# Secure Dynamic Adaptive Traffic Masking

Brenda Timmerman
California State University Northridge

*Abstract*

*As public internetworks are increasingly used for secure communications, the need grows for end-to-end protection from traffic analysis. The additional protection of Traffic Flow Confidentiality can be detrimental to performance when padding is used to mask traffic patterns. Traffic masking policies that are responsive to system service requirements can improve performance, but secure adaptive traffic masking has to balance performance requirements with system protection requirements. This paper addresses the information leaks that result from adaptations in security mechanisms.*

## Introduction and Background

Traffic flow confidentiality (TFC) is concerned with hiding communication patterns that, if exploited, could reveal or compromise sensitive information. Sources of traffic flow information that need to be protected are frequency and length of transmittals, origin/destination traffic patterns, and protocol headers [6, 7]. TFC is becoming more important as government agencies and private companies are moving away from private networks and using open data networks to meet their needs. While the security of open data networks is a concern, designers of network security are faced with an explosion of worldwide communications that includes increased data rates, universal connectivity, new services, and higher standards for performance. In such environments TFC can meet the growing need for protection from traffic analysis, but can be expensive because traffic masking involves the use of padding. Secure dynamic adaptive traffic masking (S-DATM) contributes to a global vision, providing the capability of operating in a commercial environment via traffic protected by appropriate levels of TFC with minimal impact on other traffic.

Traditionally TFC has been provided by bulk encryption between protected sites on dedicated private networks that are no longer practical for wide scale internetwork use [3]. Public networks are not only cheaper and more reliable, but have capabilities for end-to-end operational security, including confidentiality, integrity, authentication, and some privacy. Changing protection needs occur as secure hosts move outside of protective gateways, but still require operational security. Some internetwork users need the added privacy of TFC even though it is frequently considered too detrimental to performance to be considered practical [13].

It has been recognized that security mechanisms that are adaptive to changing conditions in their environment can reduce costs and improve both system and application performance [2, 4, 5, 14, 9]. It is usually acknowledged that such adjustments can cause leaks of information about the protected systems.


1999 New Security Paradigm Workshop 9/99 Ontario, Canada

Traffic masking schemes that are adaptable, i.e. make adjustments for changing network conditions or application service requirements, can cause statistical anomalies in the masked traffic patterns which are susceptible to *statistical anomaly detection* [8]. Schemes that are not adaptive are independent of original traffic characteristics and therefore not vulnerable to anomaly detection, but they do not adjust to improve performance or efficiency. Efficiency is the ratio of original traffic to masked traffic and approaches the value one as it increases. Schemes that do make adjustments for improved performance and efficiency can create traffic patterns that, while different from original traffic may imply characteristics of the original traffic that can be analyzed. Other policies have been proposed for traffic masking that outline technology for dynamic adjustments in response to changing rates of original traffic [12]. While these dynamic adjustments are shown to reduce the costs of traffic masking, the proposed policies do not address system protection, in particular, the reduction in protection resulting from the dynamic adjustments. In order to use matrix mathematics, they model the network as a matrix, requiring constraints on their models that do not scale to internetworks, such as fully connected nodes with global synchronization.

## Overview of S-DATM

In this paper, we present dynamic adaptive traffic masking techniques for secure TFC mechanisms that have the capability to dynamically adjust to changing network conditions. The techniques are based on *SMD*, a security model that precisely specifies the security requirements for S-DATM mechanisms within a network environment [10]. The goals of *SMD* are to satisfy system security requirements, minimize padding costs, and meet the throughput and delay requirements of the original traffic. In addition, *SMD* precisely specifies the relationships between protection and application performance requirements and allows trade-offs between the two that meet system security policies. The proposed techniques include *statistical anomaly masking* (SAM) that uses statistical methods to detect and prevent leaks of inference information that may occur when dynamic adjustments are allowed. [11]

The criteria for S-DATM techniques are that they satisfy system security requirements, minimize padding costs, and meet the performance requirements of the original traffic. The goals of S-DATM mechanisms are: to meet system protection requirements by masking well defined traffic characteristics in a systematic manner; to reduce processing and storage overhead of adaptive masking schemes; to improve efficiency, system performance, and the performance of the protected applications; to allow secure dynamic trade-offs between the costs of protection and application performance requirements; and to provide end-to-end protection.

