A Security Model for Dynamic Adaptive Traffic Masking*

Brenda Timmerman†
Information Sciences Institute
University of Southern California
Marina del Rey, CA 90292 USA
btimmer@isi.edu

Abstract

As mission critical communication increases on open internetworks, there is a growing need for end-to-end protection from traffic analysis. This additional protection can be expensive and detrimental to performance when padding is used to mask traffic patterns. Traffic masking policies that are responsive to system service requirements can improve performance and lower cost. However, adaptive traffic masking has to balance performance requirements with system protection requirements and thus address the information leaks that result from adaptations. This paper defines a new security model for adaptive traffic masking that satisfies system requirements for protection, efficiency, and performance. It presents secure dynamic adaptive traffic masking (S-DATM) techniques, defines a security model for dynamic adaptive traffic masking (SMD), and presents an example of the model applied to a network protocol. Mechanisms that utilize S-DATM techniques can integrate into existing security protocols to provide end-to-end protection that is scalable to internetworks. S-DATM detects and limits information leaks caused by dynamic adaptation. SMD is based on a probabilistic state machine formulation. It allows secure trade-offs between protection and performance as well as specifying the security requirements for S-DATM mechanisms.

1 Introduction

Traffic flow confidentiality (TFC) can meet the growing need for protection from traffic analysis on open internetworks, but can be expensive because traffic masking involves the use of padding. Adaptive techniques can reduce the costs of traffic masking and respond to application level service requirements. In this paper we present a security model for dynamic adaptive traffic masking, that specifies its protection requirements. The model specifies acceptable ranges of behavior that satisfy both protection and performance requirements and specifies secure trade-offs among performance, efficiency, and protection. In addition we present techniques for dynamic adaptive traffic masking that address the issues of application performance requirements versus system protection needs and the reduced protection that can result from dynamic adaptation. The proposed techniques include protection from statistical anomalies that are caused by adjustments.

1.1 Background

TFC provides protection from traffic analysis by masking frequency, length, and origin-destination traffic patterns of communications between network protocol entities [ISO84, ISO88]. Traditionally, TFC has been provided by bulk link encryption between protected sites on private networks [Bar64, Ram90]. However, dedicated private networks are too expensive for wide scale internetwork use. Public networks provide universal connectivity and are not only cheaper, but more available and reliable. Users of public networks now have capabilities for end-to-end operational security, with such features as confidentiality, integrity, and varying degrees of privacy. In addition, some users, such as government agencies, industrial corporations, and financial institutions need protection from traffic analysis. Their Internet traffic becomes vulnerable to hostile observation when they are unable to control the links that their traffic traverses.

The fundamental mechanisms of TFC are encryption, traffic padding and delays, and routing control. TFC traffic masking mechanisms introduce noise on connections between network protocol entities in the form of padding traffic and added delays. The noise protects network traffic from analysis by masking frequency and length patterns. TFC mechanisms for end-point ambiguity mask source-destination traffic patterns and obscure protocol components (such as addresses). The location of TFC mechanisms in the network architecture determine what information is masked.

TFC is frequently considered too expensive in terms of bandwidth consumption. 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 [Bad90, Bro94, HW89, Web88, NWV91]. It is also acknowledged that such adjustments can cause leaks of information about protected systems.

*This research was supported by the National Security Agency under the University Research Program contract no. MDA904-94-C-6114.

†The author wishes to acknowledge the suggestions and assistance of B. C. Neuman, G. Tsudik, and B. Tung in the preparation of this paper, with special thanks to I. Moskowitz and the other participants in the NSP '97 Workshop.

1997 New Security Paradigms Workshop Langdale, Cumbria UK

1.2 Overview of Our Approach

Internet users that want their traffic protected from analysis need adaptive traffic masking that scales to the Internet and satisfies their security requirements with minimal cost while meeting system and application performance requirements. As part of our proposed solution, we define a security model that specifies security requirements for *Secure Dynamic Adaptive Traffic Masking* (S-DATM) mechanisms. We use statistical techniques that detect and prevent (or limit) leaks of inference information which may occur when dynamic adjustments are allowed. The fundamental mechanisms of TFC are encryption, traffic padding and delays, and routing control. Our approach does not address encryption. It is assumed that all S-DATM mechanisms are protected by encryption. We also do not address routing control in this paper.

Our criteria for S-DATM techniques are:

1. satisfy system security requirements,
2. meet the performance requirements of the original traffic,
3. minimize padding costs,
4. solutions that are scalable to internetworks.

In addition to specifying system protection requirements, our security model allows trade-offs between the cost of protection and the original traffic's performance requirements, both in the design process and in dynamic adaptations.

1.3 Scope of Paper

While traffic masking that hides frequency and length traffic characteristics is a significant feature of TFC, masking origin and destination patterns is also essential. The scope of this paper is necessarily limited to masking frequency and length traffic patterns using statistical techniques. However, to effectively provide protection from traffic analysis, S-DATM mechanisms must be part of system of TFC protection that includes anonymity for senders and receivers. We will be presenting models for source and destination ambiguity in future papers. S-DATM mechanisms are modular and can fit into a system that includes already proposed mechanisms for masking origin and destination patterns such as Onion Routing [GRS96] or the sorting algorithm proposed in [RS93]. Onion Routing provides anonymous socket connections by means of proxy servers.

