

Transmission Schedules To Prevent Traffic Analysis

Balaji R. Venkatraman¹]This work is partially supported by a grant from Bell South

Computer & Information Sciences.
University Of Florida,
Gainesville, Florida - 32611.

Abstract

In this paper, we propose scheduling strategies to determine transmission schedules that prevent traffic analysis and the creation of covert channels due to temporal variation in the transmission of packets. In addition to requiring the traffic matrix be neutral as defined in [2], we require the transmission schedule be temporally neutral to eliminate that potential covert channel. The static scheduling policy generates temporally neutral transmission schedules. We extend this to develop an adaptive scheduling policy that can adapt to long term load fluctuations. The adaptive algorithm is expensive and there exists the possibility of a low bandwidth and noisy covert channel; we suggest mechanisms to reduce the bandwidth of the covert channel. The tradeoff is between the “adaptability” of the scheduling policy and the bandwidth of the covert channel.

1 Introduction

In [2] we proposed a model for the prevention of traffic analysis and did a performance analysis in [3]. In the model, the given traffic matrix is converted to a neutral traffic matrix by rerouting and padding the original traffic between given source-destination pairs. A neutral traffic matrix is one in which all the non-diagonal elements in the traffic matrix have the same numerical value. This implies that the intruder cannot derive any useful information by observing the traffic on the network as the volume and nature (packet size and type) of traffic between any source-destination pair is identical to that between any other pair. However, we are concerned with the temporal variation in traffic and the possible introduction of covert channels due to the variation in transmission

characteristics over time. Note that though the total volume of traffic communicated between any pair of nodes in the network is mandated to be the same to satisfy the neutral traffic matrix requirement (*neutrality criterion*), the source could transmit the packets in a burst or could spread out its transmissions over a period of time. There is essentially no restriction placed on this aspect by the model and therefore a knowledgeable user might be able to communicate with his accomplice by timing the transmissions, thus introducing covert channels. In [3], we suggest that with each node pair there be a “packet transmit distribution” and that all transmissions for the given source follow this distribution. While this is an intuitively appealing and feasible solution, we wish to propose a more formal approach to the problem.

We propose a mechanism to generate a transmission schedule that will satisfy our primary goal of prevention of traffic analysis and the prevention of covert channels due to temporal variation in a packet transmission schedule. We adopt a slotted time system in which nodes transmit fixed size packets in fixed size slots. By generating *temporally neutral* (defined in Section 3) transmission schedules, we eliminate certain covert channels. The scheduling policy is *static* in the sense that the transmission schedule is fixed and the nodes just follow the transmission schedule at all times, immune to load variations in the system. Since this transmission schedule is known and remains the same over a period of time, there is no possibility of a covert channel. However, such a scheduling policy is not responsive to changes in the load and can degrade system utilization and performance significantly. We therefore propose an *adaptive* scheduling policy that can adapt to long term variations in the load, but which leaves open the possibility of a covert channel. However, the covert channel has very low bandwidth and is noisy; the “adaptiveness” of the mechanism can be traded off for the bandwidth (or even existence) of

^{*}[
⁰brv@cis.ufl.edu, nemo@chameleon.cis.ufl.edu

¹This work is partially supported by a grant from Bell South

the covert channel.

In section 2, we discuss the slotted system, followed by relevant definitions. In section 3, we formalize the notion of a temporally neutral transmission schedule. In section 4.1, we present the basic transmission schedule, in section 4.2 we present the static scheduling policy followed by its extension to the adaptive scheduling policy in section 4.3. In section 5, we suggest various ways to reduce the bandwidth of the covert channel and discuss some performance considerations, and conclude in section 6 with suggestions for future research.

2 Discussion of Slotted Time

We will preface our discussion on slotted time with the following comments. We assume that all transmitted packets have the same length and that each packet requires one time unit (called a slot) for transmission. The slotted system facilitates the development of the transmission schedule and also turns the system into a discrete-time system, thus simplifying analysis. Synchronizing the transmitters for slotted arrivals at the receiver is a non-trivial problem, but may be accomplished with stable clocks, some feedback from the receiver, and some guard time between the end of a packet transmission and the beginning of the next slot.

