Decentralized Coded Caching for Wiretap Broadcast Channels

Sarah Kamel[1,2], Michèle Wigger[1] and Mireille Sarkiss[3]
[1]LTCI, Télécom ParisTech, Université Paris-Saclay, 75013, Paris, France.
Email: {sarah.kamel, michele.wigger}@telecom-paristech.fr
[2]ISEP-LISITE, 75006 Paris, France.
[3]CEA, LIST, Communicating Systems Laboratory, BC 173, 91191 Gif-sur-Yvette, France.
Email: mireille.sarkiss@cea.fr

***Abstract*—We consider a K-receiver wiretap broadcast channel where K_w receivers are weak and have cache memories and K_s receivers are strong and have no cache memories. We derive an upper bound on the secrecy rate-memory tradeoff under a joint secrecy constraint and under decentralized caching. In contrast to previous works, prefetching in our scheme is purely decentralized and receivers randomly sample from a random key stream available at the transmitter and from the files in a library. For small cache sizes, the performance of our scheme improves with increasing length of the random key stream. For moderate and large cache sizes, a small key stream suffices to perform close to the information-theoretic limit of the system.**

I. Introduction

This paper investigates the secrecy rate-memory tradeoff of wiretap erasure broadcast channels (BCs) with cache memories at the receivers in the presence of an external eavesdropper. Cache memories can be used to prefetch fragments of popular contents or secret keys during off-peak periods (called *prefetching*) with the goal to reduce and secure network traffic during subsequent peak-traffic periods (called *delivery phase*). The main challenge of these systems is that during the prefetching, the users have not yet decided which files they will download during the delivery phase, and thus cache memories should be filled with contents that are relevant for all possibly demanded files. The pioneering work in [1] proposed to diversify cache contents across users, so as to allow coded-multicast communication that can simultaneously serve multiple users during the delivery phase.

In this framework, one generally distinguishes between *centralized* and *decentralized* caching scenarios. In centralized caching [1], the set of active users is known in advance, even before prefetching starts. In contrast, in decentralized caching [2], the set of users that will be active during the delivery phase, is unknown to the transmitter during prefetching. In this scenario, users thus have to fill their cache memories in an uncoordinated way by independently downloading information that they store in their cache memories.

The focus of this work is on a decentralized caching scenario with an external eavesdropper that can access the delivery communication but not the prefetching and that is not allowed to learn anything about the set of all possibly demanded files. Such a setup has first been studied in [3] assuming that all legitimate receivers have cache memories of equal sizes and delivery communication takes place over a common noise-free bit-pipe to all receivers and to the eavesdropper. The scheme proposed in [3] is based on the coded caching scheme in [1], but where additional random keys are prefetched in the receivers' cache memories and used to secure delivery transmissions. The work in [3] deviates from the classical decentralized setup in assuming that prefetching of keys can be performed in a centralized (coordinated) manner.

In this work, we propose a purely decentralized prefetching protocol where also key distribution is performed in an decentralized manner. Moreover, we consider the more general erasure BC model for the delivery communication, where each transmitted bit is erased at each of the legitimate receivers and at the eavesdropper with a certain probability. We also backoff from the assumption that all legitimate receivers have cache memories of same size. For simplicity, we partition the set of receivers into two groups: a first group of weak receivers with same erasure probability δ_w and same cache size, and a second group of strong receivers with same erasure probability $\delta_s \leq \delta_w$ and without cache memories. Previous works [4], [5] have shown that generally, this is the most interesting scenario from a technical point of view, because it allows for new innovative coding techniques. Practically, this corresponds to a scenario where some of the users (e.g., femto base-stations) are further apart from the main BS than others, and for a given application one decides to occupy more storage space at weaker users than at stronger users.

In this paper, we present two new coding schemes for the described scenario. In our proposed prefetchings, receivers independently sample from the library of all files and/or from a random key stream available at the transmitter. The sampled key bits are then combined with wiretap BC coding to secure delivery communication. As we see, the length of this key stream (i.e., the amount of randomness available at the transmitter) influences the performance of our coding schemes. Our coding schemes further build on coded caching and the more general secure piggyback coding scheme in [5].