The rest of the paper is organized as follows:

- In Section 2 we first discuss work related to S-DATM and outline the S-DATM approach. We describe statistical anomaly detection and how it can be avoided with statistical techniques. Section 2 ends with a summary of the advantages of S-DATM techniques and a discussion of trade-offs among protection, efficiency, and performance.
- In Section 3 we focus on the S-DATM security model. We describe an S-DATM module and then define a security model, *SMD*, that is based on a probabilistic state machine formulation. We end with an example application of *SMD* that defines security for an S-DATM module in a secure IP protocol.
- In Section 4 we outline ongoing and future work that includes discussions of implementations of an S-DATM module that interface with Simple Mail Transfer Protocol and secure IP.

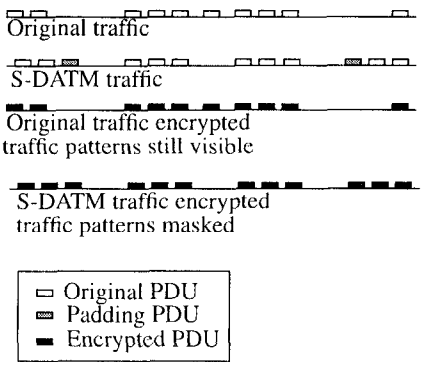

Figure 1: S-DATM masks frequency and length of traffic patterns when protected by encryption

- Our approach is summarized in Section 5.

2 Secure Dynamic Adaptive Traffic Masking (S-DATM)

S-DATM systematically introduces noise on connections between two network protocol entities in the form of padding and delays. It masks the frequency and duration of communications as well as traffic characteristics by creating fixed or random patterns of traffic transmittal, as illustrated in Figure 1. When used in conjunction with mechanisms that obscure communication end points and protect confidentiality, S-DATM provides end-to-end TFC. Dynamic adaptation can reduce the costs of traffic masking and respond to application level quality-of-service requirements. At the same time, S-DATM addresses system protection needs and the reduction in protection that results from dynamic adaptation. It also considers the costs of protection and application performance requirements. In the next section we discuss related work in this area.

2.1 Related Work

TFC has been considered too expensive and detrimental to performance to be considered practical on public networks [VK83]. The principal cost of TFC arises from the use of padding to mask traffic patterns. Recently proposed technology for protection from traffic analysis outlines adaptive policies that can reduce costs of traffic masking and respond to the changing rates of original traffic [Ven94]. Implemented in the transport layer, the rate of masked traffic increases and decreases as the rate of original traffic increases and decreases, thereby reducing the amount of padding. These policies do not address system protection needs or the reduced degree of protection that results from responsive traffic masking. Constraints on the model, such as fully connected nodes with global synchronization, are not readily scalable to internetworks.

The Network Layer Security Protocol (NLSP) [ISO92], provides a traffic padding mechanism for end-to-end TFC. The decision to use the mechanism is made by the local NLSP machine and is dependent on local security policy [Mac93]. No technology is provided for determining how to use the mechanism.

The Pump, developed at the Naval Research Laboratory [KM93], to provide good performance and security in Multiple Level Security (MLS) systems has an approach similar to

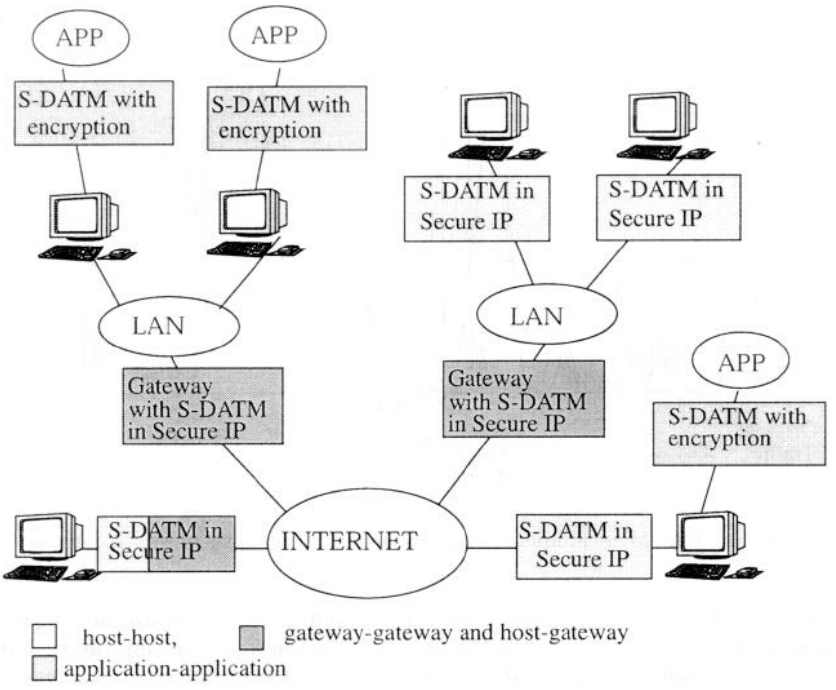

Figure 2: S-DATM modules can provide end-to-end protection between hosts in the application layer and between hosts and gateways in the IP layer

S-DATM in that it introduces noise on connections between High and Low in the form of delays. It protects against covert timing channels in the ACK stream from High to Low that can occur in the conventional store and forward buffer type of communication. The Pump has been extended to the network Pump [KML95, KML96] that protects MLS networks from denial of service attacks in addition to covert channels. The Pump is designed to thwart covert communication from High to Low in an MLS system.

2.2 The S-DATM Approach

S-DATM techniques provide an end-to-end modular system that fits into existing environments and protocols. It is constructed out of modular components (with well defined interfaces) that can be placed in existing protection systems. This facilitates adding TFC protection to established systems. In Figure 2, S-DATM modules are shown immediately below the application layer and in the IP layer to provide end-to-end protection of traffic between hosts and gateways. The communicating hosts can be in the same LAN or across the Internet.