We have a $n \times n$ (neutral) traffic matrix, with nodes numbered 1 to n . Since we assume that this traffic matrix is available at the very outset, we can treat the arrival of packets as batched arrivals. Associated with each node is a virtual queue to buffer packets before they can be transmitted. During a given period i , new packets arrive and are placed in the buffer; packets arriving in the current period time may not be transmitted until the next period. Any packet in the buffer at the beginning of the period is eligible for transmission. We will now define some important terms. See figure 1.

- **Slot:** This is the basic time unit during which a given node may send or receive at most one packet.

If we also assume that at most one node can transmit per slot, then $n(n - 1)$ slots are needed. If all n nodes can transmit in a slot, then at most n slots are needed (round robin tournament). Without the above assumptions, some number of slots between n and $n(n - 1)$ are needed. All

that really matters is that the same nodes are sending to the same destination at the same slot each period (defined below). For purposes of this discussion, we will assume that $n(n - 1)$ slots are needed and used.

- **Period:** A period is a set of successive slots during which one phase of the transmission schedule is carried out. In our model, a period consists of $n(n - 1)$ active slots and m idle slots. In the static scheduling policy, m is a constant; in the adaptive scheduling policy, m may vary over time. During the active slots of each period, the nodes may exchange packets with each other. In a fully connected network, in one period each node exchanges one packet with every other node in the system. During the idle slots, the nodes remain idle.
- **Cycle:** A set of successive periods is a cycle. Typically the transmission of the entire volume of communication as prescribed by the traffic matrix is carried out in one cycle. In the adaptive scheduling policy, negotiation to change transmission parameters such as the number of idle slots is typically carried out at the beginning of a cycle. In an environment in which the load characteristics are constant over a long period of time, the length of the cycle is so large that the adaptive policy effectively behaves like the static policy.

We will now define the arrival of a packet, the buffering of packets (due to batched arrivals) and the transmission (departure) of packets during a slot. Note that here we are accounting for independent events; we are not concerned with the “transmission schedule” as yet.

- **The arrival of packets at node i .**
The arrival of packets at a given node can be accounted for as follows:

1. The packet belongs to actual traffic (as defined by the traffic matrix) with source i , destined for node j . In each period, a given node can receive at most one packet from every node in the network (due to the neutrality criterion).
2. The packet destined for node j is rerouted via node i by another node, say k . The packet has arrived at node i from node k

Figure 1: Slotted Time System

Figure 2: Covert Channel due to Transmission Frequency

in the period and will be queued for transmission to node j in the next period.

- **Departure of packets from node i .**

Depending on the transmission algorithm, the source node makes a decision about which packet (if any) to transmit. In each period, each node will transmit exactly one packet to every node in the network. A packet may not arrive and depart in the same period; therefore all packets that arrive during the previous period ($i - 1$) are eligible for transmission and may be transmitted if scheduled. A departing packet may be transmitted either to its actual destination or might be routed via an intermediate node; the scheduling policy treats these departures in the same manner.

- **Backlog of packets at node i .**

Packets remaining in the virtual queue at the end of each period are known as backlog packets. These packets are eligible for transmission in the next period. Let $A(i)$ represent the number of arrivals in period i , $D(i)$ the departure in period i and $B(i)$ the number of packets in the virtual queue (backlog). We then have the following relation which gives the number of backlog packets: $B(i) = B(i-1) - D(i-1) + A(i)$ for $i \geq 0$. Note that $A(i)$, not $A(i - 1)$, is used since the arrivals in period i are not eligible for transmission until period $(i + 1)$.

3 Prevention of Temporal Variation in Traffic

We have alluded to the creation of a covert channel due to temporal variation in the transmission of packets. Clearly, in a secure network, we would like to prevent the creation of any covert channels. We describe several related techniques to determine a transmission schedule so that the eavesdropper may not gain any useful information by observing the traffic on the network. We define certain criteria for transmission schedules; the enforcement of each of these criterion guarantee a temporally neutral transmission schedule.

- **The Volume (V) of communication between a given pair of nodes.**

In a neutral traffic matrix, the volume of communication between each pair of nodes is the same and the model, as defined in [2][3], does not require this information to be secret. By imposing

the *neutrality criterion* on the original traffic matrix, we have effectively eliminated the volume of communication between any pair of nodes as a contributing factor to the covert channel.

