Diffie's Wireless Phone: Heterodyning-Based Physical-Layer Encryption

Jerry T. Chiang
Advanced Digital Sciences Center
Singapore
jerry.chiang@adsc.com.sg

Yih-Chun Hu
University of Illinois at Urbana–Champaign
Urbana, IL, USA
yihchun@illinois.edu

***Abstract*—We propose a physical-layer encryption scheme inspired by Diffie's telephone system. Instead of the usual *XOR-then-modulate* encryption scheme, we propose using *modulate-then-add*, and show that the decryption operation can be implemented using the circuit operation of heterodyning. We then show that a slightly modified superheterodyne receiver performs decryption at no additional cost compared to existing receivers.**

Our proposed cryptosystem is significant to the research community for the following reasons: 1) Our physical-layer encryption scheme uniquely outperforms other physical-layer security protocols by guaranteeing positive conditional secrecy capacity as long as Bob's signal-to-interference-and-noise ratio is above a threshold, even if Eve's channel condition is the same or significantly better than Bob's or if the channel between Alice and Bob is static; and 2) Our physical-layer encryption scheme shows that by removing a filter that the wireless circuit community has long considered to be necessary in the superheterodyne design, the modified receiver offers intriguing security features.

I. Introduction

In [5], Massey described a simple cryptosystem originally proposed by Diffie, which was never published:

> Diffie's secrecy system exploits the existence of a very large number, 2^K, of telephones that can be dialed by anyone. Note that a K bit telephone number suffices to identify each of these telephones. The i-th telephone, when dialed, plays back a recorded binary sequence R_i of length N (where $N >> K$) that was obtained exclusively for that telephone by coin-tossing. The secret key Z is a K-bit telephone number that is equally likely to be that of any of the 2^K telephones. The secret key is known to both the sender and intended recipient, but not of course to the enemy cryptanalyst. The system works as follows: When he wishes to send an N-bit plaintext X to the intended recipient, the sender dials the telephone number Z and obtains the random sequence R_Z. The sender then adds R_Z to X using component-wise modulo-two addition (just as in the $L = 2$ Vernam cipher) to produce the N bit cryptogram Y that he then sends over an insecure channel to the intended recipient. The recipient, upon receipt of Y, also dials the telephone number Z, obtains the same random sequence R_Z, and subtracts this from Y component-wise modulo-two (which is the same as addition) to obtain the plaintext X.

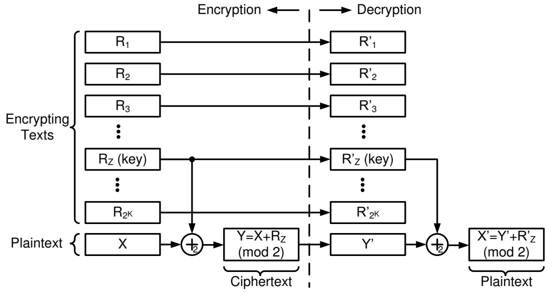

Fig. 1: Illustration of Diffie's proposed phone cryptosystem. In this paper, the "$+_2$" operator is understood to be bit-wise additive-modulo-2 (i.e. XOR).

Fig. 1 illustrates Diffie's telephone system. It is trivial to see that Diffie's telephone system satisfies what Maurer described as *conditionally provably secure* [7]: without knowing Z, and given the option to try only a small fraction of all telephone numbers (say, $\varepsilon = t/2^K$), then with probability $1 - \varepsilon$, the attacker gains no information on the plaintext[1].

In this paper, we use the wireless channels to replace the physical telephones in Diffie's telephone system. We then modify the encryption and transmission operations from "XOR-then-modulate" to "modulate-then-add." We call our modified physical-layer encryption scheme, "Diffie's Wireless Phone" (DWP). By employing "modulate-then-add," we show that the decryption operation can be implemented using a modified superheterodyne receiver, which we overview in Section II-C; however, the crypto lemma is no longer applicable, and we will discuss the security concerns later in this paper.

Our proposed protocol makes the following contributions:

1) Our physical-layer encryption scheme uniquely outperforms other physical-layer security protocols by guaranteeing positive conditional secrecy capacity as long as Bob's signal-to-interference-and-noise ratio (SINR) is above a threshold, even if Eve's channel condition is significantly better than Bob's or if the channel between Alice and Bob is static; and
2) Our scheme shows that by removing a filter that the wireless circuit community has long considered to be necessary in the superheterodyne design, the modified receiver offers intriguing security features.