S-DATM techniques have the following objectives:

- To clarify the relationships between protection, efficiency, and performance in traffic masking.
- To provide the facility for trade-offs among the three.
- To provide a modular system that fits into existing security protocols.

In S-DATM protection is the masking of the characteristics of original traffic. It can be measured as a numerical value when statistical methods are applied to masked traffic. Adaptation in masking schemes can cause *statistical anomalies* in masked traffic patterns. Statistical anomalies occur when the statistics of observed traffic deviate significantly from expected statistics. Nonadaptive schemes are independent of original traffic characteristics and, therefore, not vulnerable to statistical anomalies. However, they also do not make adjustments for improved performance or efficiency.

Adaptive schemes create changing patterns of traffic that can be analyzed. Such schemes can be attacked by *statistical anomaly detection* on the statistics of changing traffic characteristics. Abnormal traffic behavior can be observed by statistical anomaly detection without having to compare the current behavior with all historical data, but by comparing the statistics of a period of recent short-term history with the statistics of long-term behaviors using frequency tables, means, and covariances [JV91].

Statistical anomaly detection can be used to attack adaptive traffic masking schemes that rely on covert channel capacity estimation methods. In such schemes channel capacity alone is not a measure of the vulnerability of traffic characteristics [MK94]. S-DATM includes techniques that prevent statistical anomaly detection in the output of adaptive masking schemes.

The less padding a scheme produces in relation to original traffic, the more efficient it is. Efficiency is achieved by a trade-off with protection and/or performance. In traffic masking schemes modifications are possible to improve efficiency, such as basing the probability of adjustments on environmental factors like the time of day.

The performance of traffic masking schemes is evaluated by how well the masked traffic meets the throughput and delay requirements of the protected applications. Efficiency is measured by the percent of masked traffic that consists of original traffic. The less padding in the masked traffic, the

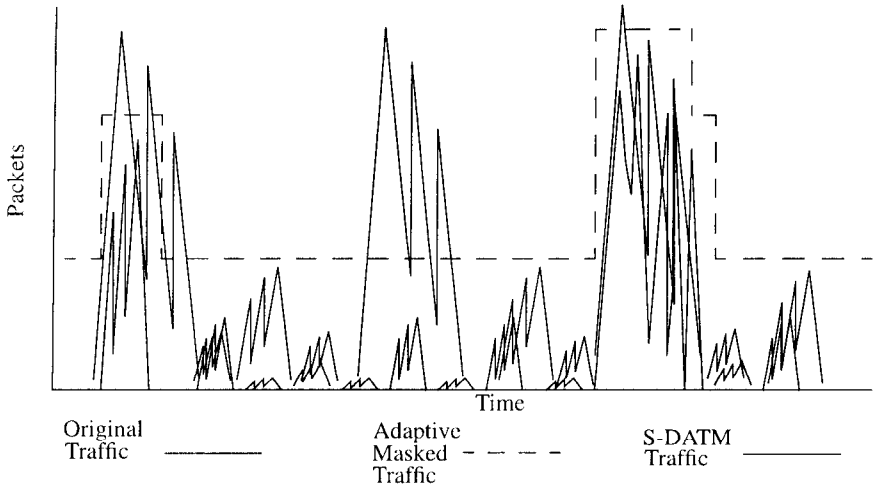

Figure 3: S-DATM techniques mask frequency and length of original traffic patterns

more efficient the scheme. SMD allows trade-offs between performance and efficiency both in the design process and in dynamic adjustments. In addition, performance can be improved by basing the probability of dynamic adjustments on traffic characteristics or network conditions. Additional fine tuning of parameters in the statistics can also result in improved performance and efficiency.

Section 2.3 goes into more detail on statistical anomaly detection and describes how it can attack adaptive traffic masking schemes. Following this, Section 2.4 discusses S-DATM's statistical techniques and how they protect adaptive schemes. Finally, Section 2.5 summarizes the advantages of applying S-DATM techniques.

2.3 Statistical Anomaly Detection

Adaptive masking schemes create patterns of traffic that, while different from original traffic patterns, may still imply characteristics of the original traffic. Even slight changes in traffic patterns may be discerned by statistical anomaly detection.

Anomaly detection systems contain three distinct phases [Fra94]:

1. Collection and abstraction of information on system behavior.
2. Evolution of background information on past system behavior using the abstracted information.
3. Establishment of background boundaries that determine anomalous behavior.

In statistical anomaly detection mechanisms statistics are usually kept as frequency tables, means, and covariances. If the statistics of observed data are sufficiently far from the expected values that are based on the statistics of the background abstraction, the data is considered anomalous. For example in [JV91], a *summary test statistic* on collected data can reflect both 1) how much the recent behavior differs from the background abstract, and 2) the degree of abnormality of the recent behavior (based on the historical probability distribution of a recently collected statistic). Whether the amount of difference or the degree of abnormality has greater influence on the size of the summary test statistic depends on the nature of information that the statistic reflects. In either case, the larger the summary test statistic, the greater the degree of abnormality of the recent behavior.

