Covert Communications in Packet Collision Channels

Azadeh Sheikholeslami*, Majid Ghaderi+, and Dennis Goeckel*,
*Dept. of Elec. and Comp. Engineering (ECE), Univ. of Massachusetts, Amherst
+Dept. of Computer Science, University of Calgary, Calgary, Alberta

***Abstract*—Covert communications, where a transmitter Alice wishes to hide the presence of her transmitted signal from a watchful adversary Willie, has been considered extensively in recent years. Those investigations have generally considered physical-layer models, where the adversary has access to a sophisticated (often optimal) receiver to determine whether a transmission has taken place, and have addressed the question of what rate can information be communicated covertly. More recent investigations have begun to consider the change in covert rate when Willie has uncertainty about the physical layer environment. Here, we move up the protocol stack to consider the covert rate when Willie is watching the medium-access control (MAC) layer in a network employing a random access MAC such as slotted ALOHA. Based on the rate of collisions and potentially the number of users involved in those collisions, Willie attempts to determine whether unauthorized (covert) users are accessing the channel. In particular, we assume different levels of sophistication in Willie's receiver, ranging from a receiver that only can detect whether there was a collision or not, to one that can always tell exactly how many packets were on the channel in the random access system. In each case, we derive closed-form expressions for the achievable covert rates in the system. The achievable rates exhibit significantly different behavior than that observed in the study of covert systems at the physical layer.**

***Index Terms*—Covert communications, decision theory, wireless system security**

I. Introduction

Security is an important topic in current and emerging wireless communications networks. Much of security research is focused on preventing a powerful adversary from deciphering the message *content*. However, there are applications where even the detection of the *presence* of a message by the adversary can have adverse ramifications for the sender. For example, the volume of military radio traffic can be used to infer the presence or size of a party. In commercial applications, one example is attempting to avoid detection by an authority who wishes to detect and then shut down any radio traffic between opponents. And, more specifically, the Snowden disclosures indicated both the value and the potential need to protect the "meta-data" (who is talking to who) rather than simply the message content. Hence, the topic of low probability of detection (LPD) communications has been of significant interest for many years. Recent work has coined the term "covert communications" for this field.

This work was supported by the National Science Foundation under grants ECCS-1309573 and CNS-1564057, and by DARPA under contract number HR0011-16-C-0111.

Covert communications has been traditionally obtained via spread spectrum, where a much larger bandwidth than that required for the data rate is employed so that the signal energy can be hidden in the wideband noise. However, despite significant military interest in hiding communication in such a manner, a characterization of the fundamental limits of covert communication systems was not conducted until recently. The work of [1] first addressed these limits; more recently, [2] independently and formally considered the rate at which covert communication can take place: communication from transmitter Alice to intended receiver Bob in the presence of a capable and attentive adversary Willie whose sole goal is to detect whether communication is taking place (see Figure 1). When there are additive white Gaussian noise (AWGN) channels between Alice and each of Bob and Willie, the authors of [2] found that $O(\sqrt{n})$ (and no more) bits can be sent in n channel uses covertly and reliably to Bob. Others have followed up on the work of [2] to consider other channels such as the binary symmetric channel (BSC) [3] or more general discrete memoryless channels (DMCs), as well as scaling constants [4], [5]. Thus, the basic case of Alice, Bob, and Willie with standard channels between the three has largely been characterized.

The limitation to only $O(\sqrt{n})$ bits in n channel uses, which results in a rate that scales as $\frac{1}{\sqrt{n}}$, is obviously disappointing. Since [2], authors have looked at other scenarios, and some allow for the rate to be improved. Since Willie is trying to detect Alice in the presence of background noise and interference, the covert rate can be improved by either keeping Alice's timing a secret [6] or by Willie not having an accurate characterization of the background environment [7], [8]. In the latter case, which can be obtained by employing time-varying jamming [9], Alice can send $O(n)$ bits in n channel uses to Bob without detection by Willie. These works indicate that the degree to which Willie can accurately view the transmission greatly impacts the covert communication rate.