S-DATM includes techniques to prevent statistical anomaly detection in adaptive masking schemes. Data is reduced to statistics that are kept as frequency tables, means, covariances, and correlation coefficients. Keeping statistics is preferable to storing all historical data because statistics require a minimal amount of storage and processing time and still provide sufficient information to be interpreted for anomaly protection. S-

DATM schemes may have to do real-time calculations based on events in system environments. Using statistical techniques, it is not necessary for schemes to store or process extensive information about past behavior. Data can be store in summary statistics in a simple additive fashion. Accumulated data about network traffic can be stored as an exponentially weighted sum of changes in traffic characteristics. These statistics can be weighted so that either recent history or far past history is dominant.

In S-DATM traffic characteristics determine whether or not a protocol data unit should be transmitted during the minimum unit of time. The decision is based on calculated throughput, inter-arrival times, and burst size statistics over a period of recent history. (the sample period). SAM employs these statistics to evaluate how closely the current characteristics conform to those in a SAM profile. The profile consists of a collection of target statistics for these traffic characteristics, and tolerance levels for the statistics. If recent history statistics are sufficiently close to those in the profile, the decision to transmit in a time slot is based on the original traffic. Over intervals of time, the module's output is constrained by the tolerance levels so that the statistics of the outgoing traffic characteristics stay sufficiently close to the profile. If original traffic is queued, it is subsequently output introducing additional delays. Padding occurs when there is no original traffic on the queue and the profile calls for a transmission. The queue's length can be a consideration in determining the rate of output. The results are that adjustments can be made without creating statistical anomalies.

The scheme can adapt to environmental changes (such as an event that increases the rate of output) by adjusting the profile's critical values that determine the tolerance levels. Allowable adjustments and interval size depend on the desired degree of protection. Longer intervals and fewer, or smaller, adjustments provide greater security. When it is determined that an adjustment is needed, the scheme checks to see if the adjusted output results in statistics that are sufficiently close to those in the profile to be accepted as expected. If not, it then checks directives from the security policy for allowable adjustments to the parameters and tolerance levels that determine the range of acceptable behavior. Adjustments to these critical values change the profile of the masked traffic. If an adjustment is not allowed by the security policy, it is not made.

The scope of the research described in this paper is necessarily limited to masking patterns of frequency and length using statistical techniques. However, masking origin and destination patterns is also essential for privacy. Therefore, to effectively provide protection from traffic analysis, S-DATM mechanisms must be part of a system of TFC protection that includes anonymity for senders and receivers. Source and destination ambiguity are the subject of a future paper. In addition, this research does not look at covert channels that can be introduced by traffic masking. This is addressed somewhat in [12] and is a worthy subject of additional research.

## Implementing a Framework for S-DATM

Figure 1 depicts a framework of an S-DATM module within SMTP. The module is placed below the User Agents (UAs) and above the Message Transfer Agent (MTA), in