2.4 S-DATM Statistical Techniques

S-DATM anticipates and exploits the phases of anomaly detection mechanisms that are described in the previous section. The pertinent characteristics of masked traffic are: burst size, PDU inter-arrival times, and throughput because these are the traffic characteristics that can reveal the *type* of traffic. In S-DATM schemes, masked traffic is manipulated so these characteristics conform to a *profile*, a background abstraction made up of statistics of traffic behavior. In the masked traffic's expected behavior, as defined by the profile, adaptations are normal. In addition, a system's security policy may allow changes in the boundaries that define tolerance levels for variations in behavior, depending on events in the environment of the system. The initial tolerance levels reflect the relative importance of protection, efficiency and performance. They can subsequently be tuned to reflect changing priorities of these factors. Figure 3 shows how output from an S-DATM scheme masks bursts of original traffic by having random bursts of masked traffic that are not dependent on the original traffic in addition to those bursts that are caused by adaptations to original traffic.

Traffic masking schemes may have to perform real time calculations based on events in the higher protocol layers. One advantage of statistical analysis is that it is not necessary for schemes to maintain logs of past behavior as this information is accumulated in the statistics. Summary statistics accumulate data in an additive fashion. Recent behavior is stored as an exponentially weighted sum of past changes. The exponent parameter of the sum can be adjusted to vary the weighting on the most recent history.

If b_n is the behavior value after the nth interval then the equation for updating the behavior is:

$$b_{n+1} = b_n \cdot 2^{-rt} + D_n$$

The initial behavior value, b_0, is assigned. D_n is the change in measured values of the characteristic from the nth to the

$(n+1)$th interval, t is the length of the sampling interval for measuring behavior values and r is the *decay rate* that determines the *half-life* of the effect of past behavior, i.e., the number of recent past states who's behavior affects the sum. If r is large, the value of b_{n+1} is mostly determined by the most recent past. If r is small, the value of b_{n+1} is more influenced by the distant past.

The half-life should be small enough so that it responds to rapid changes in behavior, but, at the same time, large enough to correctly reflect the relative normality of recent behavior. The masking scheme can create a frequency distribution for traffic behavior, such as burst size, by manipulating bogus burst sizes within appropriately selected intervals. Consequently an adjustment to adapt to a burst in the original traffic will not appear as an anomaly.

2.5 Advantages of S-DATM Mechanisms

The features of S-DATM statistical techniques are:

- precise specification of system protection requirements that satisfy a security model for adaptive traffic masking.
- adjustable protection against statistical analysis of the characteristics of masked traffic
- precise policies for adaptations by using the statistics of traffic characteristics
- secure dynamic trade-offs between protection and application performance requirements can be made if they meet the security policy specifications
- reduction in processing and storage overhead of adaptive schemes through the use of statistical techniques
- real-time evaluation of current local data
- end-to-end scalable methods

S-DATM mechanisms meet systems protection requirements by defining recognizable traffic characteristics and masking them in a systematic manner that results in improved network performance. By specifying trade-offs, it is possible to obtain a more efficient use of network resources and reduce the impact of padding on system and application performance.

2.6 Trade-offs

Schemes that are nonadaptive provide the highest degree of protection with no dynamic trade-offs for improved performance or lower bandwidth consumption. Any adjustments to masked traffic are independent of original traffic characteristics. There may be a trade-off between efficiency and performance or between efficiency and protection in the implementation phase of any scheme, adaptive or nonadaptive.

An example of trading a degree of protection for increased efficiency occurs when low throughput masked traffic is protecting a low throughput application efficiently. An inference can be made about the low throughput of the protected application. This trade-off might suffice for users that are more concerned with masking the frequency and length of their connections, than the type of application. Low throughput masked traffic protecting a high throughput application is an example of trading performance for greater efficiency. High throughput masked traffic protecting a low throughput application inefficiently, i.e. transmitting more

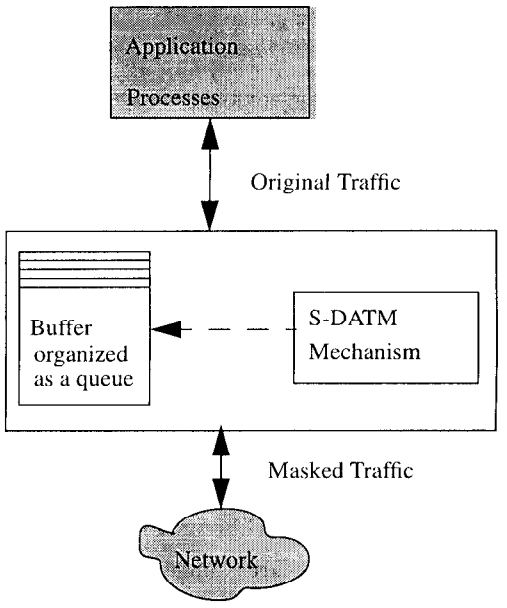

Figure 4: A transparent S-DATM module in the protocol stack between the application and network.

padding than is necessary, is an example of trading efficiency for improved performance or protection.

Adaptive schemes may lower the degree of protection by making adjustments that are based on original traffic characteristics, but these adjustments can improve performance and/or lower bandwidth consumption. For the sake of increased performance or efficiency, an adaptive scheme may leak controlled amounts of information about original traffic if allowed by the security policy of the system. A protection trade-off for improved performance is most desirable for systems with stringent performance requirements whose protection needs have room for a trade-off.

3 S-DATM Security Model

We present a S-DATM security model, *SMD*, formulated as an abstract machine description of a traffic masking module. A security model must have the properties of precision, simplicity, and generality, while still being intuitively representative of system security policies. In general, security models are design techniques that specify acceptable ranges of system behavior that satisfy a given security policy. The S-DATM security model also specifies ranges of behavior that satisfy performance requirements, and, additionally, may allow trade-offs among protection, efficiency, and performance.