Whereas much of the work on the foundations of covert communications has considered the physical layer, recent work has started to take a network view. In [10], the authors consider routing in a moderately-sized network to provide the optimal covert rate between a transmitter and receiver. In a sequence of papers [11], [12], the authors consider the ability to hide information in a packet stream without altering its characteristics such that a watchful adversary can detect such. In this paper, to our knowledge, we take the first look at

hiding packet transmissions from a monitor of the medium access control (MAC). In particular, if a network monitor knows the number of legitimate users in the network, then that monitor expects the contention process in a random access protocol such as slotted ALOHA to behave in a certain manner. If a covert user attempts to access the channel, then that behavior is changed and can potentially be detected by the network monitor. Our main goal is to consider the amount of information that can be conveyed accessing the channel by a collection of covert users without such a detection.

Critical to our investigation are the assumptions on what the network monitor can observe. In many cases, a network monitor might only be able to observe whether there was a collision in a given time slot, and not, for example, the number of packets involved in that collision. In other cases, the network monitor might have a more sophisticated receiver and be able to determine the number of packets involved in a collision. We will assume the goal of the monitor is simply to determine the number of packets involved in a collision and not to decode them, and thus we will term this "multiple-packet detection" (MPD). As developed in the next section, it makes sense to assume that the monitor can determine the number of packets involved in a collision up to some number K, and thus we will consider receivers with K-MPD capability. When the monitor can only determine whether there was a transmission or not, we will term this 0-MPD.

The rest of the paper is organized as follows. Section II presents the system model and metrics. Section III presents our main results, whereas Section IV provides numerical evaluation of those results. Finally, Section V provides the conclusions and future work.

II. System Model and Metrics

A. System Model

Consider a slotted random access system with N legitimate users $u_1, u_2, \ldots, u_N$, called system users, contending for the channel. Likewise, M covert users $c_1, c_2, \ldots c_M$ may (or may not) access the medium while trying to avoid detection by a network monitor termed Willie. In each time slot, each user independently flips a coin with probability of heads p_t and transmits a packet in that time slot if the result is heads. Hence, there is the possibility that multiple users will transmit in each time slot, causing a collision. Throughout this work, we will assume that the number of users is large enough so that the number of transmissions in each time slot for each type of user can be approximated by a Poisson random variable (for small p_t). That is, the number of packet transmissions in a given time slot from system (allowed) users is Poisson with rate $\lambda = Np_t$, and the number of packet transmissions in a given time slot from covert users (when they are present) is Poisson with rate $\lambda_a = Mp_t$.

Traditionally, packets involved in a collision would simply be discarded and re-transmission would be required. However, with advances in multi-user detection, e.g. successive interference cancellation [**?**], it is often possible for an advanced receiver to recover multiple packets from a collision. The ability of a receiver to *receive* multiple simultaneous transmissions is referred to as Multi-Packet Reception in the literature [**?**]. To detect covert communications, in our approach, the network monitor only needs to detect the number of packets transmitted in a time slot. In other words, the network monitor does not need to decode and recover the content of the packets, which requires a more sophisticated receiver. We will denote the ability to *detect* multiple simultaneous transmissions as Multi-Packet Detection (MPD). A receiver is called a K-MPD detector if it can detect up to K packet transmissions in a time slot. Specifically, a K-MPD detector can detect that one of the following $K+2$ events occurred in a given time slot:

- Idle: Event that no packets were on the channel during the time slot.
- Exactly l packets, $l = 1, 2, \ldots, K$ were detected: K events, with the l^{th} event being that there were l packets on the channel during the time slot.
- The event that more than K packets are involved in the collision.

That is, a K-MPD detector can always determine the number of packets involved in a collision as long as that number is less than or equal to K; when the number of packets involved in the collision is more than K, the receiver only knows that there were more than K packets on the channel but not the exact number. We will consider detectors ranging from a 0-MPD detector to a K-MPD detector, $1 \leq K < \infty$, to an ∞-MPD detector. The 0-MPD detector, which is the weakest detector we will consider at Willie, can only determine whether there was any packets on the channel or not, whereas an ∞-MPD detector, which is the most powerful detector we consider at Willie, can determine the exact number of packets involved in any collision.

B. Metrics

Willie's goal is to determine whether the covert users are active or not. Let H_0 be the hypothesis that the covert users are not present, and H_1 the hypothesis that the covert users are present. Let Willie's probability of false alarm be denoted by $\mathbb{P}_{FA}^W$, and his probability of missed detection be denoted by $\mathbb{P}_{MD}^W$. The transmissions are *covert* when the sum of probabilities of detection error is arbitrarily close to one, i.e. $\mathbb{P}_{FA}^W + \mathbb{P}_{MD}^W > 1 - \epsilon$ for any $\epsilon > 0$ [2].

