Covert Bits Through Queues

Pritam Mukherjee Sennur Ulukus
Department of Electrical and Computer Engineering
University of Maryland, College Park, MD 20742
pritamm@umd.edu *ulukus@umd.edu*

***Abstract*—We consider covert communication using a queuing timing channel in the presence of a warden. The covert message is encoded using the inter-arrival times of the packets, and the legitimate receiver and the warden observe the inter-departure times of the packets from their respective queues. The transmitter and the legitimate receiver also share a secret key to facilitate covert communication. We propose achievable schemes that obtain non-zero covert rate for both exponential and general queues when a sufficiently high rate secret key is available. This is in contrast to other channel models such as the Gaussian channel or the discrete memoryless channel where only $\mathcal{O}(\sqrt{n})$ covert bits can be sent over n channel uses, yielding a zero covert rate.**

I. Introduction

We consider a covert communication system, where Alice wishes to send a message W covertly to Bob using a timing channel in the presence of a warden Willie, as depicted in Fig. 1. To facilitate such covert information transfer, we allow Alice and Bob to have a shared secret key K. The channels to Bob and Willie are error-free bit pipes leading to buffers that are modeled as single-server queues with service rates of μ_1 packets per second and μ_2 packets per second, respectively. The packets arrive at both Bob's and Willie's queue simultaneously, and do not contain any covert information. When Alice does not have any covert messages to send, the arrival of the packets at the queues is modeled as a Poisson process with rate λ. In order to send the covert message W, Alice encodes W in the inter-arrival times of the packets at the queues. The warden Willie is a passive observer and tries to detect the presence of a covert message based on any *unusual patterns* in the timings of the packets observed by him. The goal of this paper is the study strategies for reliable and covert information transmission from Alice to Bob in the presence of Willie.

To that end, we first[1] assume that both queues have exponential service times; i.e., the service times of Bob's and Willie's queues are exponentially distributed with means $\frac{1}{\mu_1}$ and $\frac{1}{\mu_2}$, respectively. We exploit a result in [1], which characterizes the maximum rate of reliable information transmission, or the *capacity*, using the timing channel with an exponential service time queue. This gives us an upper bound on the *capacity* of the covert information transmission rate. Next, to ensure covertness, we want that the probability distribution of the departure times of the packets at Willie's queue remains almost the same, irrespective of whether a covert message is being sent or not. Thus, the distribution of the departure times at Willie's queue, induced by the designed codebook and the uniform choice of the covert message, closely *approximates* the distribution of the departure times when the packet arrivals are modeled as a Poisson process with rate λ. This brings us to the setting of channel resolvability [2], where we want to approximate the output distribution induced by a process with the output distribution induced by a codebook. The results of [1] and [2] together yield the proposed achievable schemes.

An interesting aspect of the result is the fact that a strictly positive rate is achievable in this case. This is in contrast to other covert channels studied in the literature, such as the covert Gaussian channel [3] and the covert discrete memoryless channel [4], [5], where only[2] $\mathcal{O}(\sqrt{n})$ bits can be sent in n channel uses, i.e., the covert rate is zero. For each of these channels, there is an *innocent* symbol which is transmitted in the absence of a covert message. In the Gaussian channel, for example, the default input symbol is zero when no communication is taking place. Analogously, for the discrete memoryless channel, it is a symbol x_0. In order to send information, one has to use at least one non-innocent symbol x_1; however, if too many[3] (more than $\omega(\sqrt{n})$) of the non-innocent symbol x_1 are sent in n channel uses, the induced output distribution differs significantly from the default output distribution induced by the input of the innocent symbol only. In the Gaussian channel, for example, using too many non-zero symbols with non-vanishing power significantly increases the output power, and the warden can detect the communication.

On the queuing timing channel, packets naturally arrive at the queue with i.i.d. exponential inter-arrival times when there is no communication using the timing channel. That is, the innocent state is a sequence of packets with i.i.d. exponential inter-arrival times. On the other hand, when communication is taking place on the covert channel, the inter-arrival times belong to a codeword drawn from a codebook. Thus, in order to detect whether communication is taking place, the warden tries to determine if the input to the timing channel, i.e., the inter-arrival times, belong to an i.i.d. exponential process, or to a codeword drawn from a fixed codebook, using its output observations. Thus, to ensure covertness, we only need to ensure *stealth* [6], which requires the output distribution induced by the codebook to approximate the default output distribution induced by an i.i.d. input process closely. Note

This work was supported by NSF Grants CNS 13-14733, CCF 14-22111, CCF 14-22129 and CNS 15-26608.