[1] This follows from the "Crypto Lemma" [2].

978-0-7695-5124-1/13 $31.00 © 2013 IEEE
DOI 10.1109/MILCOM.2013.87

TABLE I: List of cryptosystem-related variables

Variable	Definition
X	plaintext
x	plain signal
Y	ciphertext
y	cipher signal
Z	shared secret
K	length of Z (in bits)
N	length of X (in bits)
R_i	encrypting text
r_i	encrypting signal
$R_{i,e}$	portion of encrypting text used for encryption in Maurer's scheme
$r_{i,e}$	portion of encrypting signal used for encryption in Maurer's Wireless Phone scheme

TABLE II: List of communication-related variables

Variable	Definition
f_c	carrier frequency
f_{LO}	frequency of local oscillator
f_{IF}	intermediate frequency
f_{IM}	image frequency
C_s	secrecy capacity
SNR$_i$	signal-to-noise ratio observed by i
P	total transmission power
W	channel bandwidth between Alice and Bob
$d_{A,B}$	distance between A and B
α	pathloss exponent
η_Z	power allocated to transmit encrypting symbol, normalized with respected to the power allocated to ciphersymbol
β	scaling factor of the cipher signal
σ_i^2	power of the additive white Gaussian noise observed by i

II. Background

In this section, we provide the necessary background in fully understanding the DWP cryptosystem and the analysis of its security. For ease of reference, we first tabulate the variables used in this paper. Readers versed in signal processing, digital communication, and physical-layer security can freely skip Section II-B, Section II-C, and Section II-D, respectively.

A. Variable Nomenclature

In this paper, we refer to a string of data bits as *text* (e.g. plaintext and ciphertext), and refer to modulated waveform as *signal* (e.g. plain signal and cipher signal). We point out that in the original Diffie's phone system, which XORs-then-modulates, the cipher signal can be demodulated to recover the ciphertext; however, in the DWP system, which modulates-then-adds, there is no corresponding ciphertext, and demodulating the cipher signal results in nonsense. Table I and Table II list the variables used in this paper.

B. The Heterodyne Operation

Heterodyning is the operation of mixing (multiplying) a signal with a sinusoid, so as to move the signal to another frequency band. Consider an input signal with carrier frequency f_c (i.e. $s(t)e^{-j\omega_c t}$, where $\omega_c = 2\pi f_c$), and a local oscillator that produces a sinusoid with frequency f_{LO}. Mixing the input signal with the output of the local oscillator produces two identical copies of the input signal, each halved in amplitude, and with the new carrier frequencies equal to $f'_c \in \{f_c \pm f_{\text{LO}}\}$:

$$\left(s(t)e^{-j\omega_c t}\right)\cos(\omega_{LO}t) = \frac{s(t)}{2}\left(e^{-j(\omega_c-\omega_{LO})t} + e^{-j(\omega_c+\omega_{LO})t}\right).$$

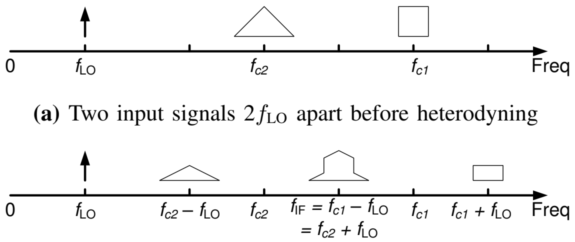

(a) Two input signals $2f_{\text{LO}}$ apart before heterodyning

(b) Output of heterodyning the input signals with local oscillator

Fig. 2: Relationship between f_c, f_{LO}, and f_{IF}. After heterodyning, the triangular input signal (s_2) at frequency $f_{c2} = f_{\text{IF}} - f_{\text{LO}}$ and the block input signal (s_1) at frequency $f_{c1} = f_{\text{IF}} + f_{\text{LO}}$ are summed together at frequency $f_{\text{IF}} = f_{c1} - f_{\text{LO}} = f_{c2} + f_{\text{LO}}$.

Of the two signal copies created by heterodyning, typically only one is desired, and its carrier frequency is called the *intermediate frequency* (f_{IF}); the other copy is filtered out before reaching other parts of the receiver system.