The probability distribution of the network state observed by Willie when the covert users are not present is denoted by $\mathbb{P}_0$, and when the covert users are present is denoted by $\mathbb{P}_1$. Let s denote the network state. The network state represents the number of packet transmissions in a time slot. For example, $s = 0$ indicates that no packets are transmitted, while $s = 1$ indicates that exactly one packet is being transmitted. In [13], it is shown that for the optimal hypothesis test performed by Willie,

$$\mathbb{P}_{FA}^W + \mathbb{P}_{MD}^W = 1 - d_{TV}(\mathbb{P}_0, \mathbb{P}_1),$$

where $d_{TV}(\mathbb{P}_0, \mathbb{P}_1) = \frac{1}{2}\sum_s |\mathbb{P}_0(s) - \mathbb{P}_1(s)|$ is the total variation distance between $\mathbb{P}_0$ and $\mathbb{P}_1$. Hence, if

$$d_{TV}(\mathbb{P}_0, \mathbb{P}_1) < \epsilon, \text{ for all } \epsilon > 0, \tag{1}$$

covertness is maintained. In this paper, we consider the total variation distance as our covertness metric, and call the transmissions covert if the condition in (1) is satisfied. We define *covert rate* as the maximum rate of packet transmission by the covert users, λ_a, such that their transmission cannot be detected by Willie, as a function of the rate λ of packet transmission by the allowed users. We will also consider the reliable throughput obtained by the covert users as a function of λ in Section IV, although we hasten to note that λ is a system parameter and hence not available for optimization by the covert users.

III. Main Results

In this section, we present the main results of this paper. We consider four different scenarios, and will find an upper bound for the transmission rate of the covert users in each case:

1) 0-*MPD detector*: Willie can only distinguish between an idle channel and a busy channel. This is the least capable Willie, and thus we predict in this case the allowable covert rates are largest compared to those of other scenarios.
2) 1-*MPD detector*: Willie can distinguish between an idle channel and one packet transmission, but when more than one packet is transmitted over the channel, he only knows that more than one packet was on the channel.
3) K-*MPD detector*: We consider a more general case of K-MPD Willie that can detect up to K simultaneous transmissions.
4) ∞-*MPD detector*: Willie can determine the exact number of packets on the channel. This is the most capable Willie, and thus we predict in this case the allowable covert rates are the smallest compared to the other scenarios.

In each case, we characterize the probability distribution of the network state conditioned on the presence or absence of covert users. Recall that the network state denotes the number of packet transmissions in a time slot. The conditional state probability distributions are then used to compute the total variation distance, and consequently the covert transmission rate.

A. Willie with 0-MPD capability

In this case, Willie can only determine if the channel is busy or not. If the channel is busy, Willie is not able to determine how many packets are being transmitted in a time slot.

The probability distributions of the network states observed by Willie, $\mathbb{P}_0$ and $\mathbb{P}_1$, are Bernoulli distributions. Let $\mathcal{S} = \{s_0, s_1\}$ denote the set of states of each Bernoulli process, where s_0 and s_1 indicate that the channel is, respectively, idle and busy. We have:

$$\mathbb{P}_0\{s = s_0\} = 1 - \mathbb{P}_0\{s = s_1\} = e^{-\lambda},$$
$$\mathbb{P}_1\{s = s_0\} = 1 - \mathbb{P}_1\{s = s_1\} = e^{-(\lambda+\lambda_a)}.$$

Therefore,

$$\begin{aligned} d_{TV}(\mathbb{P}_0, \mathbb{P}_1) &= \frac{1}{2}\sum_{s\in\mathcal{S}} |\mathbb{P}_1(s) - \mathbb{P}_0(s)| \\ &= \frac{1}{2}\left(|e^{-\lambda} - e^{-(\lambda+\lambda_a)}| + |1 - e^{-\lambda} - (1 - e^{-(\lambda+\lambda_a)})|\right) \\ &= e^{-\lambda}(1 - e^{-\lambda_a}), \end{aligned} \tag{2}$$

where the last equality holds because for $\lambda_a \geq 0$, we have $1 - e^{-\lambda_a} \geq 0$. The covertness condition is satisfied if

$$d_{TV}(\mathbb{P}_0, \mathbb{P}_1) = e^{-\lambda}(1 - e^{-\lambda_a}) \leq \epsilon \tag{3}$$