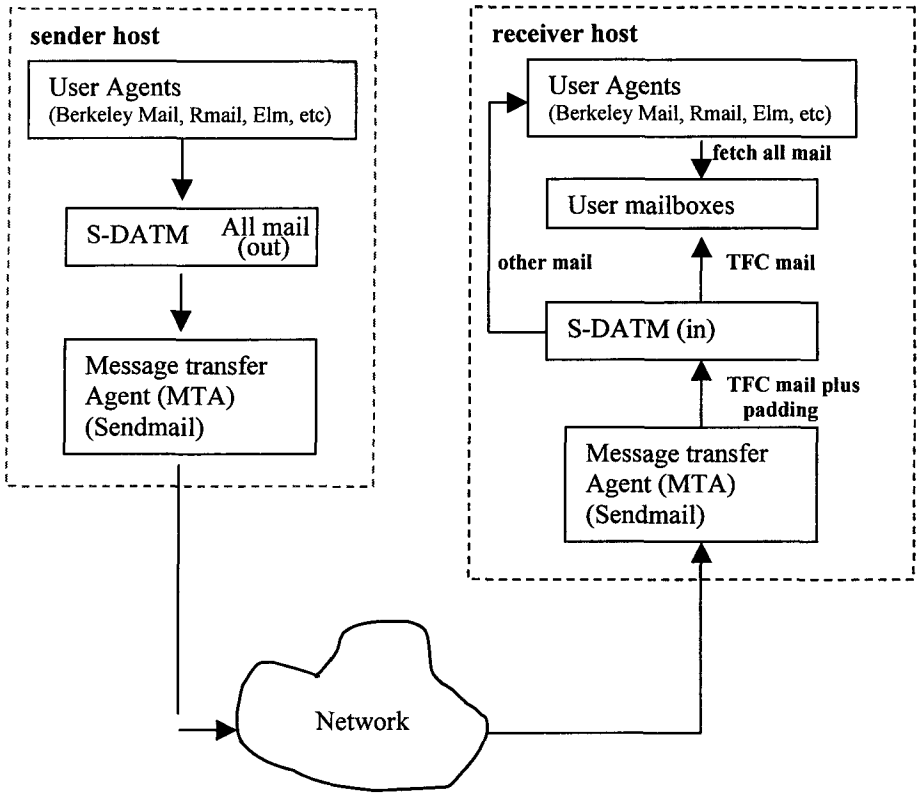


**Figure 1 S-DATM Implemented with Simple Mail Protocol**

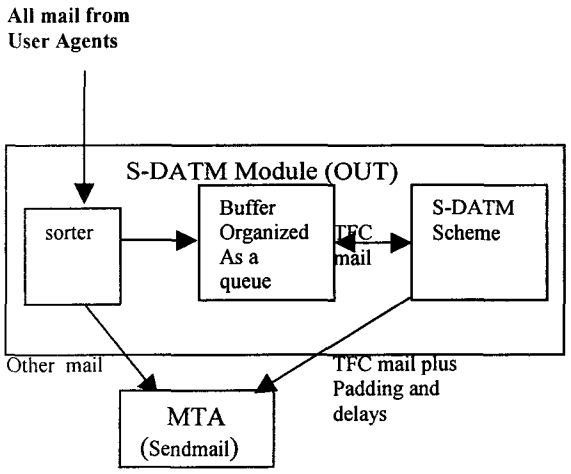


**Figure 2. S-DATM Module (in)**

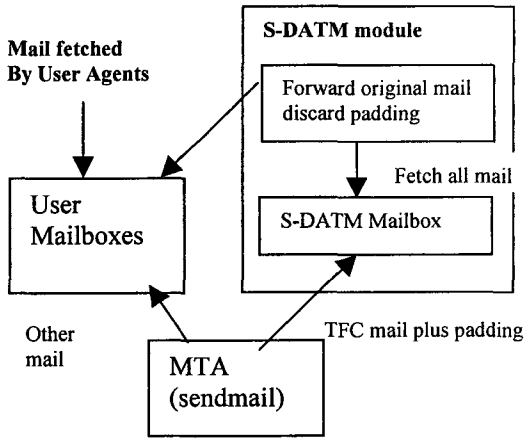


**Figure 3. S-DATM Module (in)**

this case, Sendmail, transparent to both. All mail messages from the UAs are forwarded to the S-DATM module which sorts the mail, forwards the unprotected mail to the MTA, and places the protected messages on a queue. (See figure 2) It recognizes the protected

mail by the addresses of the targeted receiving hosts. The S-DATM scheme decides when the protected messages are removed from the queue and forwarded to the MTA and determines when padding is forwarded to the MTA in the form of bogus messages for the receiving host. The MTA is unaware that these messages are bogus. Encryption of all mail to the targeted hosts, original and padding, is assumed. The module for incoming mail is shown in Figure 3. The outgoing module addresses all messages to a special userid in the receiving host when they are forwarded. Upon arrival, the bogus messages are dropped and the protected original mail is forwarded to the receiving users' mailboxes.