There are two signals on two different input frequencies, $f_{\text{IF}} + f_{\text{LO}}$ and $f_{\text{IF}} - f_{\text{LO}}$, such that heterodyning using the same local oscillator would each produce a copy at f_{IF}:

$$\begin{aligned}&\cos(\omega_{LO}t) \times \left(s_1(t)e^{-j(\omega_{IF}-\omega_{LO})t} + s_2(t)e^{-j(\omega_{IF}+\omega_{LO})t}\right)\\ &= \frac{s_1(t)+s_2(t)}{2}e^{-j\omega_{IF}t} + \\ &\quad \frac{s_1(t)}{2}e^{-j(\omega_{IF}-2\omega_{LO})t} + \frac{s_2(t)}{2}e^{-j(\omega_{IF}+2\omega_{LO})t}.\end{aligned} \tag{1}$$

To illustrate, in Fig. 2a, we show a block signal s_1 at f_{c1} and a triangular signal s_2 at f_{c2}. In Fig. 2b, we show that by mixing the signals with frequency f_{LO}, we can sum the two input signals at the intermediate frequency f_{IF}.

C. The Superheterodyne Receiver

Almost all of today's tunable radio receivers use the *superheterodyne* receiver design. A superheterodyne receiver, instead of down-converting a passband signal directly to the baseband, mixes the incoming signal with a variable-frequency local oscillator so the center frequency of the incoming signal is down-converted (or less commonly, up-converted) to a fixed intermediate frequency (f_{IF}). The signal is then down-converted to the baseband for demodulation.

Conventionally, only one of the two signals at frequencies $f_{\text{IF}} + f_{\text{LO}}$ and $f_{\text{IF}} - f_{\text{LO}}$ is desired, and a superheterodyne receiver would need to filter out the other signal (known as the *image signal*) by placing a *preselector filter* between the antenna and the first mixer. Fig. 3 illustrates the main functional blocks of a superheterodyne receiver. The preselector filter has long been considered a necessity in the receiver design in order to reject undesired image signals; we show that the removal of this filter does not necessarily render the receiver useless, but opens up the possibility to providing cryptographic features.

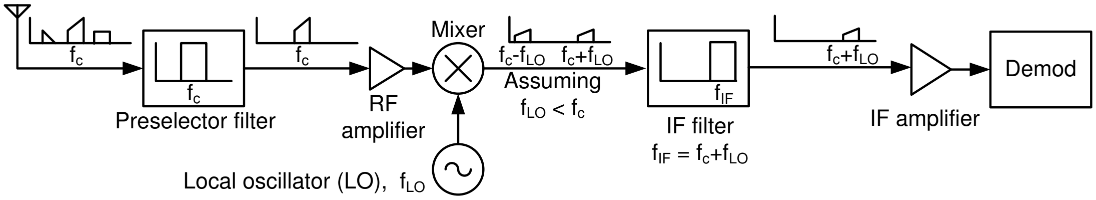

Fig. 3: Illustration of the superheterodyne receiver

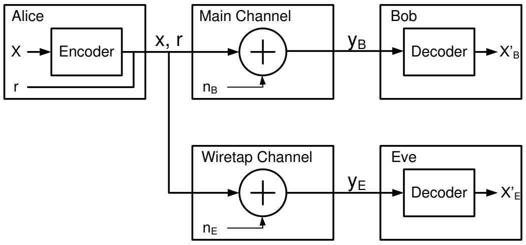

Fig. 4: Illustration of the wiretap channel

D. Physical-Layer Security

Most prior studies concerning physical-layer security focus on analyzing the *secrecy capacity* – the maximum rate achievable between Alice and Bob, two parties that wishes to communicate with each other, such that the data yields no information to an eavesdropper Eve. The secrecy capacity is an on-average measure: on average, each channel use between Alice and Bob yields some number of bits of secret data; however, one single occasion of channel use may or may not enable Alice and Bob to extract any secret bits. Fig. 4 illustrates the channel model between Alice, Bob, and Eve.

Wyner showed that, if the channel between Alice and Eve is noisier than that between Alice and Bob, then Alice and Bob enjoy a positive secrecy capacity; in Wyner's analysis model, the channel between Alice and Eve is referred to as the *wiretap channel* [11]. Leung and Hellman showed that if the channels are corrupted by Gaussian noise, the secrecy capacity between Alice and Bob is simply the channel rate of the channel between Alice and Bob minus that between Alice and Eve [4].