[1] Later, in Theorem 2, we consider queues with general service distributions.

[2] $f(n) = \mathcal{O}(g(n)) \Leftrightarrow \exists M, n_0$ s.t. $f(n) \leq Mg(n), \forall n \geq n_0$.

[3] $f(n) = \omega(g(n)) \Leftrightarrow \forall m > 0, \exists n_0 > 0$ s.t. $f(n) > mg(n), \forall n \geq n_0$.

that stealth does not imply covertness in Gaussian or discrete memoryless channels, since the warden can detect the transmission even if it cannot distinguish whether the received symbols belong to some codebook, or are outputs corresponding to i.i.d. inputs by observing the output distribution. In the timing channel, however, stealth implies covertness, and as in discrete memoryless channels, the stealth constraint can be met with a strictly positive rate [6]. Hence, we achieve covertness on the timing channel with a non-zero rate.

Related Work: Timing channels have been widely investigated in the context of both communication and computer systems. In most cases, the timing channel is not the primary intended means of information transfer, and timing is not usually considered a data object. Thus, timing channels, by themselves, are considered *covert* in most of the literature.

From a communication perspective, references [7], [8] provide achievable schemes for certain timing channel models and analyzes their performance using an information theoretic framework. Reference [1] analyzes the limits of reliable information transmission using the timings of packet arrivals and departures, i.e., the capacity of the timing channel for a single-server queue. In contrast to usual discrete memoryless channels, the timing channel model used in [1] has memory and is not stationary. Therefore, reference [1] employs information spectrum methods [9] in order to analyze its capacity. The exact capacity is derived for queues with exponentially distributed service times, and upper and lower bounds are provided for general queues. While reference [1] deals with continuous time arrivals, reference [10] analyzes discrete time queues and shows that the geometric service time distribution plays a role analogous to the exponential distribution for continuous time queues. Extensions to secure transmission of information using the timing channel are also available in the literature [11]. References [12]–[14] study practical code designs that approach the capacity of the timing channel.

In the context of computer systems, a timing channel analysis can often represent a side channel attack from a malicious adversary. For example, in shared event schedulers, a malicious user process can infer information about the legitimate user process' arrival patterns, and compromise the privacy or security of the legitimate user. References [15], [16] quantify the information leakage in the context of such shared event schedulers. Mitigation of such timing channel attacks has also been studied in the literature, e.g., in references [17]–[19]. Reference [20] studies the problem of information leakage through the timing channel in the framework of a game between the covert transmitter-receiver pair and a jammer who tries to disrupt the covert communication while being subject to buffer or delay constraints. Other mitigation techniques such as predictive mitigation [21] and time-deterministic replay [22] have also been explored in the literature.

II. System Model

We consider covert communication through two parallel single-server queues as shown in Fig. 1. The information in the packets is intended for both Bob and Willie. Alice also wants to send a covert message W to Bob without allowing the warden Willie to detect it. The covert message is encoded in the arrival times of the packets entering both Bob's queue and Willie's queue simultaneously. In the absence of a covert message, the packet arrivals are modeled as a Poisson process with rate λ, i.e., the inter-arrival times of the packets are exponential with mean $\frac{1}{\lambda}$. Bob's queue has a service rate of μ_1 packets/s, while Willie's queue has a service rate of μ_2 packets/s. We assume $\lambda \leq \min(\mu_1, \mu_2)$ to ensure stability of both queues.

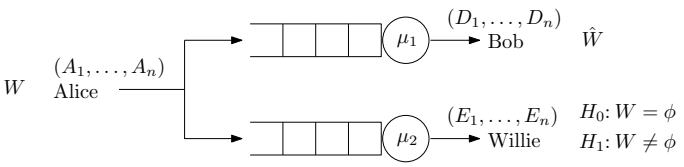

Fig. 1: Covert communication over timing channels.