We analytically derive the upper bounds on the secrecy rate-memory tradeoff (the smallest securely achievable rate for given cache sizes) corresponding to our new coding schemes. Numerical simulations show that when the key stream is sufficiently long, then these upper bounds are close to the best

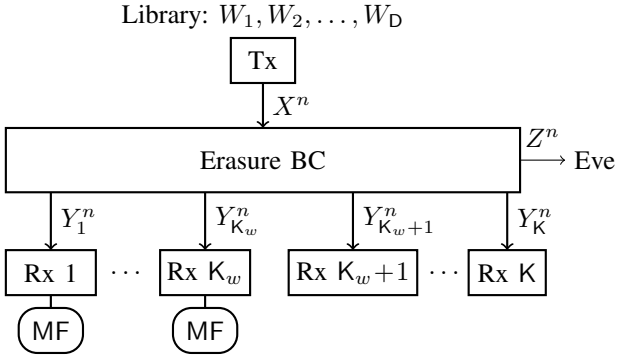

Fig. 1. Erasure BC with K_w weak receivers with cache memories of size MF bits, $\mathsf{K}_s = \mathsf{K} - \mathsf{K}_w$ strong receivers without caches, and an eavesdropper.

upper bound on the *centralized secrecy rate-memory tradeoff*. Moreover, for small cache memories, the upper bounds are also close to a lower bound on the secrecy rate-memory tradeoff, thus establishing the tradeoff almost exactly.

In a line of previous works [5], [6], we have addressed similar questions for centralized caching. Our previous work on decentralized caching [7] was based on the weaker secrecy requirement that the eavesdropper is not allowed to learn anything about any of the files in the library *individually*. In this case, there is no need for prefetching secret keys in cache memories, parts of non-demanded files can be used as onetime pads to secure delivery communication.

II. Problem Definition

Consider the erasure BC with a single transmitter, K receivers and an eavesdropper in Figure 1. The input alphabet of the BC is $\mathcal{X} := \{0,1\}$ and the K receivers and the eavesdropper have same output alphabet $\mathcal{Y} := \mathcal{X} \cup \Delta$, where Δ indicates the loss of a bit. The K receivers are partitioned into sets $\mathcal{K}_w := \{1, \ldots, \mathsf{K}_w\}$ and $\mathcal{K}_s := \{\mathsf{K}_w+1, \ldots, \mathsf{K}\}$, and the receivers in each of the sets have same channel statistics. The K_w receivers in set $\mathcal{K}_w$ are weak and have same erasure probability $\delta_w > 0$, and the $\mathsf{K}_s = \mathsf{K} - \mathsf{K}_w$ receivers in $\mathcal{K}_s$ are strong and have same erasure probability $\delta_s > 0$, where

$$0 < \delta_s \leq \delta_w < 1. \tag{1}$$

The eavesdropper is assumed to be weaker than all legitimate receivers. So, its erasure probability δ_z satisfies

$$0 < \delta_s \leq \delta_w \leq \delta_z < 1. \tag{2}$$

Assumption $\delta_z \leq \delta_w$ is only made for simplicity. Similar results hold when $\delta_z > \delta_w$.

The transmitter can access a library of $\mathsf{D} > \mathsf{K}$ independent files (messages) $W_1, \ldots, W_{\mathsf{D}}$, each consisting of F i.i.d. random bits. We assume that all files in the library have same popularity. It also has access to a *key stream* S consisting of

$$\mathsf{F}_{\text{key}} := \alpha_{\max}\mathsf{F}$$

i.i.d. random bits, for a given parameter $\alpha_{\max} \geq 0$ that is determined by the hardware in the system.

Each weak receiver is equipped with a local cache memory of size MF bits. Strong receivers have no cache memories.

Every receiver $k \in \mathcal{K} := \{1, \ldots, \mathsf{K}\}$ demands exactly one file W_{d_k} from the library. So, $d_k \in \mathcal{D} := \{1, \ldots, \mathsf{D}\}$ describes the demand of Receiver k, and $\mathbf{d} := (d_1, \ldots, d_{\mathsf{K}}) \in \mathcal{D}^{\mathsf{K}}$ the *demand vector* of all the receivers.