3.1 Description of an S-DATM Module

SMD specifies the security requirements for an S-DATM module, an input/output system where the input is original traffic and the output is masked traffic, illustrated in Figure 4. The S-DATM mechanism outputs masked traffic in the form of real PDUs from the queue, bogus PDUs (padding) and added delays. The profile of the masked traffic consists of a collection of statistics of various traffic characteristics, such as throughput and delay and the tolerance levels for these statistics. The statistics of characteristics of the outgoing traffic are constrained by the tolerance levels to stay

sufficiently close to those of the profile. Time is discrete; at each unit of time the masking scheme accepts one (or no) input PDU and emits one (or no) output PDU. The module can adapt to changes in the environment (such as an event that increases the rate of the input) by adjusting the critical values of the profile that determine the tolerance levels. Allowable adjustments and interval size depend on the desired degree of protection. Larger interval sizes and fewer, or smaller, adjustments provide greater security. In Figure 4, original traffic is queued and subsequently output with introduced delays and padding that depend on the profile. Queue length can also be a consideration in determining the rate of the output.

The goal of a secure S-DATM module is met when the system satisfies the criteria of its security model. The proposed model is based on traffic characteristics which not only facilitates providing protection from statistical analysis, but is also a means for increased efficiency and reliability.

3.2 Definition of the S-DATM Security Model (*SMD*)

A state machine-based security model includes all security-relevant state variables and functions. It specifies the variables, transition functions, constraints, and a secure initial state. If it can be shown that a system begins in a secure state, and when transition functions are applied with all constraints satisfied, it remains in a secure state, then the system is considered secure. A secure state for a state machine can be defined, but can not be proven secure; its security must be intuitively obvious [Gas88].

An S-DATM Security Model is a probabilistic state machine (time is discrete) defined by the tuple

$$SMD = \langle P, E, S, s_0, \Phi \rangle$$

where:

P, the profile of the outgoing traffic from the module, is a finite tuple $(p_1, p_2, ..., p_n)$ where n is the number of traffic characteristics being considered in the scheme and each p_k is a triple $(N_{pk}, C_{pk}, \Delta_{pk})$ such that

- N_{pk} is a positive integer that represents the size of the recent history interval over which the kth traffic characteristic is measured. (Each recent history interval includes the N_{pk} most recent states.),
- C_{pk} is the desired value of the kth characteristic of the outgoing traffic during the recent history interval,
- $\Delta_{C_{pk}}$ is the allowable variance from the profile value of the traffic characteristic C_{pk}.

E is a finite set of events in the environment of the S-DATM module that may affect the profile, i.e., the value of N_{pk}, C_{pk}, or Δ_{pk} for any p_k in set P.

S is a finite, ordered, non-empty sequence of tuples,

$$s_j = (i_j, o_j, q_j, B_j)$$

where $s_j \in S$ is the state of the machine at time j, and

i_j is the input, where

$$i_j \in \{0, 1\}$$

and 0 stands for no PDU and 1 stands for one PDU,

o_j is the pair (r_j, d_j) where

$$(r_j, d_j) \in \{(0,0), (0,1), (1,0)\}$$

and r_j is the real (original) traffic output and d_j is the bogus traffic output (padding),

q_j is the length of the queue in the module's buffer,

B_j is the tuple $(C_{j1}, C_{j2}, ..., C_{jn})$ representing the recent history of the traffic characteristic such that

- C_{jk} is the value of the kth characteristic of the outgoing traffic as measured during the recent history interval of size N_{pk},
- B_j represents the behavior of the outgoing traffic during the recent history interval, i.e., the most recent N_{pk} states (depending on traffic characteristic being measured) at state s_j.

s_0 is the initial state of the machine $s_0 = (i_0, o_0, q_0, B_0)$ such that

- i_0 and o_0 are assigned and not dependent on the environment,
- $q_0 = 0$,
- B_0 is $(C_{p1}, C_{p2}, \ldots, C_{pn})$.

Φ is the transition probability function for SMD such that

- $\Phi: P \times E \times S \times S \rightarrow [0, 1]$ where
  - Φ is defined by two component functions, Φ' and λ such that
    - $\Phi': P \times S \times S \rightarrow [0, 1]$, a probability function,
    - $\lambda: P \times E \times S \rightarrow P$, and
  - $\Phi(P, e_j, s_j, s_{j+1}) = \Phi'(P', s_j, s_{j+1})$, where
    - $P' = \lambda(P, e_j, s_j)$ and
    - $e_j \subseteq E$ is a set of events occurring in the jth interval that may affect the profile P.

The function Φ considers changes in the environment before a transition is made to the next state.

3.2.1 Secure State Invariant

SMD is secure if and only if, for each state in S, the probability of the output considering the recent history of output *and* input is sufficiently close to the probability of the output considering only the recent history of output. That is, SMD is secure iff $\forall s_j, s_{j+1} \in S$:

$$|p(o_{j+1} | i_j, B_j) - p(o_{j+1} | B_j)| \leq \epsilon$$

where p is a probability function and the value of ϵ is determined by the system's security requirements. The invariant guarantees the output is as close as necessary to conditional statistical independence from the input to meet system protection requirements.

In McLean's Flow Model (FM) [McL90] a system satisfies the security model FM only if $p(L_t | H_s, L_s) = p(L_t | L_s)$ where L_t is the value of low-level objects in state t and H_s and L_s are the sequences of values assumed by high-level and low-level objects respectively in every state preceding t. The SMD invariant is similar to FM. It agrees with FM when input is considered high-level objects, output is low-level, the value of ϵ is 0, and the recent history interval includes all previous states.