The rate of an additive white Gaussian noise (AWGN) channel can be determined by the Shannon-Hartley theorem:

$$\text{capacity} = W \times \log_2 (1 + \text{SNR}), \tag{2}$$

where SNR is the signal-to-noise ratio of the output of the channel, and W is the channel bandwidth. The secrecy capacity between Alice and Bob given a Gaussian wiretap channel is thus simply:

$$C_s = \max\left(W \times \log_2\left(\frac{1+\text{SNR}_B}{1+\text{SNR}_E}\right), 0\right),$$

where SNR_B (SNR_E) is Bob's (Eve's) observed signal-to-noise ratio.

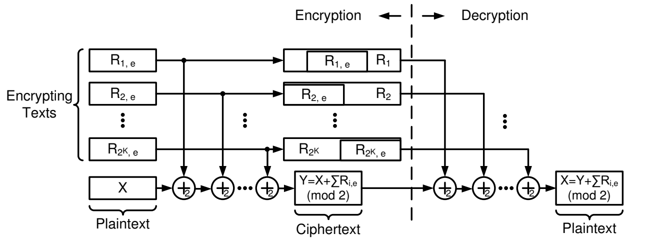

Fig. 5: Block diagram of Maurer's secure cipher.

III. Related Work

A. Provable Secrecy

Massey and Ingemarsson proposed the Rip van Winkle cipher in which Alice and Bob share a short secret key Z [6]. Alice generates a long sequence of random text R, then forms the ciphertext Y by XORing the plaintext of length N with the $(Z+1)$-th to $(Z+N)$-th bits of R. Alice then sends both the ciphertext as well as the random text to Bob. Bob, knowing Z and received both R and Y, delays R by Z, and decrypt the ciphertext by XORing Y with the first N bits of the delayed R. However, without knowing Z, Eve is forced to make random guesses on what Z is, and if the guesses are incorrect, Eve gains no information on the plaintext.

Maurer and Cachin subsequently proposed several provably-secret ciphers based on the bounded-storage adversary model [7], [1]. In these improved cryptosystems, Alice sends a $2^K \times (T+N)$ matrix of random bits to Bob, where $(T+N)$ is the length of each encrypting text and the key Z specifies for each row of the random text, the starting position of a N-bit long subsequence. The bit sequences are then XORed together to form an N-bit *encryption text*. Alice then encrypts her message by XORing the message with the encryption text. Since the attacker needs to gather all 2^K pieces in order to learn any bit of the message, the probability of key leakage decreases exponentially with respect to K. We illustrate Maurer's protocol in Fig 5.

B. Physical-Layer Security

Following Wyner's work, Maurer proposed that Alice and Bob can use a random source and a "public discussion" channel to agree on an information-theoretically secure private key, so long as the public discussion channel is authenticated, and both the channel between Alice and the random source and the channel between Bob and the random source are more reliable than the channel between Eve and the random source [8]. Maurer's protocol exploits the fact that the channel between Eve and the random source is the weakest, thus if Alice and Bob can agree, using the public channel, using only the random bits that both reliably observed, Eve likely would suffer from errors, resulting in positive secrecy capacity. The research community has proposed several theoretical extensions [3], [10], [12].

With similar concept as our proposed protocol, Negi and Goel proposed that by adding artificial noise (self-jamming), the transmitter and the receiver can increase their secrecy capacity [9]. Artificial noise is a double-edged sword that seeks to increase the secrecy capacity between Alice and Bob by degrading Eve's channel condition more than degrading Bob's. Our protocol is similar to these protocols in that we are adding noise to the plain signal to form the cipher signal; however, we also transmit a copy of our artificial noise over a secret channel, thereby enabling Bob to escape the adverse impact of self-added noise.

IV. Attacker Model

Ultimately, we assume that over the plaintext duration, the attacker can only examine and store a small fraction of all the information over the channel between Alice and Bob. This assumption implies all of the following three assumptions:

1) The attacker is able to simultaneously decrypt using random signals from $t << 2^K$ channels; otherwise the attacker has a non-negligible probability to simply search for the key by decrypting a large number of channels.
2) The frequency-adjustment time of the attacker is non-negligible; otherwise the attacker could search for the key by exhaustively trying all channels in sequence.
3) The attacker's storage or access to storage is limited. This implication is equivalent to the standard "bounded memory model" employed by Maurer and numerous prior work. We added limited "access" to storage to highlight the fact that while memory is much cheaper today than 20 years ago, the access time has not improved significantly.