We implemented a prototype traffic masking module above SMTP in the TCP/IP protocol suite in order to test the introduced S-DATM techniques. We simulated a mail User Agent (UA) in the C programming language that sent messages from the host workstation to other workstations in the local area network. All messages went through the local mail server, a SUN Sparc Server. The host station had a HP-UX operating system and the mail server had a SunOS operating system. The simulated UA sent collected patterns of real mail message traffic to the masking module. The use of a simulated UA made the realization and evaluation of the techniques feasible in the real network environment because it allowed complete control over the input to the module. Traffic characteristics such as throughput and inter-arrival times of input mail messages were controlled and could be adjusted. Traffic patterns could be repeated for scheme comparisons

## SAM Techniques

The basic approach of SAM is That it provides a means to compare the short term behavior of masked traffic output with the profile and prevent any adaptations that cause the short term behavior to be sufficiently different from the profile to cause a breach of security, i.e. statistical anomaly detection. Each potential transmittal (output) is compared with a profile of traffic behavior and only transmitted if it falls into the ranges of behavior considered acceptable by the profile. Unacceptable behavior is both short term behavior that does not appear in the profile and the absence of profile behavior in the short term histories.

With SAM, the pertinent characteristics of masked traffic, i.e. burst size, inter-arrival times, and throughput, are manipulated to satisfy the SAM profile that describes expected masked traffic behavior. The SAM profile allows anticipated adaptations to be accepted as expected behavior. Through the use of padding and introduced delays, SAM outputs traffic with patterns of behavior that satisfy the profile, and will accommodate changes in the original traffic when they occur.

Depending on the application, SAM profiles keep statistics such as frequency distributions, means, covariances, and correlation coefficients on the traffic characteristics of burst size, inter-arrival times, and throughput. Each time new data is collected, before the newly determined data is assimilated into the recent past's summary

statitic, the values in the statistic are aged by multiplying them by the exponential decay factor. The rate of decay must allow the masked traffic to respond to rapid changes in the original traffic behavior by recognizing the relative normality of recent behavior.

SAM parameters are tunable to specific system needs. Optimal values for these parameters are determined by the environment of the system including the system's security policy and the relative importance of performance and efficiency. The *half-life* of a summary evaluation statistic is determined by the value of the $r$ exponent in the formula that can be generalized as

$$b_{k+1} = 2^{-rt} \bullet b_k + D_k$$

where $b_k$ is the summary evaluation statistic and $D_k$ is the data collected in the $k^{th}$ time interval. The value of $r$ determines the decay rate of the statistic and can be set at what is appropriate for the environment. The half-life of each $b_k$ is $n$ collection periods if $r$ is set to $1/n$. Since SAM schemes collect data at each minimum time unit, $t$ has the value 1. The number of time periods that constitute a short-term history of behavior can be set through the specification of $r$. For example, when the half-life is set for 10, $r = 1/10$, the data collected 20 time periods previously has $1/4^{th}$ the influence of the most recently collected data. However, if the half life is set for 5, $r = 1/5$, the data collected 20 time units previously has only $1/16^{th}$ the influence of the most recently collected data. [11]

The selection of intervals for the frequency distribution is important. Interval size and number are critical in determining the correct bounds for each characteristic. When it is desirable for a characteristic to be in an interval then that interval is included within the bounds on the relative frequency. The values of the mean characteristics in the profile are dependent on environmental factors (as are *which* characteristics are included in the profile). These values determine the short-term history of the traffic. They reflect the relative importance of protection, efficiency, and performance. How far the short-term history mean values of traffic characteristics are allowed to vary from the profile values are determined by this critical value which must meet security policy specifications.