3.2.2 Constraint on the Transition Function

In SMD, a constraint on the transition function is needed because, in traffic masking, security is a property of sequences of states as well as a property of the current state. If $\forall s_j$ and $s_{j+1} \in S$, the following relationship holds as s_j makes the transition to s_{j+1}:

$$|p(o_{j+1}|o_j, B_j) - p(o_j|B_j)| \leq \psi$$

where p is a probability function and the value of ψ is determined by the system's security requirements, then the constraint guarantees that, when a state transition occurs, the probability distribution of the next state's output (considering the recent behavior of the outputs) comes as close as necessary to the probability distribution of the current state's output (considering the recent behavior of outputs) to meet the protection requirements of the system. Thus, only allowable changes to the output (the masked traffic) are caused by events in the environment, and a breach in security does not occur.

3.3 Example of an S-DATM Security Model

We will consider an example of an S-DATM module that interfaces with a secure IP protocol and, applying the definition outlined in the previous section, illustrate the SMD that specifies the module's security requirements. The S-DATM module is an input/output system where the input is original traffic in the form of packets from the transport layer and the output is masked traffic in the form of real packets from the S-DATM queue usually combined with padding (bogus) packets and added delays. The secure IP protocol supplies the necessary encryption.

The profile of masked traffic consists of a set of statistics from the traffic characteristics, throughput, inter-arrival delay, and burst size. During a predetermined interval of time (different for each characteristic) the statistics must stay sufficiently close to those of the profile. At each discrete minimum unit of time it accepts one (or no) input and makes one (or no) output. For the sake of simplicity, this module does not change the profile in order to adapt to changes in the environment (such as events that increase the rate of the input). The allowable adjustments and interval sizes are specified in the profile and depend on the desired degree of protection as well as the values of ϵ and ψ in the secure state invariant and the transition constraint. The input original packets are queued and subsequently output with introduced delays and padding that depend on the profile. A maximum allowable queue length could also be a consideration in determining the output. This is not done in this example, again for the sake of simplicity.

The security model for the example module is defined by $SMD = \langle P, E, S, s_0, \Phi \rangle$ where:

$$P = ((N_{p\delta}, \delta_p, \Delta_{p\delta}), (N_{p\Theta}, \Theta_p, \Delta_{p\Theta}), (N_{p\beta}, \beta_p, \Delta_{p\beta}))$$

is the profile of the outgoing traffic module where δ, Θ, and β represent average inter-packet delay, throughput, and burst-size, respectively.

E is $\{e_\emptyset\}$ a null event because, in this example, events in the environment will not cause a change to the values of the profile.

S is the sequence of states of the machine

$$s_j = (i_j, o_j, q_j, B_j)$$

at time j, where i_j, o_j, and q_j are defined as above, and B_j is $(\delta_j, \Theta_j, \beta_j)$ where δ_j, Θ_j, and β_j are respectively the inter-packet delay, throughput, and burst size, measured during the interval of recent history at time j.

s_0 is the initial state and $s_0 = (0, (0,0), 0, B_0)$, where $B_0 = (\delta_p, \Theta_p, \beta_p)$.

Φ is the transition probability function defined as follows:

$\Phi(P, e_j, s_j, s_{j+1}) = \Phi'(\lambda(P, e_j, s_j), s_j, s_{j+1})$ where

$\lambda(P, e_j, s_j) = P$ (no events affect the profile P),

$s_j = (i_j, o_j, q_j, B_j)$, and

$s_{j+1} = (i_{j+1}, U(q_j, B_j), Q(q_j, i_{j+1}, o_{j+1}), V(B_j, o_{j+1}))$,

$\Phi'(\lambda(P, e_j, s_j), s_j, s_{j+1}) = \pi_{j+1}(i_{j+1}) \cdot F(r_{j+1} + d_{j+1})$,

where $\pi_j(1) = 1 - \pi_j(0)$ is the probability of a packet arriving at time j.

U, Q, V, and F are defined as follows:

Denote $B = (\delta, \Theta, \beta)$ and

condition $A = (((\delta - \delta_p) > \Delta_{p\delta}$ AND $(\Theta - \Theta_p) < \Delta_{p\Theta})$ OR $(\beta - \beta_p) < \Delta_{p\beta})$,[1]

then

$$U(q, B) = \begin{cases} (0,1) & \text{if A is true and } q = 0 \\ (1,0) & \text{if A is true and } q \neq 0 \\ (0,0) & \text{if A is false} \end{cases}$$

and $U(q_j, B_j) = o_{j+1}$.

If $Q(q, i, (r, d)) = q + i - r$

then

$$Q(q_j, i_{j+1}, o_{j+1}) = q_{j+1},$$

$V((\delta, \Theta, \beta), (r, d)) =$

$((r + d) + 2^{-a}\Theta, (1 - (r + d))/\Theta + 2^{-b}\delta, (r + d) + 2^{-c}\beta)$

where the values of a, b, and c determine the half lives of Θ, δ, and β, respectively, and $V(B_j, o_{j+1}) = B_{j+1}$.

$$F(x) = \begin{cases} \beta_p / N_{p\beta} & \text{if } x = 1 \\ 1 - \beta_p / N_{p\beta} & \text{if } x = 0 \end{cases}$$

where F is a probability function.

4 Ongoing and Future Work

Our plans for the future are to implement an S-DATM module with clearly defined interfaces that will integrate into existing security architecture as part of a solution for traffic flow confidentiality in open internetworks. We will verify the correctness of the implementation by illustrating that it satisfies SMD.