- **The Frequency (F) of communication between a given pair of nodes.**

If a user on node i communicates with another user on node j more frequently than he does with other nodes in the network, or exchanges packets with the node at a predetermined frequency, then a covert channel could exist. For example, the user at node i and j could encode some information in the frequency of communication between them. The overt communication itself may seem to be normal and legal; it is by timing the communication that they exchange information surreptitiously, thus creating a covert channel. In Figure 2, we see that information can be encoded by timing the transmission of packets. In this case, by computing the interarrival times, i.e., the time interval between the previous message and the current message, the intruder and his accomplice succeed in creating a covert channel. Even if the average frequency is constant due to equal volume restrictions, and even if each node sends exactly one packet to every other node per period, the position of the packet transmitted to a particular node within the period the period could contain information (Pulse Position Modulation). The bandwidth of this covert channel could be as much as $\log(n(n - 1)) \approx 2 \log n$ bits/period. By requiring that each node exchange a packet with every other node in the network each period (during the active slots), if necessary by padding with dummy packets, we ensure that the frequency of communication is the same between all pairs of nodes in the system, thus preempting the use of this characteristic of the transmission schedule to create a covert channel.

- **The Order (O) of communication between a set of nodes in the network.**

Given that the long term volume of communication and the frequency of communication between each pair of nodes in the network to be the same (due to the V and F criteria discussed above), the order of communication becomes relevant. For example, if a node sends a packet to node i before sending a packet to node j versus sending a packet to node j before sending a packet to node i , then information can be en-

Figure 3: Covert Channel due to Transmission Order

coded in the order in which nodes communicate with one another. For example, in Figure 3, the node sends a packet to node A followed by a packet to node B to encode “1” and the reverse order (BA) to encode “0”. If the intruder and his accomplice(s) can affect the transmission order in k nodes, then $k!$ transmission orders are possible. This implies that bandwidth of this covert channel could be as much as $\log(k!) \geq \frac{k}{2} \log k$ bits per period. In this case, the intruder and his accomplice have encoded information in the order in which they communicate between themselves and with other nodes in the network. By requiring that each node communicate with every other node in a predetermined order during all periods, the order of communication remains the same thus preempting the use of this characteristic of the transmission schedule to create a covert channel.

- **The (extrinsic) Nature (N) of communication between a given pair of nodes.**

Given that the volume, frequency and order of communication to be the same (due to the V, F and O criteria discussed above), the nature of communication becomes relevant. Assuming that the packets are encrypted (end-to-end encryption), the intruder cannot see the contents of the packets. However extrinsic characteristics like packet size can be used to exchange information covertly. For example, a user may send his accomplice a packet of some predetermined size followed by another packet of a different size to exchange some information covertly. We can eliminate these covert channels by requiring that the extrinsic characteristic of all packets be the same, by enforcing a fixed packet size and encrypting the packets.

- **The Length (L) (or duration) of transmission.**

If the volume, frequency, order and nature of communication among nodes is the same (due to the V, F, O and N criteria discussed above), then the question that arises is: for how long do these transmission characteristics remain unchanged and if they do change, when and how do they change? For example, if a single user is able to change any of the above parameters just by performing some local operations like increasing the load on the system or choosing to ignore any of the globally accepted parameters, then he can easily create a covert channel to communi-

cate with his accomplice. For instance, he may have an agreement with his accomplice to vary the load on the system at a given time and if the transmission parameters as seen on the network change immediately by significant amounts, then this could be used to signal his accomplice. (we assume that the eavesdropper is continuously monitoring the entire network and can detect any changes in the transmission characteristics). To eliminate this possibility, we should ensure that the globally selected parameters remain the same for an extended period of time and that a single user is not capable of changing the global parameters by himself. Any changes should be done by a negotiation process involving at least a majority of nodes, if not all nodes, and that the changes should be effected in a controlled manner. By building the communication system over NTCBs we can eliminate many problems, the most important being the violation of any aspect of the transmission protocol[1][4].

From the above discussion we can describe a transmission schedule by the five tuple $\langle V, F, O, N, L \rangle$. Depending on the information that can be encoded in any of these transmission characteristic, a covert channel may exist.

Definition: A *temporally neutral* transmission schedule is one in which none of the members of the tuple $\langle V, F, O, N, L \rangle$ have any characteristic that can be used to encode any information and communicate surreptitiously via a covert channel.