Hence, we have established:

Result 1 (0-MPD): If ϵ and λ are sufficiently large such that $\epsilon e^{\lambda} \geq 1$, the transmissions of the covert users are covert for any $\lambda_a > 0$. If $\epsilon e^{\lambda} < 1$, the transmissions of the covert users are covert if

$$\lambda_a \leq \ln \frac{1}{1 - \epsilon e^{\lambda}}. \tag{4}$$

B. Willie with 1-MPD capability

In this section we consider Willie to have 1-MPD capability. He can detect an idle channel ($s = s_0$), a single packet transmission ($s = s_1$), or more than one packet transmission ($s = s_2$). The total variation distance between $\mathbb{P}_0$ and $\mathbb{P}_1$ can be written as,

$$\begin{aligned} d_{TV}(\mathbb{P}_0, \mathbb{P}_1) &= \frac{1}{2}\sum_{s\in\mathcal{S}} |\mathbb{P}_1(s) - \mathbb{P}_0(s)| \\ &= \frac{1}{2}\Big(|e^{-\lambda} - e^{-(\lambda+\lambda_a)}| + |\lambda e^{-\lambda} - (\lambda+\lambda_a)e^{-(\lambda+\lambda_a)}| \\ &\quad + |1 - (1+\lambda)e^{-\lambda} - 1 + (1+\lambda+\lambda_a)e^{-(\lambda+\lambda_a)}|\Big) \\ &= \frac{1}{2}\Big(|e^{-\lambda} - e^{-(\lambda+\lambda_a)}| + |\lambda e^{-\lambda} - (\lambda+\lambda_a)e^{-(\lambda+\lambda_a)}| \\ &\quad + |(1+\lambda)e^{-\lambda} - (1+\lambda+\lambda_a)e^{-(\lambda+\lambda_a)}|\Big) \\ &\leq |e^{-\lambda} - e^{-(\lambda+\lambda_a)}| + |\lambda e^{-\lambda} - (\lambda+\lambda_a)e^{-(\lambda+\lambda_a)}| \\ &= e^{-\lambda}(1 - e^{-\lambda_a}) + e^{-\lambda}|\lambda(1 - e^{-\lambda_a}) - \lambda_a e^{-\lambda_a}| \end{aligned} \tag{5}$$

where the inequality is the triangle inequality. We know $|x| \leq \max\{x, -x\}$. Hence,

$$\begin{aligned} &d_{TV}(\mathbb{P}_0, \mathbb{P}_1) \leq e^{-\lambda}(1 - e^{-\lambda_a}) \\ &+ e^{-\lambda}\max\{\lambda(1 - e^{-\lambda_a}) - \lambda_a e^{-\lambda_a}, -\lambda(1 - e^{-\lambda_a}) + \lambda_a e^{-\lambda_a}\} \\ &\leq e^{-\lambda}(1 - e^{-\lambda_a}) \\ &+ e^{-\lambda}\max\{\lambda(1 - e^{-\lambda_a}), -\lambda(1 - e^{-\lambda_a}) + 1 - e^{-\lambda_a}\} \\ &= e^{-\lambda}(1 - e^{-\lambda_a}) + e^{-\lambda}(1 - e^{-\lambda_a})\max\{\lambda, 1 - \lambda\}, \end{aligned}$$

where the second inequality follows from the fact that for $\lambda_a \geq 0$, $\lambda_a e^{-\lambda_a} \geq 0$, and $\lambda_a e^{-\lambda_a} \leq 1 - e^{-\lambda_a}$. The equality follows because $1 - e^{-\lambda_a} \geq 0$.