We are currently implementing a prototype of an S-DATM module in Simple Mail Transfer Protocol (SMTP) for development and evaluation. SMTP provides discrete input in the form of mail messages that can be required to be the same size without loss of generality. Delay and throughput

[1] Condition A depends on the characteristics of the protected traffic. The presence of a mandatory burst of masked traffic for each interval of length $N_{p\beta}$ allows the masked traffic to adapt to bursts of original traffic without causing detectable statistical anomalies.

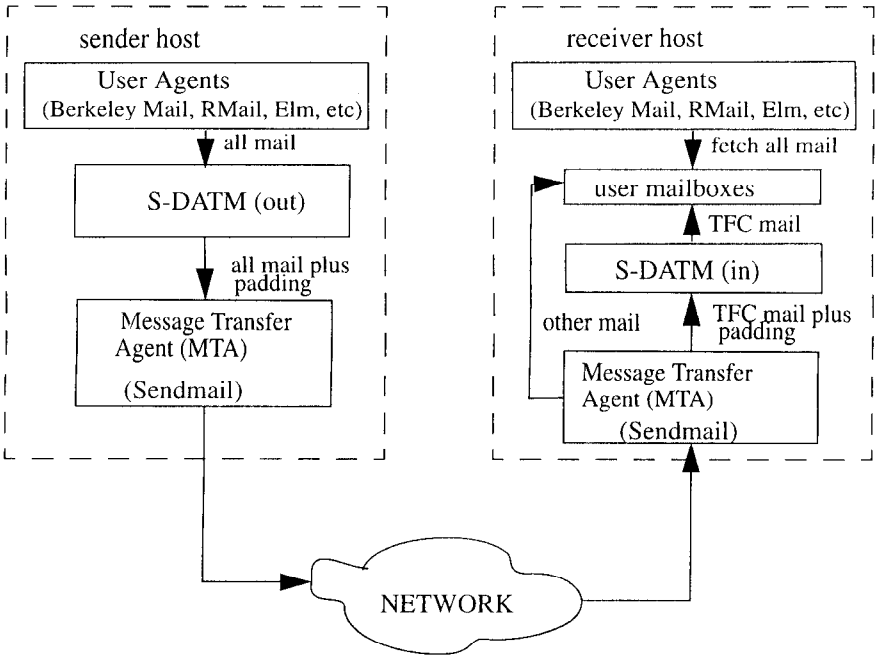

Figure 5: S-DATM implemented with Simple Mail Transfer Protocol

can be accurately measured. While SMTP has relatively lenient real time constraints, performance of masked traffic can be compared with the performance of the same traffic unmasked in order to evaluate the module.

An S-DATM module placed within SMTP is described in Figure 5. In this implementation the S-DATM Module is placed between the User Agents, such as Berkeley Mail or Rmail, and the Message Transfer Agent (MTA), such as Sendmail, with the module transparent to both. As shown in Figure 5, the module receives all the mail from the User Agents. Depending on the masking scheme in use, it then forwards the unprotected mail to the MTA and places the protected mail on a queue. It recognizes protected mail by the address of the receiving host. The masking scheme decides when protected mail is removed from the queue and forwarded to the MTA and also determines when padding is sent to the MTA in the form of bogus messages to the receiving host. The MTA is unaware that these messages are bogus. Encryption of all protected mail is assumed. The messages are addressed to a special userid in the receiving host where, upon arrival, the bogus messages are dropped and the real mail is forwarded to the receiving users' mailboxes.

The performance of masking schemes is evaluated by how well the performance of masked mail matches the performance of the same mail, unprotected. The throughput and introduced delay are compared. The S-DATM security model allows trade-offs for performance both in the design process and in dynamic adjustments. Performance is improved by basing the probability of dynamic adjustments on traffic characteristics or network conditions. Fine tuning of parameters in schemes results in improved performance and efficiency, however, usually a trade-off results. The less padding a scheme produces in relation to original traffic, the more efficient it is, but efficiency is 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. This can reduce the degree of protection, by leaking information about the throughput.

The main objectives of evaluations of S-DATM techniques used in this prototype is to clarify the relationships between protection, efficiency, and performance in traffic masking, and to analyze trade-offs among the three. Additional research will apply the same techniques and evaluation criteria to masking the destination of the mail.

As future work, we plan to implement an S-DATM module below TCP, as illustrated in Figure 6, applying the results of the SMTP implementation's evaluations. The module will interface with the secure IP protocol. For outgoing traffic, secure IP mechanisms are usually placed immediately above the IP layer, where traffic targeted for protection is routed to the secure IP for encapsulation using an IP-inside-IP protocol.

It is desirable to place a traffic masking module below the transport layer for several reasons. It is more efficient in terms of additional consumed bandwidth to have one module per host or gateway. It can interface with the secure IP protocol and does not have to provide different interfaces for each application and each transport layer protocol. However, there are issues that must be addressed when an S-DATM module is placed in the IP layer. The module must cooperate with underlying transport protocols to satisfy performance requirements of applications An example is TCP *slow start* that is used for both flow control and congestion avoidance. On the other hand, the response to TCP should not leak information about the original traffic. To provide adequate protection, a masking mechanism must also mask the commonly ignored back channels generated by TCP protocols for both interactive and bulk data flows.

