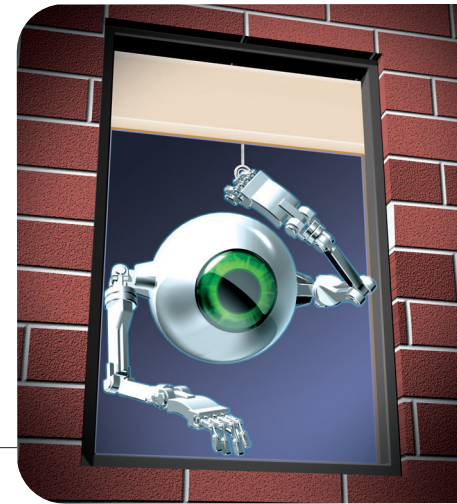


A Framework for Countering Denial-of-Information Attacks

Denial-of-information (DoI) attacks degrade a given user's ability to seek, assimilate, and process information, and are becoming more prevalent due to the Internet's rapid growth. To counter such attacks, the authors' taxonomy provides structure to this area and proposes a model for describing the information space.



GREGORY
CONTI AND
MUSTAQUE
AHAMAD
*Georgia
Institute of
Technology*

Internet users experience denial-of-information (DoI) attacks every day. Our inboxes are swamped with spam containing subject lines that sometimes fool us; our search-engine queries return many irrelevant results; and online auctions are plagued by corrupt database records filled with intentionally misleading keywords. As we attempt to uncover useful data, we waste time sifting through useless information and, in many cases, are actively tricked into taking incorrect actions.

DoI attacks go far beyond simple denial-of-service (DoS) attacks, which attempt to prevent legitimate users access to a service and typically target our machines' processing, memory, and bandwidth resources. DoI attacks include DoS attacks, but also extend to any attacks that intentionally or unintentionally consume human resources or mislead, confuse, or trick users into acting inappropriately or not acting when they should. In short, DoI attacks target humans through their computers.

Although some work in DoI has occurred, no concerted effort currently exists to coalesce and define this field. To address this problem, we propose a taxonomy that categorizes DoI attacks and countermeasures, as well as a framework that's useful for analyzing these attacks' characteristics. We hope these tools will lead to solutions that can begin countering such attacks.

DoI history

Concealing information (steganography) or attempting to make it unreadable without special knowledge (encryption) are legitimate issues related to, but distinct from, DoI. The key difference is that DoI attacks attempt to manipulate or overwhelm us with malicious information, limiting our ability to get the information we seek.

History provides a rich set of

DoI examples that illustrate our long-running quest to find the right information at the right time despite a great deal of noise within the information environment. Examples include the US's failure to predict the attacks on Pearl Harbor on 7 December 1941 and on the World Trade Center on September 11, 2001.

This quest is even more difficult today due to rapid growth in how quickly information is generated and communicated. Recent estimates place the amount of unique information produced worldwide between one and two exabytes per year. Humans exchange some 610 billion emails and annually browse 2.1 billion static pages on the Web. Internet use grew 2,000 percent between 1992 and 2000.¹ In today's information environment, this growth has outstripped our ability to find and process information, despite the increased information-processing capabilities at our command. This problem is nontrivial even when we set aside malicious entities that seek to inject noise into our environment, but for this article, we assume that in a world without malicious entities, traditional information retrieval and overload issues can be worked out.

The information environment

Myriad motivations drive malicious entities to deny us information, including political, economic, and social gain. In this dynamic environment, even determining attackers' identities is difficult. One individual's absolutely credible source is another's propaganda, but we assume there is only one base truth. The intentional and unintentional clouding of our information environment prevents

us from finding, assimilating, and acting on that truth, as several representative examples demonstrate:

- The Google SafeSearch feature, while intending to filter sexually explicit content, actually filters significantly more, including tens of thousands of pages without any such content.²
- In the late 1990s, an e-bookseller purchased the word “the” on a pay-per-click search engine, resulting in thousands of irrelevant responses to search queries.
- Web server content, when filtered at the IP address level, can block access to many other unintentional sites sharing the same hosting service.³
- A Silicon Valley programmer became so enraged by the spam he received, he threatened to torture and kill the offending company’s employees. In an interview after his arrest, he “acknowledged that he had behaved badly, but said his computer had been rendered almost unusable for about two months by a barrage of pop-up advertising and email.”⁴

As a first step to countering attacks, we must develop mechanisms to find the right place in the information space for providing the information we seek.