The covertness condition is satisfied if

$$d_{TV}(\mathbb{P}_0, \mathbb{P}_1) \leq e^{-\lambda}(1 - e^{-\lambda_a})(1 + \max\{\lambda, 1 - \lambda\}) \leq \epsilon. \tag{6}$$

Hence, we have established the following result:

Result 2 (1-MPD): If ϵ and λ are such that

$$\frac{\epsilon e^{\lambda}}{1+\max\{\lambda, 1-\lambda\}} \geq 1, \tag{7}$$

the transmissions of the covert users are covert for any $\lambda_a \geq 0$. Otherwise, the transmissions of the covert users are covert if

$$\lambda_a \leq -\ln\left(1 - \frac{\epsilon e^{\lambda}}{1+\max\{1-\lambda, \lambda\}}\right). \tag{8}$$

C. Willie with K-MPD capability

Assume that Willie is able to detect up to K simultaneous packet transmissions. If more than K packets are transmitted, then Willie will only detect a collision event. Thus, the network state observed by Willie can be one of $K+2$ states denoted by $s_0, s_1, \cdots, s_K, s_{K+1}$, where s_i (for $i = 0, 1, \ldots, K$) indicates that Willie detected i concurrent transmissions.

We have,

$$\begin{aligned} d_{TV}(\mathbb{P}_0, \mathbb{P}_1) &= \frac{1}{2}\sum_{s\in\mathcal{S}} |\mathbb{P}_1(s) - \mathbb{P}_0(s)| \\ &= \frac{1}{2}\left(\sum_{k=0}^{K}\left|\frac{\lambda^k e^{-\lambda}}{k!} - \frac{(\lambda+\lambda_a)^k e^{-(\lambda+\lambda_a)}}{k!}\right|\right. \\ &\quad \left. + \left|1 - \sum_{k=0}^{K}\frac{\lambda^k e^{-\lambda}}{k!} - 1 + \sum_{k=0}^{K}\frac{(\lambda+\lambda_a)^k e^{-(\lambda+\lambda_a)}}{k!}\right|\right) \\ &\leq \sum_{k=0}^{K}\frac{e^{-\lambda}}{k!}\left|\lambda^k - (\lambda+\lambda_a)^k e^{-\lambda_a}\right| \end{aligned} \tag{9}$$

where the inequality is the triangle inequality. Because of the absolute value in (9), it is hard to find an upper-bound for $d_{TV}(\mathbb{P}_0, \mathbb{P}_1)$ when $\lambda \geq 0$ is arbitrary. However, when $\lambda \geq K$, the term in the absolute value is greater than zero for any $\lambda_a \geq 0$, which makes the analysis of (9) easier. In this case,

$$\begin{aligned} d_{TV}(\mathbb{P}_0, \mathbb{P}_1) &\leq \sum_{k=0}^{K}\frac{e^{-\lambda}}{k!}\left|\lambda^k - (\lambda+\lambda_a)^k e^{-\lambda_a}\right| \\ &= \sum_{k=0}^{K}\frac{e^{-\lambda}}{k!}(\lambda^k - (\lambda+\lambda_a)^k e^{-\lambda_a}) \\ &\leq \sum_{k=0}^{K}\frac{e^{-\lambda}}{k!}(\lambda^k - \lambda^k e^{-\lambda_a}) \\ &= \sum_{k=0}^{K}\frac{e^{-\lambda}}{k!}\lambda^k(1 - e^{\lambda_a}). \end{aligned} \tag{10}$$

Hence, we have established:

Result 3: (K-MPD, $\lambda > K$)

$$d_{TV}(\mathbb{P}_0, \mathbb{P}_1) \leq \sum_{k=0}^{K}\frac{e^{-\lambda}}{k!}\lambda^k(1 - e^{\lambda_a}) \leq \epsilon \tag{11}$$

Thus, if λ and ϵ are such that

$$\frac{\epsilon e^{\lambda}}{\sum_{k=0}^{K}\frac{\lambda^k}{k!}} \geq 1, \tag{12}$$

the covertness condition is satisfied for any $\lambda_a \geq 0$. Otherwise, if

$$\lambda_a \leq -\ln\left(1 - \frac{\epsilon e^{\lambda}}{\sum_{k=0}^{K}\frac{\lambda^k}{k!}}\right), \tag{13}$$

the transmissions of the covert users are covert.