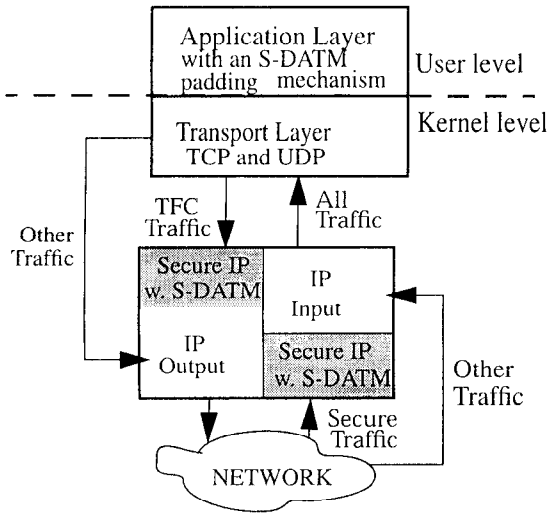

Figure 6: S-DATM implemented with secure IP

The implications are that, in implementations, both ends of a connection must be protected with a masking mechanism.

5 Summary

There is a need for TFC technology that balances the costs of protection with user performance requirements, and considers system protection needs and the reduction in the degree of protection that results from dynamic adaptation. This paper introduces new secure dynamic adaptive traffic masking (S-DATM) techniques for secure traffic masking mechanisms that adapt to changing traffic conditions. In addition, it defines a security model (*SMD*) that specifies the protection requirements for S-DATM mechanisms within an internetwork environment.

S-DATM is a tunable solution to meet specific protection requirements while minimizing overhead and meeting application performance requirements. S-DATM techniques have the capability to dynamically adapt to changing environmental conditions in order to meet the throughput and delay requirements of the original traffic. They use statistical methods to detect and prevent statistical anomalies that contain information about original traffic characteristics.

References

[Bad90] Lee Badger. Providing a flexible security override for trusted systems. In *Proceedings of Computer Security Foundations, Workshop III*, Franconia, NH, June 1990.

[Bar64] Paul Baran. On distributed communications, vol ix. In *Security Secrecy and Tamper Free Considerations, Memo RM-3765-PR.* Rand Corp., Santa Monica, CA, August 1964.

[Bro94] Randy 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.

[Fra94] Jeremy Frank. Machine learning and intrusion detection: Current and future directions. In *Proceedings of the 17th National Computer Security Conference*, October 1994.

[Gas88] Morrie Gasser. *Building a Secure Computer System.* Van Nostrand Reinhold Company, New York, NY, 1988.

[GRS96] D.M. Goldschlag, M.G. Reed, and P.F. Syverson. Hiding routing information. In *Workshop on Information Hiding*, Cambridge, UK, May 1996.

[HW89] 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.

[ISO84] ISO. Information Process Systems - Open System Interconnection - Basic Reference Model - ISO 7498. American National Standards Association, Inc., New York, NY, 1984.

[ISO88] ISO. Information Process Systems - Open System Interconnection Proposed Draft Addendum 2 ISO 7498. American National Standards Association, Inc., New York, NY, 1988.

[ISO92] ISO. ISO/IEC JTC1/SC6, Network Layer Security Protocol. ISO-IEC DIS 11577 International Standards Organisation, November 29 1992.

[JV91] Harold R. Javitz and Alfonso Valdes. The SRI IDES statistical anomaly detector. In *Proceedings of the IEEE Symposium on Research in Security and Privacy*, Oakland, CA, April 1991.

[KM93] Myong H. Kang and Ira S. Moskowitz. A pump for rapid, reliable, secure communication. In *Proceedings of the ACM Conference on Computer and Communication Security*, Fairfax, VA, 1993.

[KML95] Myong H. Kang, Ira S. Moskowitz, and Daniel C. Lee. A network version of the pump. In *Proceedings of the IEEE Symposium on Security and Privacy*, Oakland, CA, May 1995.

[KML96] Myong H. Kang, Ira S. Moskowitz, and Daniel C. Lee. A network pump. *IEEE Transaction on Software Engineering*, 22(5), May 1996.

[Mac93] Betty Mackman. Reorganized NLSP. DRA/CIS/(SE2)/BM/93/3/1, Defence Research Agency, March 1993.

[McL90] John McLean. Security models and information flow. In *Proceedings of the IEEE Symposium on Research in Security and Privacy*, Oakland, CA, May 1990.

[MK94] Ira S. Moskowitz and Myong H. Kang. Covert channels – here to stay? In *Proceedings of COMPASS '94*, Gaithersburg, MD, June 27 - July 1 1994. IEEE Press.

[NWV91] R. E. Newman-Wolfe and B. R. Venkatraman. High level prevention of traffic analysis. In *Seventh Annual Computer Security Applications Conference*, San Antonio, Texas, December 1991.

[Ram90] Raju Ramaswamy. Traffic flow confidentiality security service in OSI computer network architecture. In *Proceedings of the IEEE Region 10 Conference on Computer and Communication Systems*, Hong Kong, September 1990.

[RS93] Charles Rackoff and Daniel R. Simon. Cryptographic defense against traffic analysis. In *Proceedings of the 25th Annual Symposium on the Theory of Computing*, CA, USA, May 1993.

[Ven94] B. Venkatraman. *Prevention of Traffic Analysis and Associated Covert Channels.* PhD thesis, University of Florida, Gainsville, Florida, 1994.

[VK83] V. L. Voydock and S. T. Kent. Security mechanisms in high-level network protocols. *ACM Computing Surveys*, 15(3), June 1983.

[Web88] Douglas Weber. Security policies for army tactical C^2 systems. Technical report, Odyssey Research Associates, Inc., Ithaca, NY, 1988.