Communication takes place in two phases: a *decentralized prefetching phase* where each weak receiver fills its cache memory with randomly chosen bits from the library and the random key stream S, and a *centralized delivery phase* where the demanded files W_{d_k}, for $k \in \mathcal{K}$, are conveyed to the receivers. During the placement phase, the demand vector $\mathbf{d}$ is unknown to the transmitter and the receivers. As is standard for decentralized caching, the cache placement at a given receiver cannot depend on the number of receivers K (or K_w) in the system. That means, each weak receiver $i \in \mathcal{K}_w$ computes its cache content V_i by means of a universal prefetching function $g : \left\{1, \ldots, 2^{\mathsf{F}}\right\}^{\mathsf{D}} \times \left\{1, \ldots, 2^{\mathsf{F}_{\text{key}}}\right\} \times \Theta \to \left\{1, \ldots, 2^{\mathsf{MF}}\right\}$:

$$V_i := g(W_1, \ldots, W_{\mathsf{D}}, \mathsf{S}, \theta_i), \tag{3}$$

where θ_i is a random seed stored locally at Receiver i.

Prior to the delivery phase, the demand vector $\mathbf{d}$ as well as the realization of all random seeds $\theta_1, \ldots, \theta_{\mathsf{K}_w}$ are learned by the transmitter, all legitimate receivers, and the eavesdropper. The centralized delivery phase takes place over

$$n = R_{\text{sec}}\mathsf{F}$$

uses of the erasure BC, so R_{sec} denotes the secrecy delivery rate. For a given demand vector $\mathbf{d}$, the transmitter thus sends

$$X^n = f_{\mathbf{d}}(W_1, \ldots, W_{\mathsf{D}}, \mathsf{S}, \theta_1, \ldots, \theta_{\mathsf{K}_w}), \tag{4}$$

for some choice of the encoding function $f_{\mathbf{d}} : \left\{1, \ldots, 2^{\mathsf{F}}\right\}^{\mathsf{D}} \times \left\{1, \ldots, 2^{\mathsf{F}_{\text{key}}}\right\} \times \Theta^{\mathsf{K}_w} \to \mathcal{X}^n$ that can depend on the demand vector $\mathbf{d}$ as well as on the realizations of the random seeds.

Each weak receiver $i \in \mathcal{K}_w$ decodes its demanded message W_{d_i} based on its observed binary erasure channel (BEC) outputs $Y_i^n := (Y_{i,1}, \ldots, Y_{i,n})$ and its cache content V_i:

$$\hat{W}_i := \varphi_i(Y_i^n, V_i), \qquad i \in \mathcal{K}_w, \tag{5}$$

for some function $\varphi_i : \mathcal{Y}^n \times \mathcal{V} \to \left\{1, \ldots, 2^{\mathsf{F}}\right\}$. Each strong receiver $j \in \mathcal{K}_s$ decodes its demanded message based only on the observed outputs Y_j^n:

$$\hat{W}_j := \varphi_j(Y_j^n), \qquad j \in \mathcal{K}_s, \tag{6}$$

for some function $\varphi_j : \mathcal{Y}^n \to \left\{1, \ldots, 2^{\mathsf{F}}\right\}$. Notice that all functions $\varphi_1, \ldots, \varphi_{\mathsf{K}}$ can depend on the demand vector $\mathbf{d}$ and the realizations of the random seeds $\theta_1, \ldots, \theta_{\mathsf{K}_w}$.

A decoding error occurs whenever $\hat{W}_k \neq W_{d_k}$, for some $k \in \mathcal{K}$. We consider the worst-case probability of error over all feasible demand vectors

$$\mathrm{P}_e^{\text{Worst}} := \max_{\mathbf{d} \in \mathcal{D}^{\mathsf{K}}} \mathrm{P}\left[\bigcup_{k=1}^{\mathsf{K}} \left\{\hat{W}_k \neq W_{d_k}\right\}\right]. \tag{7}$$

In the described communication, the set of *all files*

$W_1, \ldots, W_{\mathsf{D}}$ needs to be kept secret from an external eavesdropper that observes the (delivery) channel outputs Z^n, but has no access to the cache memories. So, here Z^n is the result of passing the transmitters' (delivery) channel inputs X^n through a BEC with erasure probability $\delta_z > 0$.