D. Willie with ∞-MPD capability

Now, we consider Willie with ∞-MPD capability, i.e. a Willie who can determine how many packets are being transmitted simultaneously over the channel. In this case, the probability distributions of the network states observed by Willie are given by two Poisson probability distributions: $\mathbb{P}_0 = Poisson(\lambda)$, and $\mathbb{P}_1 = Poisson(\lambda + \lambda_a)$. Using the upper-bound for total variation distance between two Poisson distributions in [14], it then follows that

$$\begin{aligned} d_{TV}(\mathbb{P}_0, \mathbb{P}_1) &\leq \min\left\{\lambda_a, \sqrt{\frac{2}{e}}(\sqrt{\lambda+\lambda_a} - \sqrt{\lambda})\right\} \\ &\leq \sqrt{\frac{2}{e}}(\sqrt{\lambda+\lambda_a} - \sqrt{\lambda}) \\ &\leq \epsilon. \end{aligned} \tag{14}$$

Thus, we have established:

Result 4 (∞-MPD): In the presence of an ∞-MPD Willie if

$$\lambda_a \leq \frac{e\epsilon^2}{2} + \epsilon\sqrt{2e\lambda}, \tag{15}$$

the transmissions of the covert users are covert.

Willie with ∞-MPD capability is the most capable Willie, and thus the covert rates in this case are the smallest compared to the covert rates in other scenarios. In fact, in contrast to the other scenarios, we see a square root law similar to that in [2] and subsequent work [15] at the physical layer.

IV. Numerical Examples

In this section, we numerically evaluate the bounds obtained in Results 1-4. In Fig. 1, the covert transmission rate is plotted versus the system nodes' transmission rate for different levels of Willie's MPD capabilities. For small values of λ, the covert rates of ∞-MPR are larger than the covert rates of 1-MPD and 5-MPD. This contradicts the expectation that when Willie has ∞-MPD capability, the covert rates should be the smallest compared to other scenarios. This is due to the fact that the curves in Fig. 1 are only upper bounds and not exact values. As expected from Results 1-4, when Willie employs 0-MPD, 1-MPD, or 5-MPD detection, the covert rate λ_a can be arbitrarily large for λ sufficiently large. Whereas at first this seems surprising (or even erroneous), the reason is clear: if λ is much larger than K, then Willie's detector observes the event of more than K packets occurring with high probability in each time slot, even when the covert users are not present. When the covert users are present, they are unlikely to cause a change in the observation at Willie, hence providing covertness. A more interesting observation is that the square root law of [2] is not observed at small λ for the 0-MPR and 1-MPR detector,

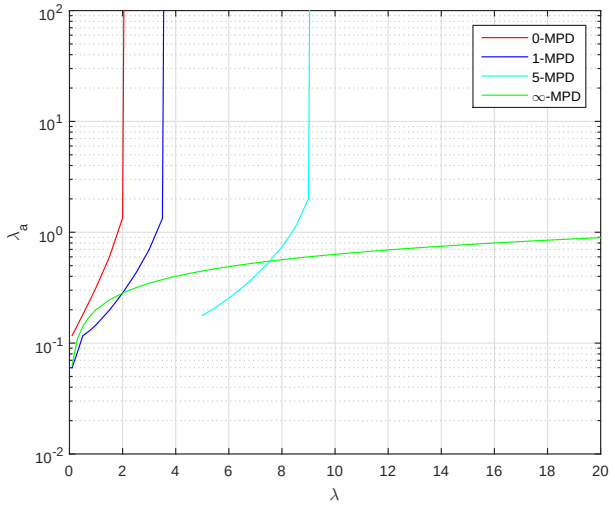

Fig. 1: Covert rate λ_a versus the rate of system nodes λ. As the MPD capability of Willie increases, the achievable covert rate decreases. Specifically, 0-MPD, which is the weakest Willie, results in the highest covert rate.

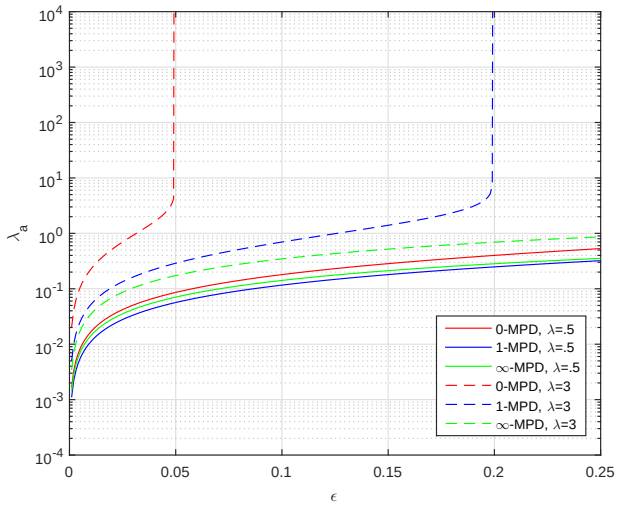

Fig. 2: Covert rate versus ϵ for different system nodes' packet transmissions rates, $\lambda = 0.5$ and $\lambda = 3$.

hence indicating the degree to which only a noisy view of the collision status of the channel can hide the presence of covert users.