The transmission schedule that satisfies each of the restrictions in $\langle V, F, O, N, L \rangle$ is temporally neutral. In other words, the intruder cannot gain any useful information regarding the traffic matrix, source and/or destination user identity, etc., just by observing the flow of packets on the network. In effect, a user in collusion with an accomplice should not be able to use the volume, frequency, order, nature of communication, and the duration (length) of transmission in the network to exchange information surreptitiously.

4 The Basic Model

We consider the set of packets that each node has to transmit according to the traffic matrix, T , as batch arrivals. Some order of transmission of packets for all pairs (i, j) of nodes is selected. The specific order does not matter as long as all pairs appear exactly once in the order and the order and timing of transmission of

packets is fixed. Such a schedule satisfies some of the requirements of the tuple $\langle V, F, O, N, L \rangle$, in particular the V, F, O and N restrictions. We can then bring upon the transmission schedule various scheduling policies to smooth out temporal variations and by satisfying the L restriction, we achieve our objective of eliminating the covert channel due to temporal variation in traffic.

The restriction that each source-destination pair exchange exactly one packet per period, ensures both the neutrality criterion and the V and F criteria of $\langle V, F, O, N, L \rangle$ are satisfied. The N restriction can be satisfied as described earlier in section 3. However, we have not imposed any restriction on the order or length of the transmission period. Therefore we need to develop transmission schedules that will satisfy the O and L restrictions of $\langle V, F, O, N, L \rangle$. We propose such scheduling policies below.

We will discuss the static policy first followed a discussion of the adaptive policy. The static policy generates transmission schedules that are temporally neutral; we give an intuitive proof to show that the transmission schedule is temporally neutral. We then discuss the adaptive policy that could adapt to the long term load variations in the system, but could lead to the creation of a covert channel. We will show that this is a low bandwidth and noisy covert channel and suggest mechanisms to reduce its bandwidth.

4.1 The Basic Transmission Schedule

In Figure 4, we see an example of the generic transmission schedule schematic. The horizontal lines represent the arrival, scheduling and transmission of packets. Note that all new arrivals in the current period i are eligible for transmission only during the next period ($i+1$). All new arrivals are entered into the buffer along with previously backlogged packets (if any). Depending on the scheduling policy, a packet may or may not be scheduled for transmission. This is because even if there are more than one packet in the buffer, we may select only one packet to transmit in order satisfy the V and F restriction of $\langle V, F, O, N, L \rangle$. If a packet is not scheduled for transmission, then it is backlogged at this period, as indicated in the figure by a packet marked “**B**” which has been enqueued and not scheduled. The scheduled packets are transmitted. Observe that some of the packets scheduled in the current period have been rerouted (actually destined for another node). This creates a packet that needs to be rerouted from the intermediate node to the actual

destination as shown in the figure by a dotted line. As soon as the packet is received by the intermediate node, it enqueues the packet for transmission to the actual destination. Since we know the final traffic matrix, we can foresee such an event; therefore we can indicate the creation of the new packet in the same period. In reality, this packet may be scheduled for transmission at the earliest during the next period, which is the case shown in the figure.

Dummy packets are transmitted by the nodes if there are no new arrivals or backlogged packets destined for a particular node. To satisfy the V and F restriction of the tuple $\langle V, F, O, N, L \rangle$, a dummy packet has to be transmitted to that node; if a dummy packet is to be transmitted, it is generated and scheduled dynamically. Note that the dummy packet is never enqueued into the “true queues”. This ensures that an actual packet, if available, will never be backlogged, awaiting the transmission of a dummy packet. We can see that each node exchanges exactly one packet with every other node in the network each period and therefore the frequency of communication is the same over the entire cycle. This satisfies the F restriction of the tuple $\langle V, F, O, N, L \rangle$.

Note that the order of transmission is not the same as the order of arrival. In fact the actual order of transmission does not matter; we just have to guarantee that the order and timing of transmission is the same for each period in the cycle. The order of transmission may be predetermined or may be negotiated by the nodes. This is to satisfy the O restriction of the tuple $\langle V, F, O, N, L \rangle$.

The N criterion of a transmission is not directly affected by the transmission schedule and has to be enforced by the nodes at a higher level. Given that the transmission schedule obeys the V, F, and O restrictions, we will assume that we can enforce the N restriction by requiring that the extrinsic characteristics of the transmitted packets remain the same for the entire duration of this transmission schedule.