V. Diffie's Wireless Phone

A. *Diffie's Wireless Phone*

Diffie's phone system can be generalized as a system in which Alice and Bob communicate over 2^K+1 *data channels*. Instead of actual phones, we can divide a wireless frequency spectrum into 2^K+1 wireless channels. Alice then modulates 2^K different random texts $R_i, \forall i \in [1, 2^K]$ and send the i-th modulated encrypting signal r_i over the i-th channel. Alice then modulates the ciphertext $Y = X \oplus R_Z$ and sends the cipher signal y over the last channel. Bob demodulates both r_Z and y to recover R_Z and Y, respectively. Bob then XORs R_Z and Y to recover the plaintext.

We modify the above generalized Diffie's phone system: instead of *XOR-then-modulate*, we use *modulate-then-add*; we call this proposed scheme *Diffie's Wireless Phone* (DWP). We first let Alice construct an encryption signal r_Z, which does not necessarily follow any particular distributions, i.e. it can be a Gaussian noise-like signal or the modulated output of a random codebook. Alice then modulates the plaintext X to get the plain signal x, and subtracts from it a weighted version of the encryption signal to form the cipher signal $y = x - \eta_Z r_Z$. Alice then scales the encryption signal (and cipher signal) by multiplying it with $\sqrt{1-\beta^2}$ (β), and transmits it over the Z-th (2^K+1-th) channel. Alice should fill the other channels with noise to prevent Eve from recovering the key Z by sensing the spectrum.

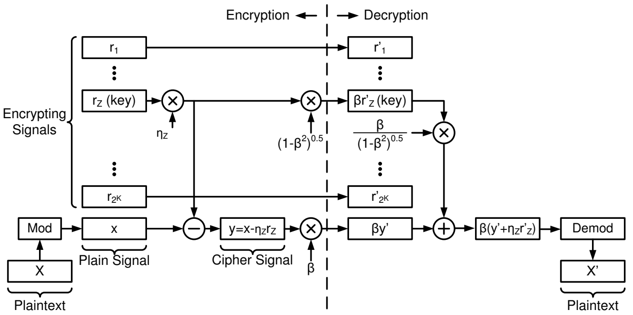

Fig. 6: Block diagram of DWP.

Without considering the noise or pathloss, to recover the plaintext, Bob multiplies the scaled encrypting signal on the Z-th channel by $\frac{\beta}{\sqrt{1-\beta^2}}$. Bob then converts both the scaled encrypting signal and the cipher signal to the same frequency, sums the two signals and demodulates to recover the plaintext X. Fig. 6 illustrates the DWP protocol.

In the original Diffie's phone system, $\beta^2 = 1 - \beta^2 = 0.5$ and $\eta_Z^2 = 1$. By reusing these parameters, Bob can eliminate the first multiplication and can readily implement the DWP cryptosystem using the superheterodyne receiver design. On Bob's side, let the center frequency of the k-th channel be f_k. Then by setting the intermediate frequency at $f_{\text{IF}} = \frac{f_{2^K+1}+f_Z}{2}$, the cipher signal and the encrypting signal are located at each other's image frequency. Thus, Bob can recover the plain signal x at the intermediate frequency by prefiltering neither y nor r_Z, and setting $f_{\text{LO}} = \frac{f_{2^K+1}-f_Z}{2}$. Bob then down-converts the output to the baseband and demodulates to recover the plaintext X. We observe that knowing f_{LO} is equivalent to knowing Z. Through the rest of this paper, unless otherwise specified, we assume the DWP decryption is implemented using the modified superheterodyne receiver.

B. *Practical Limitations of DWP*

The single decryption step that is critical to successfully recovering the plaintext is the heterodyne operation. The first obstacle we face is that Alice and Bob's local oscillators must be frequency and phase synchronized. Frequency drifts and phase incoherence both result in higher demodulation error and lower delivery ratio.

The removal of the preselector filter also incurs some real-life hurdles. In particular, without rejecting out-of-band noise before all signals reach the RF amplifier, the RF power can saturate the analog-to-digital converter (ADC), thereby reducing the usable dynamic range of the ADC, and significantly degrading Bob's observed signal-to-noise ratio.

Another known physical-layer issue in wide band communication is that the RF channel characteristics vary depending on carrier frequency because the same physical displacement equates to different RF path lengths on different frequencies. Thus, the encrypting signal and the cipher signal may lose synchronization over distance regardless of the quality of Bob's local oscillator.