## Prototype Overview

In the prototype model, the profile, a collection of statistics about the traffic, is represented as a *profile structure*, containing a *profile characteristic structure*, an *event structure*, and *a frequency structure*. (Figure 4) The *profile characteristic structure* is an array of tuples where each tuple consist of the decay rate, the target mean value (measured over the sample period determined by the half-life), and the tolerance level of the variance from the target mean value for a traffic characteristic. The *event structure* is an array of events that could cause an adjustment in the characteristic structure. The *frequency distribution structure* is an array of upper and lower bounds for permissible relative frequencies of the throughput, inter-arrival times, and burst size.

The *current state structure* shown in Figure 4 consists of the current *input from the masking scheme* in the form of a logical value, the *queue length*, a *current state*

*characteristics structure*, an array of means measured over the sample period (the short-term history) of the characteristics in the profile, and a *relative frequency structure*, an array of relative frequencies for values of characteristics in the profile.

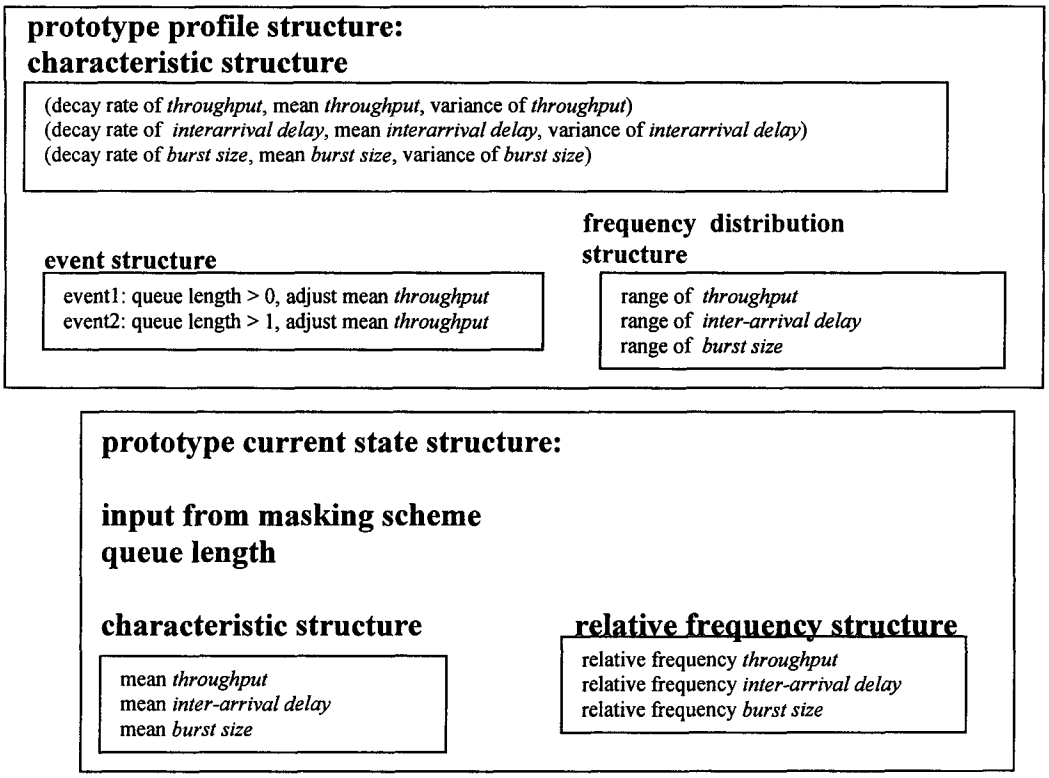


**Figure 4. Profile and Current State Statistics**

## Evaluation

The evaluation of S-DATM schemes, i.e. traffic masking schemes that use S-DATM techniques, includes measuring and comparing the protection they offer, their performance, and their efficiency. Protection is the masking of the characteristics of original traffic and can be measured precisely as a summary statistic. Performance includes introduced delay, inter-arrival delay, and throughput of the original traffic. Efficiency is the ratio of masked traffic to original traffic. A traffic masking module should satisfy the protection and performance requirements of the applications whose traffic it is masking. It should consume at most only enough bandwidth to meet these requirements. Introduced delays should meet the delay and throughput requirements of the original traffic