Limitations in today’s addressing mechanisms and semantic search technology often result in a search returning unwanted information (noise) that conceals the information sought (signal). Consider the signal-to-noise ratio (S/N) of your email inbox; messages you value are the signal, whereas spam is the noise. Other problems you might encounter include the lack of an answer in the information space or an inability to adequately specify a search such that it includes the desired information.

The Observe, Orient, Decide, and Act (OODA) loop model, developed by US Air Force Colonel John Boyd,⁵ can help us analyze the effects of the signal-to-noise ratio on people. In this model, humans repeatedly go through these four steps as they assimilate information and act on it. In a competitive environment, individuals who can execute these steps faster have the advantage. DoI attacks can slow each step—an email user, for example, might take time deciding to delete a given piece of spam with a particularly compelling subject line. Another class of attack can force the user to perform extra cycles through the loop. Email users, for example, must spend time scanning their inboxes, deciding to delete (or not delete) each message and then physically performing the delete operation. Ultimately, such attacks can exhaust peoples’ available resources (such as time for the task) so that the loop halts as they move on to other tasks. A third class of attack includes those that so cloud the information environment that incorrect decisions are made, also forcing extra cycles. We think the OODA loop is useful for quantitatively modeling the specific human impact of DoI attacks.

Although we don’t predict a panacea in the near future, we aim to help tame the information space. To this end, we first present terminology for defining the problem.

Terminology

Our proposed framework has six primary concepts:

- *Node*—a machine information producer or consumer. Each node has finite amounts of processing, I/O, and storage resources available. In most cases, nodes are interconnected through communication channels.
- *Communication channel*—the medium used to transmit a message from a transmitting node to a receiving one. A communication channel has a finite ability to transmit information.
- *Information object*—all individually addressable objects within the data universe. These objects fall into either the signal or noise categories. Malicious entities might either design inappropriate information objects so that they appear as signal or design sought-after information objects such that they appear as noise.
- *Information universe*—the set of all possible nodes, information objects, and communication channels.
- *Human*—a node user. Humans interact with nodes using vision and hearing for input as well as motor skills and speech for output. Each human has available finite amounts of cognitive processing, perceptual capabilities, and memory.
- *DoI*—reduction of the ability of an information node or human to acquire desired information, resulting in deception or loss of resources. When an attack is written to target only a node’s computational resources, we view it as a traditional DoS attack.

Working from these concepts, we created our taxonomy and framework.

DoI taxonomy and framework

Our taxonomy and framework considers both objects and actions associated with the DoI domain. We’ve consulted with domain experts who analyzed the design for completeness and applicability, and observed users in DoI scenarios. Although proving the taxonomy’s absolute validity and completeness is difficult, we conducted 12 months of analysis at Slashdot.org, primarily during 2004 and early 2005, and other metanews Web sites, searching for DoI instances in any form. This let us see patterns in the DoI domain, and each new instance helped us iteratively improve and verify our taxonomy. To further refine it, we reviewed the proceedings from recent antispam conferences and other relevant knowledge-management conferences and verified that we reflected each relevant object and action. Although we can’t verify our taxonomy’s absolute completeness, we do believe we’ve paid due diligence to making it comprehensive and accurate.

Table 1. Denial-of-information attack scenarios.

	SIGNAL (S)	NOISE (N)	S/N	IMPACT
Scenario 1	Low	Low	Parity	Marginal to good ability to find information
Scenario 2	High	Low	Good	Good to excellent ability to find information
Scenario 3	Low	High	Bad	Denial of information
Scenario 4	Very high	Very high	Parity	DoI or processing, I/O, or storage capability exceeded (denial of service)

Overview

Objects in this framework fall into four of the main concepts we defined earlier: information nodes, their interconnecting communication channels, humans, and information objects. Actions include production, consumption, attacks, and defenses. Each information node can be both an information producer and consumer. Humans interact with information nodes via communication channels to search for and, in some cases find, information objects. In this model, humans are also consumers and producers of information, which they feed into or receive from information nodes. Essentially, they work through information nodes to accomplish their tasks.