In Fig. 2, the covert transmission rate versus the covertness factor, ϵ, for two different system nodes' transmission rates is shown. For small system nodes' transmission rates (e.g. $\lambda = 0.5$), as ϵ increases the covertness constraint becomes less restrictive and thus larger λ_a's can be achieved. For large system nodes' transmission rates (e.g. $\lambda = 3$), for the 0-MPD and 1-MPD detectors when ϵ is assumed to be sufficiently large again the presence of covert users is hidden for any λ_a.

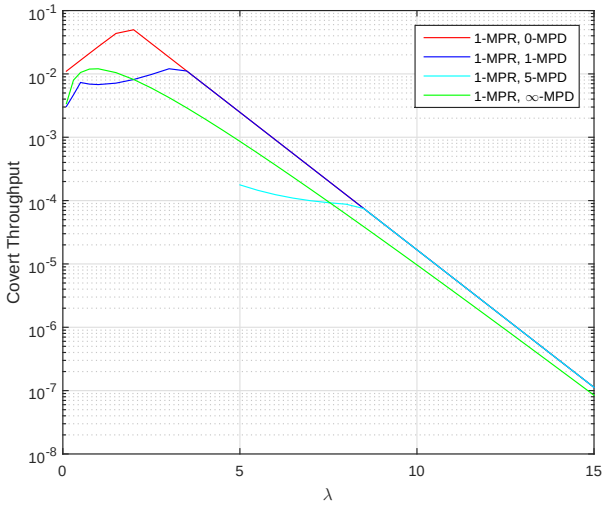

Fig. 3: Throughput of covert users versus the transmission rate of system nodes λ. The receiver for covert users has 1-MPR capability. In all cases, as λ increases, so does the covert throughput up to a certain λ after which the covert throughput decreases. For large values of λ, the covert throughput is restricted by packet collisions in the network.

Note that the covert rate depicted is restricted by collisions. Hence, in order to calculate the covert throughput, in addition to the covertness constraint we should consider the success probability of packet transmissions. In the case of a receiver with 1-MPR capability, the success probability is the probability that only one covert transmitter transmits. Fig. 3 shows the covert throughput of the covert users versus system nodes' transmission rate. Note that when λ is sufficiently large, users' transmissions are covert for any $\lambda_a \geq 0$, and thus the achievable throughput is restricted only by the receiver's multi-packet reception capability.

Finally, whereas we have provided closed-form analytical expression for the covert rate achievable by the covert users, it is possible to evaluate the covert rate numerically. This is shown in Figure 4. Compared to Figure 3, it shows the accuracy of our analytical bounds, in particular in demonstrating the difference in behavior when Willie has a K-MPD receiver, $K < \infty$.

V. Conclusions and Future Work

The fundamental limits of covert communications have been considered extensively in recent years at the physical layer for the scenario of covert transmitter Alice, receiver Bob, and capable and attentive warden Willie who attempts to detect Alice. Here, we consider for the first time the medium access control (MAC), where a number of covert users are attempting to access the channel without detection by warden Willie, who is observing the channel collision process. We consider a variety of receivers at Willie, ranging from one that can only determine whether the channel was idle or busy (0-MPD),

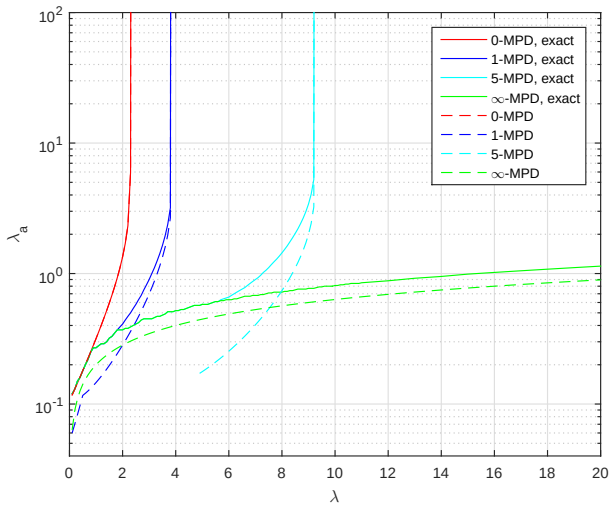

Fig. 4: Exact covert rate versus rate of system nodes' packet transmissions when $\epsilon = 0.1$.