We determined the relative protection of masking schemes by measuring the correlation between characteristics of the output masked traffic and characteristics of the original traffic. The *correlation coefficient* [1] is a measure of the degree of linearity between two sample distributions. Its value lies between +1 and -1. A value near +1 or -1 indicates a high degree of linearity between the two distributions, and a value near 0 indicates a lack of such linearity. A positive value of the coefficient indicates that both distributions tend to increase and decrease at the same time, while a negative indicates that as one

distribution increases, the other decreases and vice versa. Any value under .3 is considered to indicate a doubtful correlation. A desirable correlation between the current statistics and the real traffic is close to zero, but system security policy specifications determine how close to zero is satisfactory. The closer the correlation coefficient is to zero, the more independent the masked traffic is from the original traffic. For sample input (original traffic) distribution and a sample output (masked traffic) distribution, the sample correlation coefficient is:

$$r = \frac{\sum_{i=1}^{k}(x_i - \mu_x)(y_i - \mu_y)}{\left(\sum_{i=1}^{k}(x_i - \mu_x)^2 \sum_{i=1}^{k}(y_i - \mu y)^2\right)^{1/2}}$$

where $x_i$ and $y_i$ are the output of the masked traffic and the original traffic respectively at time $i$. The system security policy determines how large a value of $r$ is acceptable.

## Comparison of Masking Schemes

Table 1 compares the performances of the nonadaptive and adaptive random schemes with that of a similar S-DATM scheme. The performance is measured by the average introduced delay per message. Table 1 also compares the efficiencies and total padding of the same schemes. The nonadaptive scheme randomly transmits a message (original or padding) every two seconds. (For a greater challenge, we chose this nonadaptive scheme for comparison because it had the best combination of performance and efficiency of the nonadaptive schemes we tested.) For this scheme the average introduced delay per message is 30 seconds. The adaptive random scheme transmits a message every two seconds unless there is an original message left on the queue waiting to be transmitted. In that case, it transmits a message every two seconds until there are no more messages on the queue and it then returns to random mode. The average introduced delay of the adaptive random scheme is 5 seconds, a reduction of about 80%. The efficiency drops from .62 for the nonadaptive scheme to .47 for the adaptive scheme with an increase of 139 padding messages.

Comparing Performance with Introduced Delay, Efficiency and Padding

| *Type of scheme* | *Average introduced delay per message* | *efficiency* | *Total padding* |
|---|---|---|---|
| Nonadaptive random | 30 secs | .62 | 170 |
| Adaptive random | 5 secs | .47 | 309 |
| S-DATM | 17secs | .46 | 322 |

**Table 1: Compares the average delay per message and compares the cost in terms of efficiency and additional padding messages of the same schemes.**

In the S-DATM scheme, a message is transmitted every two seconds as long as the throughput is less than 15 messages per minute with a burst of 30 messages sent every 120 seconds if the average introduced delay per message exceeds 10, or a burst is

scheduled randomly otherwise. The average added delay per message is 17 seconds, a reduction of about 43% from the nonadaptive scheme. The efficiency is .46 with an increase of 152 more padding messages than the nonadaptive scheme.

From these statistics alone, it would appear that the adaptive random scheme is the best. Previous work in the area of traffic padding [9, 12] goes only this far in its evaluations. However, Table 2 has important additional information. It lists the correlation coefficients of the three scheme's output correlated with the original traffic. The nonadaptive scheme and the S-DATM scheme show no correlation. However, the correlation coefficient of the throughput and burst size of the adaptive random scheme with the original traffic indicates its output *is leaking more information about the original traffic than the output of the other two schemes, a significant decrease in protection.* The S-DATM scheme has a 43% decrease in introduced delay over the nonadaptive random scheme, and the correlation coefficients with the original traffic's throughput and burst size are small enough to indicate no correlation. *S-DATM improves the performance by a considerable margin without decreasing protection.*