In order to send the covert message W, Alice encodes it in the inter-arrival times of the packets at the queues, i.e., the covert message is encoded as a vector of n non-negative inter-arrival times $\mathbf{A} = (A_1, \ldots, A_n)$, such that the kth packet enters both queues at time $\sum_{i=1}^{k} A_i$. The intended receiver Bob and the warden Willie both observe the sequence of departure times of the packets at their respective queues. Therefore, the output of each channel is a vector of n non-negative inter-departure times, which we denote by $\mathbf{D} = (D_1, \ldots D_n)$ and $\mathbf{E} = (E_1, \ldots, E_n)$ for Bob's and Willie's channels, respectively. The departure times of the kth packet from Bob's and Willie's queues are $\sum_{i=1}^{k} D_i$ and $\sum_{i=1}^{k} E_i$, respectively.

The service times for the kth packet in Bob's and Willie's queues are denoted by S_k and T_k, respectively. S_k and T_k are mutually independent of each other and of $\mathbf{A}$, D^{k-1} and E^{k-1}. The *idling time* for the kth packet is the time elapsed between the $(k-1)$th departure and the kth arrival; if the kth arrival occurs before the $(k-1)$th departure, the idling time is zero. Denoting the idling time for the kth packet in Bob's and Willie's queues by U_k and V_k, respectively, we have,

$$U_k = \max\left(0, \sum_{i=1}^{k} A_i - \sum_{i=1}^{k-1} D_i\right) \tag{1}$$

$$V_k = \max\left(0, \sum_{i=1}^{k} A_i - \sum_{i=1}^{k-1} E_i\right) \tag{2}$$

which are deterministic functions of (A^k, D^{k-1}) and (A^k, E^{k-1}), respectively. Also the inter-departure times can now be expressed as

$$D_k = U_k + S_k \tag{3}$$

$$E_k = V_k + T_k \tag{4}$$

We have the following Markov chains:

$$D_k \to U_k \to (A^k, D^{k-1}) \tag{5}$$

$$E_k \to V_k \to (A^k, E^{k-1}) \tag{6}$$

Note that this model of the single-server queue has memory, a non-linear input-output relation and is non-stationary.

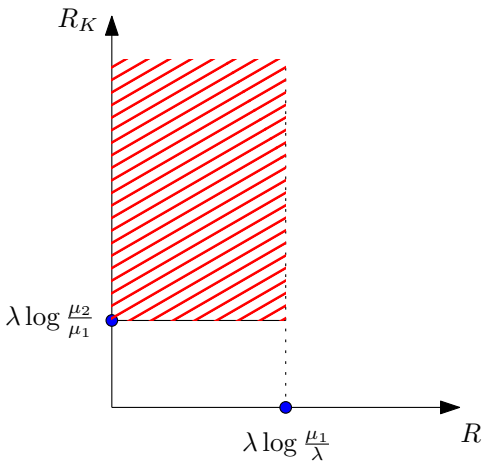

Fig. 2: Achievable (R, R_K) region when $\mu_2 > \mu_1$.

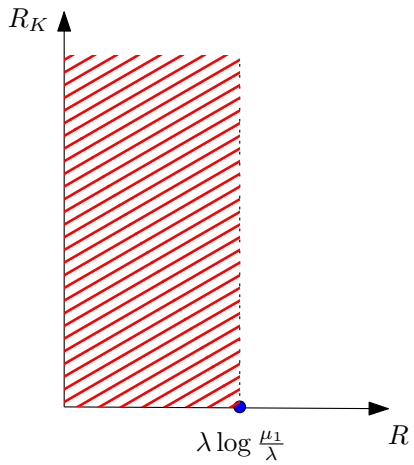

Fig. 3: Achievable (R, R_K) region when $\mu_2 < \mu_1$.

The objective in the covert communication setting is to ensure that the warden Willie cannot detect the presence of a covert message with its observations of the departure times, while also allowing the intended receiver Bob to decode the covert message with probability of error approaching zero. To facilitate this, we also allow Alice and Bob to share a secret key K which is not available to Willie. Formally, we have the following definition.

Definition 1 *An $(n, M_n, L_n, T_n, \epsilon_n, \delta_n)$ secret key assisted* covert *timing code consists of the following:*