The lengths of successive periods are the same to satisfy the L restriction of the tuple $\langle V, F, O, N, L \rangle$. In the static policy, a period has $n(n-1)$ active slots and m idle slots and in the adaptive policy a period has $n(n-1)$ active slots and m idle slots, which can vary over time. As a result, the behavior of the L characteristic of a transmission schedule differs in the static and adaptive policies and we explain the role of active and idle slots in greater detail when we discuss

Figure 4: The Generic Transmission Schedule Schematic

Figure 5: Transmit Schedule Generated by the Static Policy

the individual policies.

From the intuitive discussion, we can see that the general transmission policy generates transmission schedules that satisfy $\langle V, F, O, N, L \rangle$, i.e., are temporally neutral.

4.2 The Static Scheduling Policy

We are given an $n \times n$ neutral traffic matrix and our goal is to develop a temporally neutral transmission schedule. The solution we propose uses slotted time to transmit packets. The period contains $n(n-1)$ active slots and m idle slots. The arrival, buffering and departure of packets in a period was discussed earlier in section 2 and the static transmission schedule is an extension of the basic transmission schedule discussed above.

Figure 5 shows the transmission schedule for a period as per the static scheduling policy. As in Figure 4, the horizontal lines indicate the arrival of new packets, scheduling and transmission of packets. In addition, we show the backlog queues associated with each node and the arrival of rerouted packets from intermediate nodes. The new arrivals and the rerouted packets received during period i are eligible for transmission during period $(i+1)$. Nodes 1 and 4 have at least one packet in the buffer eligible for transmission at the beginning of period i and are scheduled for transmission. Note that at the beginning of period i , there are no packets in the backlog buffer for nodes 2 and 3. This implies that there were no new arrivals or rerouted packets for either of the nodes during the period $(i-1)$; dummy packets are generated on behalf of nodes 2 and 3. The figure also shows the arrival of packets from nodes 3, 2, 1 and 4 destined for the local node. Thus exactly one packet is received and transmitted by each node in period i . Note that in period i , a packet is generated in the local node destined for node 4, routed via node 2. As explained in the Section 4.1, the local node immediately enqueues a packet on (the intermediate) node 2's queue and marks it "destined for node 4". In the next period, node 2 transmits this packet to the appropriate destination (node 4). This is shown in the figure as 4(via 2). The queuing of packets in the virtual queues is shown in dotted lines.

The status of queues at the beginning of period $(i+1)$ can be explained as follows: the packets that arrived for node 1, 3 and 4 in period i are enqueued in the queues associated with nodes 1,3 and 4 respectively.

The queue for node 4 also shows a packet backlogged from the previous period. The 4(via 2) arrival in period i enqueues a packet in the node 2's queue. Thus each node has at least one packet to transmit in its queue.

The timing and order of transmission (4,3,2,1) and the order of packet arrival (3,2,1,4) remain the same over all periods. The actual order is not important; the order of transmission could be something as simple as round robin or tournament order. The real issue is that we need to decide an order and ensure that it is strictly followed in all periods of the cycle. The volume of communication between each node has to be the same to satisfy the neutrality criterion. Thus we have satisfied the V, F, and O restrictions of the tuple $\langle V, F, O, N, L \rangle$. This satisfies the L restriction of $\langle V, F, O, N, L \rangle$. If we build the communication protocol on a NTCB and fix extrinsic packet characteristics like the packet size and encryption algorithm, we can be reasonably confident of satisfying the N restriction of $\langle V, F, O, N, L \rangle$. Since this is a static policy, there are exactly $n(n-1)$ active slots in the period, one for each pair of nodes in the system and m idle slots. This implies that the length of the period is fixed at $n(n-1) + m$ and remains the same at least until the end of this cycle.

We have shown that the transmission schedule generated by the policy satisfies each of the restrictions in the tuple $\langle V, F, O, N, L \rangle$. Therefore by our definition, the schedule is temporally neutral.