If Alice and Bob are close to each other, Bob can receive the signals with high SNR and the signals traveled over relatively small distance so the difference in RF paths is not significant. That is, the physical proximity ameliorates both the degraded ADC dynamic range and the loss in signal synchronization.

VI. Evaluation

In this section we study the effect of frequency and phase disagreements between Bob's and Alice's local oscillators. We simulated the proposed DWP system using GNU Radio[2].

In our simulations, we send a 100 kB file, modulated using DBPSK, over an AWGN channel; the modulated output is x. We then modulate a random binary source also using DBPSK, and use the modulated output as the encrypting signal r_Z. We let $f_{\text{LO}} = 100$ kHz. We form the baseband cipher signal $y = x - r_Z$, and up-convert r_Z to another channel with center frequency $2f_{\text{LO}}$. We sum the cipher signal in baseband and r_Z in passband before sending over the channel.

In our first experiment, we let Bob's local oscillator be 0 to 1.5 Hz faster than Alice's local oscillator, and vary the noise level of the AWGN channel from 0 to $\frac{P}{4}$, where P is Alice's total transmission power. We define the *nominal SNR* to be Bob's total received power divided by Bob's observed noise ($\text{SNR}_N = \frac{Pd_{AB}^{\alpha}}{\sigma_B^2}$); since our simulation does not take pathloss into account, the AWGN level corresponds to nominal SNR of infinity to 6 dB. In our second experiment, we fix the nominal SNR of the AWGN channel to 20 dB, and let Bob's local oscillator be zero to a quarter period lagging in phase compared to the received signal.

Fig. 7a shows our simulation results from the first experiment. We plot a different line for each different SNR, and plot the delivery ratio versus the offset in local oscillator frequency. The delivery ratio is defined as

$$\frac{\text{the number of bytes successfully decoded}}{\text{the number of bytes transmitted (100 kB)}}.$$

Our result shows that if Bob's oscillator is more than 1.5 Hz faster than Alice's oscillator, the delivery ratio reaches 0. Additionally, low SNR intuitively also deteriorates the performance.

Fig. 7b shows our simulation results from the second experiment. Our results show that if Bob's oscillator is lagging the incoming signal by less than 0.2 period, he can still successfully receive nearly all data. However, as the phase offset increases to around 0.25 period, the decryption operation stops working, and the delivery ratio quickly drops to 0.

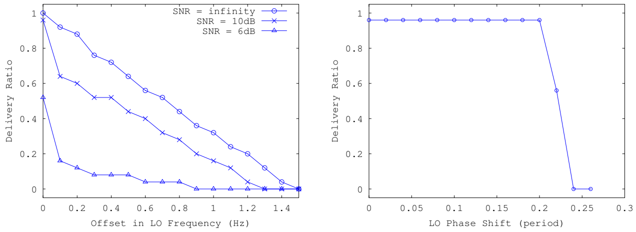

(a) The delivery ratio versus offset in local oscillator frequency for different nominal SNRs

(b) The delivery ratio versus local oscillator phase offset

Fig. 7: Effects on delivery ratio due to imperfect local oscillator

VII. Security Analysis

A. Conditional Secrecy Capacity of Diffie's Wireless Phone

With probability $1 - \frac{1}{2^K}$, in the original Diffie's phone system, Eve obtains no useful information; however, in the DWP cryptosystem, Eve obtains some information of the plaintext due to the use of modulate-then-add. We turn to information theory to capture the secrecy of the DWP cryptosystem.

In the original Diffie's phone system, the secrecy capacity simply equals the channel capacity between Alice and Bob since Eve obtains no information about the plaintext without correctly guessing Z. To determine the channel capacity, we examine Bob's observed signal-to-noise ratio, which Alice maximizes by splitting her power P evenly between y and r_Z. Without considering fading or shadowing, Bob's received signal power is $||y||^2 d_{AB}^{-\alpha}$ where d_{AB} is the distance between Alice and Bob, and α is the pathloss exponent. Bob thus observes $\text{SNR}_{B,D} = \frac{\frac{P}{2} d_{AB}^{-\alpha}}{\sigma_B^2} = \frac{1}{2}\text{SNR}_N$.