With our framework, we aim to characterize the components and characteristics in which DoI attacks occur. Our model recognizes three levels of granularity, but only addresses two: the coarse node level and the medium-grain transaction level. We don't attempt to deal with a given piece of information's fine-grained semantic meaning. As an example of this granularity, consider the following questions:

- Node level: Do you trust a given Web site?
- Transaction level: Do you trust a given Web page?
- Semantic level: Do you trust a given fact on a Web page?

We leave the integration of semantic-level techniques for future work.

Categorization of actions

Actions take the form of information attacks and defenses humans initiate, as well as information production or consumption. Attacks target the producer, consumer, or communication channel, exploiting inherent vulnerabilities in humans or machines as they process data. We should put things into perspective here: one individual's defense is another person's attack. Consider the case of a spammer who initiates a wave of emails to sell a product. Some of those targeted people might use a technological defense such as email filtering or, perhaps, a legal defense such as a cease-and-desist lawsuit; what the recipients consider defenses are perceived as attacks from the perspective of the spammer, who's seeking information on potential customers.

Production. Producers are information nodes that provide information to other nodes, either directly or by gathering information from other nodes, humans, or sensors. The information they provide falls along a spectrum of absolute accuracy to absolute inaccuracy. Nodes might intend to provide accurate content, misinformation, or a combination of the two. They might take countermeasures to bypass consumers' protection measures, and in some cases, might deliberately cloak misinformation by surrounding it with accurate information. Such countermeasures could also include modifying information to bypass these protections. In the case of spam, attackers might manipulate subject lines to penetrate the consumer's filtering mechanisms. Producer nodes typically have several outbound channels to consumers.

Consumption. Consumers are information nodes that receive information from producers, due either to direct requests, as with a Web search engine query, or to having information pushed to them, as with email, which they might or might not have requested. Typically, consumers aim to acquire a certain amount of accurate information while filtering extraneous noise. In most cases, the more noise that penetrates their defensive mechanisms, the greater the likelihood of a successful DoI attack.

Attacks. Attacks can be intentional or unintentional, but they always attempt to degrade the quality of information humans receive. This is apparent in a lower S/N. Table 1 illustrates four main attack scenarios.

An ideal Web search by a consumer using a reputable search engine will return results where noise is low and the signal is high (scenario 2). With a malicious node, the results might include pay-per-placement advertisements in response to the query, which would increase the noise (scenario 3). Consumers might incorrectly formulate queries, thereby receiving only marginal results (scenario 1) or, through no malicious action at the producing node, receive incorrect information (also scenario 3). Attacks succeed when they prevent consumers from receiving desired information in the necessary time window or by transmitting noise that appears to be valid information. This usually occurs when humans' information-processing capabilities are overloaded, but can also hap-

Table 2. Denial-of-information attack taxonomy (by target). Each category has an example of an intentional attack.

EXAMPLE						
HUMAN	Processing	Cognition	Memory	Perceptual buffers	Magician's sleight of hand (things moving faster than the eye can see)	
				Short term	Choosing a long password of random characters	
				Long term	Spam from "familiar" sounding source Remembering multiple passwords	
			Cognitive processing		Using advanced technical jargon to exclude newcomers	
	Input	Vision	Blinking advertisements			
		Hearing	Playing very loud rock music outside Manuel Noriega's home			
	Output	Speech	Speaking too softly for eavesdropping			
		Motor	Monitoring keystroke dynamics for authentication			
	MACHINE	Processing			An attacker floods an operating system with inbound packets, causing some packets to be dropped	
		Input			Distributed denial-of-service (DDoS) attack	
Output			Encrypting output to avoid eavesdropping			
Storage		Primary storage (ROM/RAM)	Corrupting a machine's basic I/O system (BIOS)			
		Secondary storage (Hard drive)	Sending network traffic to flood an intrusion detection system's logs			

pen if the capabilities of their automated information-processing hardware and network cause the problem (scenario 4). Both attackers and defenders might take countermeasures to increase the likelihood of their success. We measure attack success by the influence it has on the target's decision-making abilities. In other words, how much change is there in what the target does or doesn't do? Table 2 shows these attacks in more detail.