4.3 Adaptive Scheduling Policy

We are given an $n \times n$ neutral traffic matrix and our goal is to develop a temporally neutral transmission schedule. The solution we propose uses slotted time to transmit packets. The period contains $n(n-1)$ active slots and m idle slots. The arrival, buffering and departure of packets during a period was discussed in section 2. All aspects of the adaptive transmission policy are the same as the static policy except that we now have a variable number of idle slots in a period. The purpose of the idle slots in the period is that the scheduling algorithm can now adapt to variations in load to satisfy increased bandwidth requirements. The traffic matrix is required to be neutral and the nodes will receive one packet from every other node in the network per period. The order of transmissions is maintained the same for the entire cycle and the extrinsic characteristic of the packets do not change. Since the adaptive scheduling policy is similar to the

static scheduling policy, we can see that the V , F , O and N restrictions of $\langle V, F, O, N, L \rangle$ are satisfied in the adaptive scheduling policy. The only restriction that we cannot satisfy is the L restriction of the tuple $\langle V, F, O, N, L \rangle$ because the nodes may change the number of idle slots, changing L , the duration (length) of transmission. Though this could potentially introduce a low bandwidth, noisy covert channel, we feel that this tradeoff may be acceptable in order to have an adaptive scheduling policy.

The slot and period sizes are predetermined and are global values. We define a period to consist of the active slots where all nodes are transmitting packets according to a predetermined schedule and idle slots, the number of which the nodes may change to satisfy the (temporarily increased) bandwidth demands of the node. The number of active and idle slots (and therefore the bandwidth allocated) is decided by a global negotiation process in which all nodes participate. Depending on the current load on the system, the nodes can negotiate the active and idle time components of a period as shown in Figure 6. Once these parameters are decided, they remain constant for the duration of a cycle as opposed to the duration (length) of a period.

When a node or a group of nodes see a need to change the number of active slots to accommodate additional traffic, they initiate the negotiation process. To understand the model, details of the negotiation process are not relevant and we will not discuss it further in this paper. It is sufficient to realize that it is possible for the nodes to agree upon a new number of idle slots for future periods of the same cycle. As seen in Figure 6, after a sustained increase in the load, the nodes negotiate and decide decrease the number of idle slots per period. The number of active slots in the period remains the same, but the total period length decreases (by one slot), thereby increasing the utilization from $\frac{n(n+1)}{n(n+1)+m}$ to $\frac{n(n+1)}{n(n+1)+m-1}$. Note that the length of the period L and therefore the transmission characteristic has changed. It is this possibility that prevents us from guaranteeing the L restriction in $\langle V, F, O, N, L \rangle$ and leaves open the possibility of a covert channel.

It should be noted that no single node can affect the active and idle slot times significantly without reaching a consensus with other nodes in the network. Therefore the potential of a single node to change the transmission schedule is very limited. For example, an user may try to change the load on a particular node

in an attempt to change the transmission characteristics, which could be observed by the accomplice on the network, thus creating a covert channel. However, in response to the variation in the load, the scheduling policy initiates the negotiation protocol to decide on new transmission characteristics. Since the negotiation for new transmission characteristics is not done frequently and is a global activity, the bandwidth of this covert channel is very low. Also any eventual changes to the transmission schedule after the negotiation process is due to the cumulative effects of several individual node's (user's) actions and view of the network and the effects of any single node on the transmission characteristics is relatively minor. If a node is using all its capacity and wants to increase its traffic to a particular node by k packets, then due to the neutrality criterion, it must increase its traffic by a factor of kn . Also due to the non-local effect of rerouting[2], traffic on other nodes are affected as well and there might exist some excess capacity after negotiation. Therefore the covert channel has low bandwidth and is very noisy.

Having shown the possibility of existence of a covert channel, we now suggest mechanisms to reduce the bandwidth of the covert channel, if not eliminate it. These mechanisms are arranged in the order of increasing bandwidth of covert channel (or increasing adaptability of the scheduling policy).

- **No idle slots**

If we use the network at full capacity as allowed by the protocol, we can completely do away with the idle slots in a period and thus eliminate any possibility of covert channels according to our definition of temporally neutral transmission. However the scheduling policy degenerates to a simple static scheduling policy. Secondly, if a node is using all its capacity, i.e., there are no idle slots in a period then the scheme is very costly because the volume of true traffic may be only a small fraction of the capacity being used.