In the DWP scheme, suppose $||x||^2 = ||r_Z||^2$, i.e. $\eta_Z = 1$. Assuming x and r_Z are independent, then $||y||^2 = 2||x||^2$ and $||x||^2 = \frac{P}{3}$. In DWP, after summing y and r_Z, x' is corrupted by *twice* the observed noise power, thus Bob's observed signal-to-noise ratio is $\text{SNR}_{B,W} = \frac{\frac{P}{3} d_{AB}^{-\alpha}}{2\sigma_B^2} = \frac{1}{6}\text{SNR}_N$. For ease of analysis, we let the random encrypting signal be AWGN, and thus Eve's observed signal-to-noise ratio without correctly guessing Z is $\text{SNR}_{E,W} = \frac{||x||^2 d_{AB}^{-\alpha}}{\eta_Z^2 ||r_Z||^2 d_{AB}^{-\alpha} + \sigma_E^2} \le \frac{1}{\eta_Z^2} = 1$.

Thus, compared to Diffie's original phone system, DWP *reduces* the secrecy capacity by

$$\begin{aligned} &W \log_2 (1 + \text{SNR}_{B,D}) - W \log_2 (1 + \text{SNR}_{B,W}) + \\ &\quad W \log_2 (1 + \text{SNR}_{E,W}) \\ \le &W (\log_2 (3) + 1) . \end{aligned}$$

I.e., we are losing at most $\log_2(6) \approx 2.585$ bits per hertz (of channel bandwidth) per second (of time) of secrecy capacity. This loss is substantial in a high-noise environment, but may be tolerable if Alice and Bob are close.

B. Optimal Power Allocation between the Cipher Signal and the Encrypting Signal

In this section, we explore the tradeoff in secrecy capacity of allocating power between y and r. In order to decrypt using our modified superheterodyne receiver, the plain signal must be

[2] http://gnuradio.org

equal to the sum of the encryption signal and the cipher signal. We show that the secrecy capacity is maximized between Alice and Bob when $\beta^2 = 0.5$, and thus the superheterodyne receiver readily provides the optimal performance.

Lemma 7.1: Given $y = \beta x - \sqrt{1-\beta^2}\eta_Z r_Z$, the secrecy capacity between Alice and Bob is maximized when $\beta^2 = 1 - \beta^2 = 0.5$.

Proof: As illustrated in Figure 6, we fix a scaling factor for r_Z so that the power of the encrypting signal is some η_Z^2 times the power of x. As discussed in last section, the signal-to-noise ratio observed by Eve is at most $\frac{1}{\eta_Z^2}$, and is irrelevant to β; thus the maximum secrecy capacity is achieved when the channel capacity between Alice and Bob is maximized, which is equivalent to maximizing the SNR observed by Bob.

Bob's observed SNR is

$$\text{SNR}_B = \frac{d_{AB}^{-\alpha}\frac{P}{2\eta_Z^2+1}\beta^2}{\left(1+\frac{\beta^2}{1-\beta^2}\right)\sigma_B^2} = \beta^2(1-\beta^2)\cdot\frac{\text{SNR}_N}{2\eta_Z^2+1}$$

The last term is constant with respect to β, and thus SNR_B is maximized when the product of the first two terms is maximized at $\beta^2 = 1 - \beta^2 = 1/2$. ■

C. Power Allocation between the Plain Signal and the Encrypting Signal

In this section, we discuss the relationship between η_Z and the secrecy capacity between Alice and Bob.

Lemma 7.2: There exists a threshold to Bob's observed SNR ($\text{SNR}_{B,T}$) such that for all $\text{SNR}_B > \text{SNR}_{B,T}$, there exists $0 < \eta_Z^2 < \infty$ that yields positive secrecy capacity between Alice and Bob.

Proof: Given that y and r are equally scaled (i.e. $\beta^2 = 0.5$), the SNR observed by Bob is $\text{SNR}_B = \frac{1}{4}\frac{\text{SNR}_N}{2\eta_Z^2+1} = \frac{\text{SNR}_N}{8\eta_Z^2+4}$, and the SNR observed by Eve is $\text{SNR}_E \leq \frac{1}{\eta_Z^2}$. Prior work has shown that Alice and Bob enjoy positive secrecy capacity as long as the channel condition between Alice and Bob is better than that between Alice and Eve, thus it suffices to show that there exists a threshold above which Bob's observed SNR can be greater than Eve's with a positive η_Z^2.