- *a message set $\mathcal{W}_n = \{1, \ldots, M_n\}$ at Alice, from which a covert message W is picked uniformly,*
- *a key set $\mathcal{K}_n = \{1, \ldots, L_n\}$ available at Alice and Bob, from which a secret key K is picked uniformly,*
- *an encoding function (possibly stochastic) at Alice $\phi_n : \mathcal{W}_n \times \mathcal{K}_n \to \mathbf{A}$ that maps the message and the secret key into a codeword, which is a vector of n non-negative inter-arrival times such that the kth arrival occurs at $\sum_{i=1}^{k} A_k$,*
- *a decoding function at Bob $\psi_n : \mathbf{D} \times \mathcal{K}_n \to \mathcal{W}_n$ that maps the observed codeword of inter-departure times $\mathbf{D}$ and the secret key K to the decoded message $\hat{W}$, such that the probability of error*

$$\mathbb{P}(W \neq \hat{W}) \leq \epsilon_n \tag{7}$$

- *satisfies the covertness constraint at Willie,*

$$d(P_{\mathbf{E}}, Q_0^n) \leq \delta_n \tag{8}$$

where $d(P, Q)$ denotes the variational distance $||P - Q||_1$, $P_{\mathbf{E}}$ denotes the distribution of $\mathbf{E}$ in the presence of covert message, Q_0 is the default distribution of inter-departure times when no covert message is present; in our case, if the arrival process has rate $\lambda < \min(\mu_1, \mu_2)$, Q_0 will also be a Poisson process with rate λ,

- *the nth departure from Bob's queue occurs, on average, no later than T_n.*

A *covert rate* R is achievable with *secret key rate* R_K, if there exists a $(n, M_n, L_n, T_n, \epsilon_n, \delta_n)$ *covert timing code* with

$$\liminf_{n\to\infty} \frac{\log M_n}{T_n} \geq R \tag{9}$$

$$\limsup_{n\to\infty} \frac{\log L_n}{n} \leq R_K \tag{10}$$

$$\limsup_{n\to\infty} \epsilon_n = 0 \tag{11}$$

$$\limsup_{n\to\infty} \delta_n = 0 \tag{12}$$

A *covert rate* R is said to be achievable with *secret key rate* R_K at *output rate* λ, if R is achievable using a sequence of $(n, M_n, L_n, n\lambda, \epsilon_n, \delta_n)$ covert timing codes. The *covert capacity at output rate* λ, denoted by $C(\lambda)$ is the supremum of all rates R that are achievable with output rate λ.

III. Main Results

The main results of this paper are the following two theorems.

Theorem 1 *Assume both queues are $M/M/1$. Then, covert communication is possible with output rate $\lambda < \min(\mu_1, \mu_2)$ if (R, R_K) lies in the following region:*

$$R \geq 0 \tag{13}$$

$$R < \lambda \log \frac{\mu_1}{\lambda} \tag{14}$$

$$R_K > \max\left(0, \lambda \log \frac{\mu_2}{\mu_1}\right) \tag{15}$$

Fig. 2 shows the (R, R_K) region when $\mu_2 > \mu_1$, while Fig. 3 shows the (R, R_K) region when $\mu_2 < \mu_1$. Note that when $\mu_2 < \mu_1$, i.e., when the service rate of the Willie's queue is less than the service rate of Bob's queue, no positive rate secret key is required to achieve the maximum covert rate. Intuitively, Willie has a *worse* timing channel in this case than Bob. On the other hand, when $\mu_2 > \mu_1$, a secret key of sufficient rate is required for covertness, as shown in Fig. 2.

Further, note that even when $\mu_2 > \mu_1$, given a sufficient rate of secret key to enable covert communication, Alice can

communicate with Bob with the full capacity at output rate λ. In other words, there is no loss of rate due to the extra covertness constraints, as long as a secret key of sufficient rate is available to Alice and Bob. Intuitively, a secret key of sufficient length increases the size of the input codebook such that the output distribution induced by this codebook is close to that induced by an i.i.d. input, thus, ensuring covertness. If a key of sufficient rate is not available, our achievable scheme cannot guarantee covertness. Thus, the minimum rate R_K^{min} of the secret key, if required, can be considered the price of covertness, and is given by

$$R_K^{min} = \max\left(0, \lambda \log \frac{\mu_2}{\mu_1}\right) \tag{16}$$

Also note the contrast of this result with the corresponding result for Gaussian channels [3] or discrete memoryless channels [4]. In both these cases, no covert communication with positive rate is possible in general. The best achievable rate in n channel uses is $O(\omega_n\sqrt{n})$ in each case, where