Defenses. To put this category in perspective, again, note that one person's attack is another's defense. Both attacking and defending nodes can use defensive countermeasures to increase or reduce the likelihood of a successful DoI attack. Defenders use legal, regulatory, moral, cultural, organizational, financial, technological, and perhaps even violent countermeasures to reduce noise or increase the signal in their information environment. An attacker might use similar measures to increase the likelihood of a successful DoI attack. Table 3 lists a taxonomy of such countermeasures.

Because our work focuses on technical countermeasures, Table 4 goes into much greater detail in this area. Each technique attempts to increase the S/N ratio, thereby reducing the likelihood of a successful DoI attack:

- *Filtering* is a widely employed technique that removes

noise from the environment. Current examples include keyword, Bayesian, and collaborative filtering; another example is human-in-the-loop systems that attempt to increase information quality by requiring a human representative to review information before it's entered into a database.

- *Resource-consumption-based* defenses force potential attackers to pay a given amount of a finite resource—typically money, processing, or time—to reduce their attacks' effectiveness. Adaptive and agile systems let humans better acquire information by applying only the appropriate amount of computing resources required for a task. This frees resources for other information countermeasures. Agents shift information-processing requirements from humans to information technology, where computational resources can handle them more efficiently.
- *Meta-information* countermeasures attempt to match externally computed information with a given data set. Examples include data-quality measurements and trust metrics.
- *Trusted computing* attempts to ensure an information node's integrity. If a system is compromised, any information it provides is suspect.
- *Data fusion* attempts to reduce the amount of data presented to a user by merging data from disparate sources to produce useful information and lessen the cognitive load.

Table 3. Denial-of-information defense taxonomy (the big picture).

TYPE	COUNTERMEASURE	EXAMPLE
Legal	Lawsuits New laws	Recording Industry Association of America sues small-scale MP3 distributors New York State's no-call database
Regulatory	Government regulation	US Federal Communications Commission approves media consolidation
Moral	Public relations campaign Code of ethics	Movie industry campaign against movie piracy A state bar association imposes code of ethics on its members
Cultural	Communities	Exclusion from the group
Organizational	Topical counter-DoI groups	Antispam consortium founded
Financial	Increasing cost of DoI operations	Cost of postage increases
Violence	Violence against DoI perpetrators	Sending threatening messages to telemarketing firm
Technology	See Table 4	

- *Database keys/indices* place emphasis on quality. An example is a Web search engine that bans a Web site that attempts to use inappropriate keywords to gain better placement.
- *Online communities* let individuals interact about mutually interesting topics. These communities can apply techniques such as excluding counterproductive members and enforcing cultural standards to provide highly relevant information to group members.
- *Source-evaluation* techniques attempt to identify authoritative information sources as well as those that generate noise to provide higher-quality information. This evaluation might include challenge-response techniques or a test to determine whether the consumer is a human or a machine. This category also considers source anonymity, which is a double-edged sword. In some cases, it can lower information quality if the source believes it will remain anonymous. In other cases, anonymity will improve information quality by letting individuals speak freely.
- *Structuring data.* Improved data structuring allows more efficient retrieval of relevant information. Technologies such as XML embed knowledge into information and subsequently improve encoding.
- *Restricted connectivity* helps organizations prevent or reduce opportunities for data corruption. For example, an “air-gapped” wide area network that provides connectivity to only organizational machines and no open Internet access will reduce direct data corruption to insider threats only.
- *Translating data* provides information that better meets human needs. XML provides a rich set of tools to support transformations. Other examples include middleware that converts between disparate formats and language translation systems such as BabelFish.com.
- *Human-computer interfaces* give users an efficient way to interact with their machines. An excellent example is information visualization, which reduces humans’ cognitive burden by presenting data in a more efficient

graphical representation. Humans can then make better decisions or discover new information.