- **Renegotiate transmission characteristics at cycle boundaries**

In this option, we restrict the times at which the scheduling policy can respond to variations in the load. Since the cycle length is considerably longer than the period length, the nodes will have to buffer all the packets generated due to the additional load (in this cycle) and dispatch them at the usual rate. The nodes have to wait until the beginning of a new cycle before the period characteristics can be changed. This could introduce severe queuing delays and adversely

affect the Quality Of Service (QOS) requirements. In fact it is entirely possible that by the time a cycle terminates, the load on the network has smoothed out and there is no necessity to renegotiate the active and idle time slots. In this case, the user tried to create a covert channel, but was unsuccessful and no information was communicated at all. Since the cycle boundaries are far apart, the bandwidth of the covert channel is considerably reduced. The key to the success of this mechanism is that the transmission parameters be constant over long durations, i.e., the negotiations be few and far apart and the nodes decide to use additional bandwidth in small increments.

5 Performance Considerations

The application of the basic transmission schedule to develop the static transmission schedule is very intuitive and efficient. The computation costs are negligible and is done once only. Also there is no cost overhead in terms of the number of packets transmitted as no additional information is exchanged between the hosts to implement the protocol during actual transmission of the packets. However, at beginning of transmission, the nodes need to negotiate the transmission parameters and for long transmission sequences (i.e., for a long cycle), this cost is amortized over successive phases of the transmission, i.e., each period of the cycle. In the static scheduling policy, the tradeoff is between wasted capacity and poor QOS.

As against a completely secure static transmission policy, we can adopt a more responsive adaptive transmission policy with the penalty that there might exist a covert channel. However, by ensuring that the transmission parameters remain the same over long cycles implies that there are fewer chances to change the transmission parameters, thereby reducing the observable information on the network. Thus we have a low bandwidth, noisy covert channel.

In times of crises, the V, F, N, O and L characteristics could change significantly and the scheduling policies could suffer performance degradation. Under both scheduling policies, when the load is low or constantly varying or when the number of nodes in the network is large, the penalty incurred could be quite high. In the static scheduling policy, we have to pad the actual traffic with dummy packets. This implies that we may see too high a total load relative to the actual traffic. In the adaptive scheduling policy, variation in load causes additional traffic to be backlogged and the packets would suffer significant transmission delays.

However we feel that by assigning dynamic priorities to packets, we can ensure speedy delivery of certain packets at the expense of regular traffic.

6 Conclusion

In this paper, we start to formalize the notion of a temporally neutral transmission schedule and propose two scheduling policies: the static scheduling policy that generates temporally neutral transmission schedules and the adaptive scheduling policy that generates near temporally neutral transmit schedules. The adaptability of the adaptive scheduling policy can be traded off for the existence of a low bandwidth and noisy covert channel. Depending on specific operating environment and the degree of perceived threat to an installation by an intruder, an appropriate scheduling policy may be selected. Three useful metrics to select and evaluate a transmission schedule are the cost, the required responsiveness and the ability to tolerate covert channels.

One primary limitation in the above approach is the fact that the computation of the neutral traffic matrix (reroute quantities) is not distributed and for any real time system with reasonable expectations of fault-tolerance, we would prefer to distribute computation of the traffic matrix. Ideally we should be able to make all (reroute/padding) decisions locally.

The feasibility of assigning dynamic priority to packets and its effects on the scheduling policy must be studied and appropriate techniques to guarantee better quality of service (QOS) must be developed. Further research is also required in developing optimal scheduling policies for specific types of networks and validation of analytical results by simulation studies.

References

- [1] DOD, 1985, "Department of Defense Trusted Computer System Evaluation Criteria", DOD 5200.28-STD, Dec 1985.
- [2] Newman-Wolfe R. E, Balaji R. Venkatraman, 1991, "High Level Prevention of Traffic Analysis," *Seventh Annual Computer Security and Applications Conference*, Dec 2-6, 1991, San Antonio, Texas.
- [3] Newman-Wolfe R. E, Balaji R. Venkatraman, 1992, "Performance Analysis of a Method for High Level Prevention of Traffic Analysis,"

Eighth Annual Computer Security and Applications Conference, Nov 30 - Dec 4, 1992, San Antonio, Texas.

- [4] Ward Richard, 1989, "OSI Network Security and the NTCB," *Lecture Notes in Computer Science*, 396, Springer-Verlag, 1989, pp 67-74.

Figure 6: Changing Transmission Characteristic in the Adaptive Scheduling Policy