$$\lim_{n\to\infty} \omega_n = 0 \tag{17}$$

$$\lim_{n\to\infty} n\omega_n = \infty \tag{18}$$

In each case, there is an *innocent* symbol, the zero symbol for the Gaussian channel or some symbol x_0 for the discrete memoryless channel, which is transmitted in the absence of a covert message, and a corresponding *default* output distribution. Hence, the default distribution at the input is the degenerate distribution with all the probability mass at the *all-zero* (or, the all x_0) codeword. This degenerate distribution, by itself, cannot be used to transmit any information. In order to transmit information, at least one non-innocent symbol x_1 must be used. Using too many (more than $\omega(\sqrt{n})$) such non-innocent symbols in n channel uses, however, results in the induced output distribution to differ significantly from the default output distribution, and covertness is lost. In our case, however, we exploit the fact that packets naturally arrive at the queue even in the absence of a covert message, i.e., in our case, the *default*, or *innocent*, setting is a sequence of packets with i.i.d. exponential inter-arrival times. Thus, we ensure *stealth* [6] in the timing channel. As in the discrete memoryless channel with a stealth constraint [6], we achieve a non-zero rate in the timing channel. Unlike discrete memoryless channels, however, stealth implies covertness in the timing channel.

While the result in Theorem 1 holds for timing channels with exponential service times, similar achievability results can be obtained for general service distributions as well. We have the following theorem.

Theorem 2 *Assume that the queues are $M/G/1$, such that Bob's queue has a service distribution P_B and Willie's queue has a service distribution P_W. Then, covert communication is possible with output rate $\lambda < \min(\mu_1, \mu_2)$ when (R, R_K) lies in the following region:*

$$R \geq 0 \tag{19}$$

$$R < \lambda \log \frac{\mu_1}{\lambda} \tag{20}$$

$$R + R_K > \max\left(0, \lambda \log \frac{\mu_2}{\mu_1} + \lambda D(P_W || e_{\mu_2})\right) \tag{21}$$

where e_{μ_2} denotes the exponential distribution with mean $\frac{1}{\mu_2}$.

Note that the rates in the theorem above are sufficient conditions to ensure covert communications; they are not necessary. As in the case of the exponential queue, no secret key is required if μ_1 is sufficiently larger than μ_2: when $\log \frac{\mu_1}{\mu_2} > D(P_W||e_{\mu_2})$.

IV. Proofs of Theorems 1 and 2

A. Proof of Theorem 1

An achievable scheme for covert communication is as follows:

Encoding: First, fix any (R, R_K) in the achievable region stated in Theorem 1. Introduce a dummy message $\tilde{W} \in \tilde{\mathcal{W}} = \left\{1, \ldots, 2^{nR_0}\right\}$, where R_0 is chosen such that

$$R + R_0 \leq \lambda \log \frac{\mu_1}{\lambda} \tag{22}$$

$$R + R_0 + R_K \geq \lambda \log \frac{\mu_2}{\lambda} \tag{23}$$

Note that this is possible as long as (R, R_K) lies in the region specified in Theorem 1. Bob will try to decode $(W, \tilde{W})$. Then:

- The transmitter generates a random codebook with $2^{n(R+R_0+R_K)}$ n-length codewords, with each code symbol being drawn in an i.i.d. fashion from an exponential distribution with mean $\frac{1}{\lambda}$. The codewords are indexed as $\mathbf{a}(w, \tilde{w}, k)$, $w \in \left\{1, \ldots, 2^{nR}\right\}$, $\tilde{w} \in \left\{1, \ldots, 2^{nR_0}\right\}$ and $k \in \left\{1, \ldots, 2^{nR_K}\right\}$.
- To send a message $w \in \left\{1, \ldots, 2^{nR}\right\}$, when the realization of the secret key is $k \in \left\{1, \ldots, 2^{nR_K}\right\}$, the transmitter chooses $\tilde{w}$ uniformly from $\left\{1, \ldots, 2^{nR_0}\right\}$, and sends $\mathbf{a}(w, \tilde{w}, k)$, i.e., the inter-arrival time of the kth packet is the kth code symbol $a_k(w, \tilde{w}, k)$.