to one that always knows the number of packets involved in a collision (∞-MPD). In the latter case, the results follow much of what has been found at the physical layer, where the rate of the covert users is restricted roughly to the square root of the rate of the system users. However, for a K-MPD detector, $K < \infty$, the throughput grows much faster than the square root of λ, thus indicating the degree to which Willie's blindness to the channel state allows for covert transmission.

REFERENCES

[1] V. Korzhik, G. Morales-Luna, and M. H. Lee, "On the existence of perfect stegosystems," in *Proc. 4th IWDW*, Siena, Italy, Sep. 2005, pp. 30–38.

[2] B. Bash, D. Goeckel, and D. Towsley, "Limits of reliable communication with low probability of detection on AWGN channels," *IEEE Journal on Selected Areas in Communications*, vol. 31, no. 9, pp. 1921–1930, 2013.

[3] P. H. Che, M. Bakshi, and S. Jaggi, "Reliable deniable communication: Hiding messages in noise," in *Proc. IEEE ISIT*, 2013, pp. 2945–2949.

[4] M. R. Bloch, "Covert communication over noisy channels: A resolvability perspective," *IEEE Transactions on Information Theory*, vol. 62, no. 5, pp. 2334–2354, 2016.

[5] L. Wang, G. W. Wornell, and L. Zheng, "Fundamental limits of communication with low probability of detection," *IEEE Transactions on Information Theory*, vol. 62, no. 6, pp. 3493–3503, 2016.

[6] B. A. Bash, D. Goeckel, and D. Towsley, "Covert communication gains from adversary's ignorance of transmission time," *IEEE Transactions on Wireless Communications*, vol. 15, no. 12, pp. 8394–8405, 2016.

[7] S. Lee and R. Baxley, "Achieving positive rate with undetectable communication over AWGN and Rayleigh channels," in *Proc. IEEE ICC*, June 2014, pp. 780–785.

[8] S. Lee, R. J. Baxley, M. A. Weitnauer, and B. Walkenhorst, "Achieving undetectable communication," *IEEE Journal on Selected Topics in Signal Processing*, vol. 9, no. 7, pp. 1195–1205, Oct 2015.

[9] T. V. Sobers, B. A. Bash, S. Guha, D. Towsley, and D. Goeckel, "Covert communication in the presence of an uninformed jammer," *IEEE Transactions on Wireless Communications*, vol. 16, no. 9, pp. 6193–6206, 2017.

[10] A. Sheikholeslami, M. Ghaderi, D. Towsley, B. A. Bash, S. Guha, and D. Goeckel, "Multi-hop routing in covert wireless networks," *IEEE Transactions on Wireless Communications*, vol. 17, no. 6, pp. 3656–3669, 2018.

[11] R. Soltani, D. Goeckel, D. Towsley, and A. Houmansadr, "Covert communications on Poisson packet channels," in *Communication, Control, and Computing (Allerton), 2015 53rd Annual Allerton Conference on*, 2015, pp. 1046–1052.

[12] ——, "Covert communications on renewal packet channels," in *Communication, Control, and Computing (Allerton), 2016 54th Annual Allerton Conference on*, 2016, pp. 548–555.

[13] E. L. Lehmann and J. P. Romano, *Testing statistical hypotheses*. Springer Science & Business Media, 2006.

[14] P. Ruzankin, "On the rate of Poisson process approximation to a Bernoulli process," *Journal of Applied Probability*, vol. 41, no. 1, pp. 271–276, 2004.

[15] B. A. Bash, A. H. Gheorghe, M. Patel, J. L. Habif, D. Goeckel, D. Towsley, and S. Guha, "Quantum-secure covert communication on bosonic channels," *Nature communications*, vol. 6, 2015.