Comparing Protection with Correlation Coefficients

| *Type of scheme* | *throughput* | *burst size* |
|---|---|---|
| Nonadaptive random | 0.011 | -0.090 |
| Adaptive random | 0.721 | 0.595 |
| S-DATM | 0.028 | 0.042 |

**Table 2: Compares the correlation of the three schemes output with the original traffic**

In terms of relative protection, S-DATM now appears to be the best scheme. Its performance in terms of inter-arrival times of the messages satisfies the performance requirements of unprotected e-mail as can be seen by Table 3 which shows that the average inter-arrival time of the original messages when the traffic is masked by S-DATM is sufficiently close to that of the original messages when the traffic is not masked. In addition, the correlation coefficients of the S-DATM scheme's output with that of the original traffic (from Table 2) indicates they are not correlated, and it offers a 40% increase in performance over the nonadaptive random scheme (from Table1).

## Trade-off Analysis

The S-DATM security model allows protection trade-offs both in implementation phase and in dynamic adjustments. Performance can be improved by basing the probability of dynamic adjustments on traffic characteristics or network conditions. Fine-tuning of parameters can result in improved performance and efficiency. The less padding a scheme produces in relation to original traffic, the more efficient it is, but efficiency is usually achieved at the expense of protection and/or performance. For example, modifications are possible to improve efficiency, such as basing the probability of adjustments on the throughput of the original mail. However, these adjustments can reduce the degree of protection, by leaking information about this throughput. Thus

adaptive traffic masking schemes can improve performance and efficiency, but are constrained by the fact that adjustments to accommodate changes in the original traffic can leak information about the original traffic. For each individual system, a careful study of trade-offs between protection and improved performance or efficiency is necessary. Each system's security policy should determine if the gains are worth the lowered degree of protection.

Performance of Original Traffic

| Type of traffic | Average inter-arrival delay |
|---|---|
| Original traffic unmasked | 6.3 seconds |
| Original traffic masked by S-DATM | 6.4 seconds |

**Table 3: Compares the inter-arrival delay of masked and unmasked traffic**

Table 4 illuminates trade-offs among protection, efficiency, and performance for four masking schemes, a nonadaptive interval scheme, an adaptive interval scheme, and two S-DATM schemes that are based on intervals.[1] Table 4A compares the average introduced delay per message, the efficiency and the total padding for the four schemes. Table A shows a reduction of the average introduced delay from 20 seconds per message for the nonadaptive scheme to one second for the adaptive scheme and 8 and 12 seconds for the S-DATM schemes. The improved performance of the adaptive and S-DATM schemes is accompanied by a smaller, but significant, decrease in efficiency and increased padding. For these schemes, the decrease in efficiency is about 18-25% while the increase in performance is about 90-95%. It would appear that the trade-off for improved performance was reasonable. However, for the adaptive interval scheme, Table 4B shows a high correlation with original traffic indicating a reduction in protection that would probably not meet the specifications of system security policies, except under extreme conditions when the priority of performance was much higher than that of protection. Both of the S-DATM schemes show improved performance with no loss of protection, showing no correlation with the original traffic. The S-DATM scheme that is tuned for performance, shows better performance at the expense of efficiency and increased padding.

For applications with stringent performance requirements, the probability distribution of output from an S-DATM module profile can be based on the peak rate of the application, assuring minimal buffering of the original traffic. For applications with less stringent requirements such as file transfers, the probability distributions of the output can be

[1] The nonadaptive interval scheme transmits three mail messages every ten seconds. The adaptive scheme transmits three mail messages every ten seconds unless the queue length exceeds zero, in which case it increases the number of transmittals, up to ten. The two S-DATM schemes increase the number of transmittals both randomly and when the introduced delay would otherwise exceed what is allowed by the profile. The first S-DATM scheme is tuned for performance and the second S-DATM scheme is tuned for more of a balance between performance and efficiency.

adjusted for smaller, but more efficient output that may cause an increase in delay, but consume less bandwidth (assuming the same level of protection).