Decodability: Bob tries to decode $(w, \tilde{w})$. Since the secret key is available at Bob, decoding with vanishing probability of error is possible as long as

$$R + R_0 \leq \text{p-}\liminf_{n\to\infty} \frac{1}{n} i_{\mathbf{A};\mathbf{D}}(\mathbf{A}; \mathbf{D}) \tag{24}$$

where

$$i_{X;Y}(X;Y) = \frac{p_{XY}(X,Y)}{p_X(X)p_Y(Y)} \tag{25}$$

is the *mutual information density* random variable, and

$$\begin{aligned} &\text{p-}\liminf_{n\to\infty} \frac{1}{n} i_{X^n;Y^n}(X^n; Y^n) \\ &= \sup\left\{r : \lim_{n\to\infty} \mathbb{P}\left[\frac{1}{n} i_{X^n;Y^n}(X^n; Y^n) < r\right] = 0\right\} \end{aligned} \tag{26}$$

It is shown in [1] that in our case when $\lambda < \mu_1$

$$\text{p-}\liminf_{n\to\infty} \frac{1}{n} i_{\mathbf{A};\mathbf{D}}(\mathbf{A}; \mathbf{D}) = \lambda \log \frac{\mu_1}{\lambda} \tag{27}$$

Therefore, the decodability constraint is satisfied.

Covertness: We note that the secret key is not available at Willie. However, we want to approximate the i.i.d. random Poisson process with rate λ by the output induced by our codebook in order to ensure covertness. From the result of [2], this is possible if

$$R + R_0 + R_K \geq \text{p-}\limsup_{n\to\infty} \frac{1}{n} i_{\mathbf{A};\mathbf{E}}(\mathbf{A};\mathbf{E}) \tag{28}$$

Again, it is shown in [1] that in our case when $\lambda < \mu_2$

$$\text{p-}\limsup_{n\to\infty} \frac{1}{n} i_{\mathbf{A};\mathbf{E}}(\mathbf{A};\mathbf{E}) = \lambda \log \frac{\mu_2}{\lambda} \tag{29}$$

Therefore, the covertness constraint is also satisfied. This completes the proof of Theorem 1.

B. Proof of Theorem 2

The achievable scheme for the general queue is similar to the scheme for the exponential queue. The encoding involves a random codebook with $2^{n(R+R_0+R_K)}$ n-length codewords, with each code symbol being drawn in an i.i.d. fashion from the exponential distribution with mean $\frac{1}{\lambda}$. To send a message $w \in \left\{1, \ldots, 2^{nR}\right\}$, when the realization of the secret key is $k \in \left\{1, \ldots, 2^{nR_K}\right\}$, the transmitter chooses a dummy message $\tilde{w}$ uniformly from $\left\{1, \ldots, 2^{nR_0}\right\}$ and sends $\mathbf{a}(w, \tilde{w}, k)$, i.e., the inter-arrival time of the kth packet is the kth code symbol $a_k(w, \tilde{w}, k)$.

Decodability: As in the $M/M/1$ case, since the secret key is available at Bob, decoding $(w, \tilde{w})$ with vanishing probability of error is possible as long as

$$R + R_0 \leq \text{p-}\liminf_{n\to\infty} \frac{1}{n} i_{\mathbf{A};\mathbf{D}}(\mathbf{A};\mathbf{D}) \tag{30}$$

It is known from [1] that capacity of the timing channel with exponential service time is the least among all service distributions. Hence, decoding is guaranteed when

$$R + R_0 \leq \lambda \log \frac{\mu_1}{\lambda} \tag{31}$$

Covertness: As in the $M/M/1$ case, covert communication is possible as long as

$$R + R_0 + R_K \geq \text{p-}\limsup_{n\to\infty} \frac{1}{n} i_{\mathbf{A};\mathbf{E}}(\mathbf{A};\mathbf{E}) \tag{32}$$

Again from [1], it is known that for the $M/G/1$ channel, the capacity of the timing channel is bounded by $\lambda \log \frac{\mu_2}{\lambda} + \lambda D(P_W || e_{\mu_2})$. Therefore, covertness is guaranteed as long as

$$R + R_0 + R_K \geq \lambda \log \frac{\mu_2}{\lambda} + \lambda D(P_W || e_{\mu_2}) \tag{33}$$

Eliminating R_0 from the decodability condition in (31) and the covertness condition in (33) yields the rate constraints in Theorem 2. This completes the proof of Theorem 2.

V. Conclusions

We introduced the notion of covert communication using a queuing timing channel in the presence of a warden. The covert message is encoded using the inter-arrival times of the packets and the legitimate receiver and the warden observe the inter-departure times of the packets from their respective queues. The transmitter and the legitimate receiver also share a secret key to facilitate covert communication. We proposed achievable schemes that obtain a non-zero covert rate for both $M/M/1$ and $M/G/1$ queues. This is in contrast to other channel models such as the Gaussian channel or the discrete memoryless channel where only $\mathcal{O}(\sqrt{n})$ covert bits can be sent over n channel uses, and therefore, the achievable covert rate is zero. We exploit the fact that *stealth* implies *covertness* in the timing channel.

