** that are of interest to the adversary, **N** be the set of network locations of interest to the adversary, and **F** be all of the traffic created and used by members of **U** using technologies in **T**. The SAFER communicators of interest to the adversary are a subset U_T of **U**. Similarly, members of U_T use a subset N_T of **N** and create a subset F_T of **F**. Finally, let $U' = \mathbf{U} \setminus U_T$ be the complement of U_T in **U** (*i.e.*, those users of interest to the adversary who are *not* using the SAFER technology *T*), $N' = \mathbf{N} \setminus N_T$ be the complement of N_T in **N**, and $F' = \mathbf{F} \setminus F_T$ be the complement of F_T in **F**.

o <u>Unobservability:</u>
- *Users of T are unobservable* if U_T and *U'* are indistinguishable by the adversary.
- *Traffic produced by T is unobservable* if F_T and *F'* are indistinguishable by the adversary.

- *Locations at which T is used are unobservable* if N_T and N' are indistinguishable by the adversary.
  - ○ <u>Disruption:</u> We assume that, for *T* to be *disruptable*, it must be *observable*, *i.e.*, it must fail to satisfy the previous definitions[2]. Moreover, there are at least three forms that disruption by the adversary could take:
    - *SAFER technology T is blockable* if the traffic that it generates is observable and, once observed, can be blocked by the adversary.
    - *SAFER technology T is corruptible* if the traffic that it generates is observable and, once observed, can be corrupted by the adversary.
    - *SAFER technology T is malleable* if the traffic that it generates is observable and, once observed, can be malled by the adversary.

Note that the ability to observe SAFER traffic does not imply the desire to disrupt it. An adversary who observes a SAFER-traffic stream may prefer to continue observing it, to attempt to identify the SAFER communicators who produced it, and, if he succeeds, to attempt to infiltrate a group of SAFER communicators and learn what they are saying (even if he cannot attribute specific messages to their senders and receivers).

5. Use Cases

We now use the terminology developed in Sections 2, 3, and 4 to describe three SAFER-program use cases that were put forth in the SAFER CONOPS document [4]; all address the need for SAFER communicators to browse websites to which adversaries want to restrict access. Ideally, SAFER browsing technology would provide client anonymity, server anonymity, and resistance to all forms of disruption in all three scenarios. Even if this ideal is not achievable, SAFER technology may be able to provide some valuable protection; thus, in our discussion below, we point out the highest-priority goal or goals in each scenario.

Figure 1 depicts the Department of Defense (DoD) use case. SAFER communicators are working in a protected network enclave E inside a surrounding, potentially adversarial environment A. For example, E may be an American military base, and A may be an enemy country (or, if not an outright enemy, perhaps a country with which the US has ambiguous relations). A communicator C in E needs to browse websites in A (but outside of E). Assuming that there is just one communicator of interest (as depicted in Figure 1) and that the adversary knows who he is, the highest priority in this scenario is for the SAFER browsing technology that C uses to provide server anonymity; that is, the adversary should be unable to attribute particular traffic streams sent to or from C to particular sensitive websites. If there are multiple communicators of interest, say C_1 and C_2, then it would also be desirable for the SAFER browsing technology to provide client anonymity; this would ensure that, even if there were only one sensitive website of interest, say Site 1, the adversary could not reliably determine which of C_1 or C_2 was responsible for particular traffic streams sent to or from that site. If, in the case of multiple sites and multiple communicators, it is infeasible to provide both server anonymity and client anonymity, it may still be feasible to provide relationship anonymity, *i.e.*, to prevent the adversary from reliably determining which communicators are accessing which sites.

The rationale for prioritizing anonymity in this scenario is that the adversary is less likely to block or otherwise disrupt than he is to try to break anonymity: E is known to be a DoD

---

[2] This assumption may be relaxed or dropped in future versions of this briefing paper.

enclave, and presumably the adversary is aware that the DoD could retaliate quite forcefully if its operations are disrupted.


Figure 1: The Department of Defense Use Case

Figure 2 depicts the Strategic-Communication use case. Communicators C_1 through C_n, for some large value of *n*, are located in a country that is (at least partly) hostile to the US and wish to receive information from a website outside of that country. Here, the highest priority is to provide SAFER browsing technology that is not blockable and, preferably, is not disruptable in any fashion. The assumption is that it is well known that many people in the country access this website and that the adversary would thus face resistance or cause social unrest if he started punishing people for doing so; on the other hand, blocking (or just degrading the quality of) the communication may be something that the adversary can get away with (*e.g.*, because users blame the website operator or the Internet Service Provider for poor service rather than blaming the adversary). If one were interested in anonymity as well in this scenario, the goal would be client anonymity: Everyone, including the adversary, knows the name and location of the strategic website; the SAFER browsing technology would thus be charged with preventing the adversary from linking the set $\{C_1, \ldots, C_n\}$ and the traffic streams from this site.


Figure 2: The Strategic-Communication Use Case

Figure 3 depicts the dissident use case. Communicators C_1 through C_n in a repressive environment want to visit k public websites outside of that environment. In this scenario, they really need SAFER browsing technology that provides client anonymity and that is not blockable (or, better yet, not disruptable at all). If the traffic is disrupted, it will not serve its political purpose. If particular traffic streams can be linked to individual communicators, then those dissident individuals will be punished harshly by the repressive regime.

Figure 3: The Dissident Use Case

6. Additional Sources

The most thorough and frequently cited paper about anonymity terminology is the one by Pfitzmann and Hansen [6]. Two standard approaches to the design and analysis of anonymous-communication systems are Dining-Cryptographer nets [1, 2] and MIX nets [3]; both are in use by SAFER performers, along with newer and more novel approaches. TOR is an open-source, widely deployed, flexible, anonymous-communication system [7]. The Privacy Enhancing Technologies Symposium is an annual conference that features many current works on anonymity [5].

7. References
[1] D. Chaum, "Security without Identification: Transaction Systems to make Big Brother Obsolete," *Communications of the ACM* **28:10** (1985), pp. 1030-1044.
[2] D. Chaum, "The Dining Cryptographers Problem: Unconditional Sender and Recipient Untraceability," *Journal of Cryptology* **1:1** (1988), pp. 65-75.
[3] D. Chaum, "Untraceable Electronic Mail, Return Addresses, and Digital Pseudonyms," *Communications of the ACM* **24:2** (1981), pp. 84-88.
[4] DARPA SAFER Program Concept of Operations (CONOPS), prepared by the SAFERlab Project of USC Information Sciences Institute and Cobham, Inc., Version 1.04 (rev 23), September 23, 2011.

[5] Privacy Enhancing Technologies Symposium, http://petsymposium.org/2011/ (and, more generally, http://petsymposium.org/20XY/)

[6] A. Pfitzmann and M. Hansen, "A terminology for talking about privacy by data minimization: Anonymity, Unlinkability, Undetectability, Unobservability, Pseudonymity, and Identity Management," http://dud.inf.tu-dresden.de/Anon_Terminology.shtml

[7] TOR: Anonymity online, https://www.torproject.org/