Table 4A: Introduced Delay, Efficiency and Padding

| *Type of scheme* | *Average introduced delay per message* | *efficiency* | *Total padding* |
|---|---|---|---|
| Nonadaptive interval | 20 seconds | .51 | 261 |
| Adaptive interval | 1 second | .37 | 466 |
| S-DATM performance | 8 seconds | .31 | 625 |
| S-DATM balance | 12 seconds | .40 | 415 |

Table 4B: Correlation Coefficients

| *Type of scheme* | *throughput* | *bursts* |
|---|---|---|
| Nonadaptive interval | 0.000 | 0.000 |
| Adaptive interval | 0.885 | 0.941 |
| S-DATM performance | -0.001 | -0.016 |
| S-DATM balance | 0.007 | -0.055 |

**Table 4: Table 4A compares average delay per message, efficiency, and total padding of four schemes. Table 4B compares the correlation coefficients of traffic characteristics of the same schemes with original traffic.**

## Conclusions

Data collected from the prototype allowed the evaluation of S-DATM techniques, measuring the protection they offer, their performance and their bandwidth consumption. Currently employed and other proposed technology for TFC does not adequately address the cost of protection or system performance requirements, and they do not consider system protection needs and the reduction in protection that results from dynamic adaptation. We have shown that S-DATM schemes can improve performance and efficiency by allowing dynamic adjustments without a loss of protection. We have also shown that they can be tuned to allow secure trade-offs between the costs of protection and application performance.

## Reference List

[1] A. O. Allen. *Probability, Statistics, and Queuing Theory with Computer Science Applications*. Academic Press, Inc., 1990

[2] L. Badger. Providing a flexible security override for trusted systems. In *Proceedings of Computer Security Foundations, Workshop III*, Franconia, NH, June 1990.

[3] P. Baran. On distributed communications, vol. ix. In *Security Secrecy and Tamper Free Considerations, Memo RM-3765-PR*. Rand Corp., Santa Monica, CA August 1964.

[4] R. Browne. Mode Security: An infrastructure for covert channel suppression. In Proceedings of the IEEE Symposium on Research in Security and Privacy, Oakland, CA, May 1994.
[5] M. P. Herlihy and J.M. Wing. Specifying security constraints with relaxation lattices. In Proceedings of Computer Security Foundations Workshop II, Franconia, NH, June 1989.

[6] ISO. Information Process Systems - Open System Interconnection - Basic Reference Model - ISO 7498. American National Standards Association, Inc., New York, NY, 1984.

[7] ISO. Information Process Systems - Open System Proposed Draft Addendum 2 ISO 7498. American National Standards Association, Inc., New York, NY, 1988.

[8] B. Mukherjee, L.T. Heberlein, and K.N. Levitt. Network intrusion detection. *IEEE Network Magazine*, 8(3):26-41, May-June 1994.

[9] R.E Newman-Wolfe and B.R. Venkatramen. Performance analysis of a method for high level prevention of traffic analysis. In *Eighth Annual Computer Security Applications Conference*, San Antonio, Texas, November- December 1992.

[10] B. Timmerman. A security model for dynamic adaptable traffic masking. In *Proceedings of New Security Paradigms Workshop*, Langdale, Cumbria, United Kingdom, September 1997.

[11] B. Timmerman. Secure Dynamic Adaptive Traffic Masking for Traffic Flow Confidentiality on Internetworks. PhD thesis, University of Southern California, Los Angeles, California, 1997.

[12] B. Venkatramen. Prevention of Traffic Analysis and Associated Covert Channels. PhD thesis, University. University of Florida, Gainesville, Florida, 1994.

[13] V.L. Voydock and S.T. Kent. Security mechanisms in a transport layer protocol. *Computer Networks*, 8, 1984.

[14] D. Weber. Security policies for Army tactical $C^2$ systems. Technical report, Odyssey Research Associates, Inc., Ithaca, NY, 1998.