Assume Bob's observed SNR is $\text{SNR}_B = 8 + \varepsilon$ for some positive ε. Bob's observed SNR is larger than Eve's if $\eta_Z^2 > \frac{4}{\varepsilon}$. Since $\varepsilon > 0$, η_Z is finite. On the other hand, if $\text{SNR}_B \leq 8$, Bob's observed SNR can be greater than Eve's only if $\eta_Z^2 < 0$, which has no physical meaning. Thus, $\text{SNR}_{B,T} = 8$. ■

D. Impersonation Attacks

One danger of our protocol is that without the key Z, Mallory can still impersonate as Alice by sending plain signal over the cipher channel, and nothing over the keying channel. Bob would then sum the plain signal on the cipher channel with the pure noise from the keying channel, and receive Mallory's plaintext intact; an equivalent problem exists in the original Diffie's telephone system if the Z-th telephone plays an all-0 encrypting text. Because of this attack, the plaintext should be authenticated at higher layers.

VIII. Conclusion

In this paper, we propose a physical-layer cryptosystem, inspired by Diffie's telephone system. In exchange of the usual *XOR-then-modulate* encryption scheme, we propose using *modulate-then-add*, and show that the corresponding decryption scheme can be implemented by slightly modifying the popular superheterodyne receiver. Our physical-layer encryption scheme uniquely outperforms other physical-layer security protocols by guaranteeing positive conditional secrecy capacity even if Eve's channel condition is the same or significantly better than Bob's, or if the channel between Alice and Bob is static. Additionally, being a physical-layer scheme, our proposed protocol is orthogonal to any source-encrypting cryptosystem, thereby adding another layer of security.

We simulate our cryptosystem using GNU Radio, and outlined several physical constraints that the receiver must meet in order to decrypt correctly. We carefully examine the secrecy capacity of our proposed protocol compared to the original Diffie's telephone system, and provide an upper bound on the loss of secrecy capacity.

References

[1] C. Cachin and U. Maurer, "Unconditional security against memory-bounded adversaries," in *Proceedings of the 17th Annual International Cryptology Conference on Advances in Cryptology (CRYPTO '97)*, ser. Lecture Notes in Computer Science, vol. 1294, Aug. 1997, pp. 292–306.

[2] J. Forney, G. D., "On the role of MMSE estimation in approaching the information-theoretic limits of linear Gaussian channels: Shannon meets Wiener," in *Proceedings of the Allerton Conference on Communication, Control, and Computing*, 2003.

[3] I. Hero, A.O., "Secure space-time communication," *IEEE Transactions on Information Theory*, vol. 49, no. 12, pp. 3235–3249, Dec. 2003.

[4] Y.-C. Leung and M. Hellman, "The "gaussian" wire-tap channel," *IEEE Transactions on Information Theory*, vol. 24, no. 4, pp. 451–456, Jul. 1978.

[5] J. L. Massey, "The relevance of information theory to modern cryptography," in *Proceedings of the 1990 Bilkent International Conference on Communications, Control, and Signal Processing*, 1990, pp. 176–182.

[6] J. L. Massey and I. Ingemarsson, ""the rip van winkle cipher" - a simple and provably computationally secure cipher with a finite key," in *Proceedings of the IEEE International Symposium on Information Theory (ISIT)*, 1985, p. 146, (Abstract).

[7] U. Maurer, "Conditionally-perfect secrecy and a provably-secure randomized cipher," *Journal of Cryptology*, vol. 5, no. 1, pp. 53–66, 1992.

[8] ——, "Secret key agreement by public discussion," *IEEE Transactions on Information Theory*, vol. 39, no. 3, pp. 733–742, May 1993.

[9] R. Negi and S. Goel, "Secret communication using artificial noise," in *Proceedings of the IEEE 62nd Vehicular Technology Conference, (VTC Fall '05)*, vol. 3, Sep. 2005, pp. 1906–1910.

[10] A. Thangaraj, S. Dihidar, A. Calderbank, S. McLaughlin, and J.-M. Merolla, "Applications of LDPC codes to the wiretap channel," *IEEE Transactions on Information Theory*, vol. 53, no. 8, pp. 2933–2945, Aug. 2007.

[11] A. D. Wyner, "The wire-tap channel," *Bell System Technical Journal*, vol. 54, pp. 1355–1387, Oct. 1975.

[12] A. Zuquete and J. Barros, "Physical-layer encryption with stream ciphers," in *IEEE International Symposium on Information Theory (ISIT)*, 2008, pp. 106–110.