References

[1] V. Anantharam and S. Verdu. Bits through queues. *IEEE Trans. on Inf. Theory*, 42(1):4–18, Jan. 1996.

[2] T. S. Han and S. Verdu. Approximation theory of output statistics. *IEEE Trans. on Inf. Theory*, 39(3):752–772, May 1993.

[3] B. A. Bash, D. Goeckel, and D. Towsley. Limits of reliable communication with low probability of detection on AWGN channels. *IEEE JSAC*, 31(9):1921–1930, Sep. 2013.

[4] M. R. Bloch. Covert communication over noisy channels: A resolvability perspective. *IEEE Trans. on Inf. Theory*, 62(5):2334–2354, May 2016.

[5] L. Wang, G. Wornell, and L. Zheng. Fundamental limits of communication with low probability of detection. *IEEE Trans. on Inf. Theory*. To appear. Also available at [arXiv:1506.03236].

[6] J. Hou and G. Kramer. Effective secrecy: Reliability, confusion and stealth. In *IEEE ISIT*, Jun. 2014.

[7] I. S. Moskowitz and A. R. Miller. The channel capacity of a certain noisy timing channel. *IEEE Trans. on Inf. Theory*, 38(4):1339–1344, Jul. 1992.

[8] I. S. Moskowitz, S. J. Greenwald, and M. H. Kang. An analysis of the timed z-channel. In *IEEE S & P*, May 1996.

[9] T. S. Han. *Information-Spectrum Methods in Information Theory*. Springer-Verlag Berlin Heidelberg, 2003.

[10] A. S. Bedekar and M. Azizoglu. The information-theoretic capacity of discrete-time queues. *IEEE Trans. on Inf. Theory*, 44(2):446–461, Mar. 1998.

[11] B. P. Dunn, M. Bloch, and J. N. Laneman. Secure bits through queues. In *IEEE ITW*, June 2009.

[12] T. P. Coleman and N. Kiyavash. Practical codes for queueing channels: An algebraic, state-space, message-passing approach. In *IEEE ITW*, May 2008.

[13] N. Kiyavash, T. P. Coleman, and M. R. D. Rodrigues. Novel shaping and complexity-reduction techniques for approaching capacity over queuing timing channels. In *IEEE ICC*, Jun. 2009.

[14] R. Sundaresan and S. Verdu. Sequential decoding for the exponential server timing channel. *IEEE Trans. on Inf. Theory*, 46(2):705–709, Mar. 2000.

[15] X. Gong, N. Kiyavash, and P. Venkitasubramaniam. Information theoretic analysis of side channel information leakage in FCFS schedulers. In *IEEE ISIT*, Jul. 2011.

[16] X. Gong and N. Kiyavash. Quantifying the information leakage in timing side channels in deterministic work-conserving schedulers. *IEEE/ACM Trans. on Netw.* To appear. Also available at [arXiv:1403.1276].

[17] S. Kadloor, N. Kiyavash, and P. Venkitasubramaniam. Mitigating timing based information leakage in shared schedulers. In *IEEE INFOCOM*, Mar. 2012.

[18] S. Kadloor, X. Gong, N. Kiyavash, and P. Venkitasubramaniam. Designing router scheduling policies: A privacy perspective. *IEEE Trans. on Signal Process.*, 60(4):2001–2012, Apr. 2012.

[19] S. Kadloor, P. Venkitasubramaniam, and N. Kiyavash. Preventing timing analysis in networks: A statistical inference perspective. *IEEE Signal Process. Mag.*, 30(5):76–85, Sep. 2013.

[20] J. Giles and B. Hajek. An information-theoretic and game-theoretic study of timing channels. *IEEE Trans. on Inf. Theory*, 48(9):2455–2477, Sep. 2002.

[21] A. Askarov, D. Zhang, and A. C. Myers. Predictive black-box mitigation of timing channels. In *ACM CCS*, Oct. 2010.

[22] A. Chen, W. B. Moore, H. Xiao, A. Haeberlen, L. T. X. Phan, M. Sherr, and W. Zhou. Detecting covert timing channels with time-deterministic replay. In *USENIX OSDI*, Oct. 2014.