- *Data protection* via systems that detect changes in information or intrusions help build trust in information from a given source. Examples include the TripWire system, which can detect changes in an information system’s files or the use of public-key cryptography to protect a given document’s integrity. More recently, detection systems are used in conjunction with “self-healing” technologies to repair damage from an attack. The related topic of information preservation seeks to safeguard information that might degrade or change over time. An excellent example is Web page caching by Google, or the Internet WayBack machine, which stores snapshots of Web pages.
- *Locating data* assists humans in initially finding information as well as locating information that they’ve located before. An example is a Web browser that provides bookmarks and history functions.

Examples

Now that we’ve developed this DoI taxonomy and framework, we can look at two examples of its application: spam and a news Web site.

Spam. Spam is a prime example of a DoI attack. The human mail recipient wants to receive emails (information objects) that contain content of interest (signal) and exclude others (noise). The spammer’s goal is to penetrate the countermeasures the recipients have put in place to promote their product, service, or agenda. Spammers typically send blanket email broadcasts to hundreds of thousands, perhaps millions, of addresses. This process results in inboxes filled with irrelevant email. As we mentioned, humans must spend time cycling through the OODA loop to find messages that interest them.

For our example, let’s assume that a given human, Alice, has received 47 emails, of which six are of actual interest. The remaining 41 constitute noise. Of these 41 emails, 35

Table 4. Denial-of-information defense taxonomy (technological).

		EXAMPLE
Filtering	Collaborative filtering Filtering algorithms Human-in-the-loop	Slashdot.org Link analysis, proximity searches, page-rank computation Yahoo's human reviewers
Resource consumption	Money Time Memory Processing Bandwidth Resource allocation	Charging fee for each email Sending artificial-intelligence-generated email responses to spammers Adding more RAM to a workstation Systems that enact a small computational charge to prevent abuse Data-compression algorithms Agents Adaptive/agile systems
Meta-information	Data quality measurements Trust metrics Currency	Organizational guidelines on data quality management Advogato.org Googlebot/Web crawlers
Data fusion	Reduction and merging of data	Air traffic control, weather forecasting
Online communities	Exclusivity Cultural standards of behavior	Orkut.com Slashdot.org
Source evaluation	Anonymous input Community input Authoritative input Hardware/software trust Authentication/testing	The Freenet project Communities of trust and reputation systems Your college professor's list of recommended links Trusted computing Use of Completely Automated Public Turing Test to tell Computers and Humans Apart (CAPTCHA) to determine if a consumer is human
Structuring data	Keywords Improved encoding	Including keywords with academic papers XML
Restricted connectivity	Air-gapped network P2P communities	Power plant control network Waste by WinAmp.com
Translating data	Converting data to more useful forms	XML transforms
Human computer interface	Presenting data Interface design	Information visualization Effective location of interface elements
Data protection	Integrity protection Preservation	Tripwire Historical archives, such as the Internet WayBack Machine and Google cache
Locating data	Keeping things found Searching Indices Naming	Bookmarks, browsing histories, shortcut links Google Yahoo Providing unique names for information objects

are “obvious” spam, and the other six are convincing enough that Alice will need to examine them more closely. Alice's strategy to review and process email is to

- scan email headers,
- delete the obvious spam, and
- read the remaining emails, then decide which to delete or keep.

Given this strategy, Alice works with the mail client running on her PC (information node) to read her mail. She'll cycle through the OODA loop 47 times during this process. Each step in the loop (observe, orient, decide,

and act) requires her time and other resources. She'll spend significantly more time analyzing the convincing spam than she will deleting the obvious spam. If Alice's wasted resources exceed those she has available for her current tasks, the DoI attack has been successful.

Frustrated by the wasted time, Alice employs several countermeasures to help protect her. She attempts to deter spam by getting a new email address and sharing it with only her personal friends (restricting connectivity). She explicitly doesn't place it on the Web (addressing and naming). She builds rules that examine her incoming mail and moves messages addressed to her old email address to a lower-priority folder that she scans only occasionally (fil-

tering). She plans to delete the old address after several months, when she feels confident that the new address has been disseminated to the right people (exclusivity and restricted connectivity). She also considers switching mail clients to a text-based program that runs quickly and has efficient keyboard shortcuts (interface design).

Alice is heartened to hear that several laws have been passed making spamming illegal and that metered mail is being explored to make spamming less attractive as a business (financial and legal). She is concerned, however, that the fee for email being considered, albeit small, could lead to unwelcome charges.

In this scenario, spammers aren't static enemies. They employ counter-countermeasures of their own to keep up with their target's defenses, a kind of cold-war one-upmanship. They attempt to bypass filters by using compelling but innocuous subject lines and spoofed return addresses and names, as well as creative ways to construct their messages that are easy for humans to understand, but difficult for a computer to analyze.

Metanews Web site. DoI attacks also affect groups of individuals and organizations, as we can see by examining the metanews Web site Slashdot.org. Site users submit short articles for viewing and comment. If left unprotected from DoI attacks, this site's information quality would degrade quickly. In comparison, consider the low quality of information available on unprotected Usenet news groups and unmoderated mailing lists. Slashdot employs various DoI countermeasures to great effect. The Slashdot staff (human-in-the-loop) reviews all story submissions for relevance and interesting content, selecting only the best for publication. Once published, users comment on the story. A randomly selected subset of users can rate each comment's quality (collaborative filtering). Comments submitted by registered users (authentication) receive higher initial ratings than those submitted by anonymous users (anonymous input). Readers can easily filter stories at any rating threshold from -1, meaning very poor quality, to 5, meaning excellent quality (data quality metrics, filtering algorithms, and interface design). Standards of acceptable usage have evolved over time and contribute to the site's unique culture and quality (cultural standards of behavior). Long after a story is published, it's still available (preservation) and can be found via a search mechanism on the Web site (searching) or with a URL linking directly to the story (keeping found things found).

DoI is a critical problem, but current research is scattered and application-specific. Our definitions, taxonomy, and model provide some clarity and structure, but only as a starting point for holistic research in countering DoI attacks. Much work remains open in this field.

Our next step is to define, implement, and validate

prototypes that counter DoI attacks. Our validation plan must include empirical measurements that demonstrate such prototypes' effectiveness in helping people with realistic tasks. In particular, these prototype systems must reduce our perceptual, cognitive, and motor burden. An information infrastructure that supports the exchange of meta-information is necessary for implementing such prototypes. Attack-resistant trust metrics, combined with collaborative filtering and trust communities, might prove useful models to emulate. Finally, tools that better understand the semantic level of information would allow improved resolution of information presented to humans. A markup language could prove useful in this area. □

Acknowledgments

We thank the anonymous reviewers for their constructive comments and suggestions as well as the staff of IEEE Security & Privacy. In addition, we thank the Georgia Tech Information Security Center (www.gtisc.gatech.edu) and the US Military Academy's Information Technology and Operations Center (www.itoc.usma.edu) for their ongoing support. US National Science Foundation Information Technology Research award 0121643 partially supported this work. The views expressed in this article are those of the authors and do not reflect the official policy or position of the US Military Academy, the Department of the Army, the Department of Defense, or the US government.

References

1. "How Much Information Research Study," School of Information Technology, Univ. of California at Berkeley, 18 Oct. 2000; www.sims.berkeley.edu/research/projects/how-much-info/.
2. B. Edelman, *Empirical Analysis of Google SafeSearch*, law report, Berkman Center for Internet & Society, Harvard Law School, 14 Apr. 2003.
3. B. Edelman, *Web Sites Sharing IP Addresses: Prevalence and Significance*, law report, Berkman Center for Internet & Society, Harvard Law School, 12 Sept. 2003.
4. "Man Arrested for Spam Rage," *Wired*, 21 Nov. 2003; www.wired.com/news/culture/0,1284,61339,00.html.
5. G. Hammond, *The Mind of War*, Smithsonian Institution Press, 2004.

Gregory Conti is an assistant professor of computer science with the US Military Academy and a PhD student at the Georgia Institute of Technology. His areas of expertise include network security, information visualization, and information warfare. Conti has a BS from the US Military Academy and an MS from Johns Hopkins University, both in computer science. Contact him at conti@acm.org.

Mustaque Ahamad is a professor of computer science at the Georgia Institute of Technology. His research interests are in distributed operating systems, computer security, and fault-tolerant systems. Ahamad has a BE in electrical and electronics engineering from the Birla Institute of Technology and Science, Pilani, India, and an MS and PhD in computer science from the State University of New York, Stony Brook. Contact him at mustaq@cc.gatech.edu.