Definition 1. *A tuple* $(R_{\text{sec}}, \mathsf{M}, \alpha_{\max})$ *is* securely achievable*, if for every* $\epsilon > 0$ *and sufficiently large* F*, there exist prefetching, encoding, and decoding functions so that*

$$\mathrm{P}_e^{\text{Worst}} \leq \epsilon, \tag{8a}$$

$$\frac{1}{n} I(W_1, \ldots, W_{\mathsf{D}}; Z^n, \mathbf{d}, \theta_1, \ldots, \theta_{\mathsf{K}_w}) < \epsilon. \tag{8b}$$

Definition 2. *Given cache memory size* M *and* $\alpha_{\max}$*, the* secrecy rate-memory tradeoff $R^\star_{\text{sec}}(\mathsf{M}, \alpha_{\max})$ *is the smallest rate* R_{sec} *so that the tuple* $(R_{\text{sec}}, \mathsf{M}, \alpha_{\max})$ *is securely achievable:*

$$R^\star_{\text{sec}}(\mathsf{M}, \alpha_{\max}) := \inf \big\{ R_{\text{sec}} \colon \ (R_{\text{sec}}, \mathsf{M}, \alpha_{\max}) \text{ securely achievable} \big\}. \tag{9}$$

Without cache memory, i.e., $\mathsf{M} = 0$, the secrecy rate-memory tradeoff does not depend on $\alpha_{\max}$ and is [8]:

$$R^{(0)}_{\text{sec}} := R^\star_{\text{sec}}(0, \alpha_{\max}) = \frac{\mathsf{K}_w}{\delta_z - \delta_w} + \frac{\mathsf{K}_s}{\delta_z - \delta_s}. \tag{10}$$

III. Main Results

If

$$\alpha_{\max} \geq \frac{\mathsf{K}_w(1-\delta_z)}{1-\delta_w},$$

define the secrecy rate-memory pair

$$\mathsf{M}^{(\text{Key})} := \alpha_{\max} \left(1 - \sqrt[\mathsf{K}_w]{1 - \frac{\mathsf{K}_w(1-\delta_z)}{\alpha_{\max}(1-\delta_w)}} \right), \tag{11a}$$

$$R^{(\text{Key})}_{\text{sec}} := \frac{\mathsf{K}_w}{1-\delta_w} + \frac{\mathsf{K}_s}{\delta_z - \delta_s}. \tag{11b}$$

For any $0 \leq \alpha \leq \alpha_{\max}$, define the secrecy rate-memory pairs

$$\mathsf{M}^{(\text{Mixed})}(\alpha) := \frac{(\mathsf{D}+\alpha)(1-\delta_z)}{\alpha(1-\delta_w) + (1-\delta_z)}, \tag{12a}$$

$$R^{(\text{Mixed})}_{\text{sec}}(\alpha) := \frac{\frac{1-p}{p}\left[1-(1-p)^{\mathsf{K}_w}\right]}{1-\delta_w} + \frac{\mathsf{K}_s(1-p)^{\mathsf{K}_w}}{\delta_z - \delta_s} + \frac{\left[1-(1-p)^{\mathsf{K}_w}\right]\left[\mathsf{K}_s(1-\delta_w) - (\delta_w - \delta_s)\frac{1-p}{p}\right]^+}{(1-\delta_w)(\delta_z - \delta_s)}, \tag{12b}$$

where $[x]^+ := \max\{0, x\}$, and

$$p := \frac{1-\delta_z}{\alpha(1-\delta_w) + (1-\delta_z)}. \tag{13}$$

Theorem 1. *The secrecy rate-memory tradeoff is upper bounded by the lower convex hull of the secrecy rate-memory pairs satisfying* (11) *and* (12)*:*

$$R^\star_{\text{sec}}(\mathsf{M}, \alpha_{\max}) \leq \text{lower hull} \Big\{ \Big(R^{(0)}_{\text{sec}}, \mathsf{M} = 0\Big), \Big(R^{(\text{Key})}_{\text{sec}}, \mathsf{M}^{(\text{Key})}\Big), \text{ and } \Big(R^{(\text{Mixed})}_{\text{sec}}(\alpha), \mathsf{M}^{(\text{Mixed})}(\alpha)\Big) \ \forall \alpha \in [0, \alpha_{\max}] \Big\}. \tag{14}$$

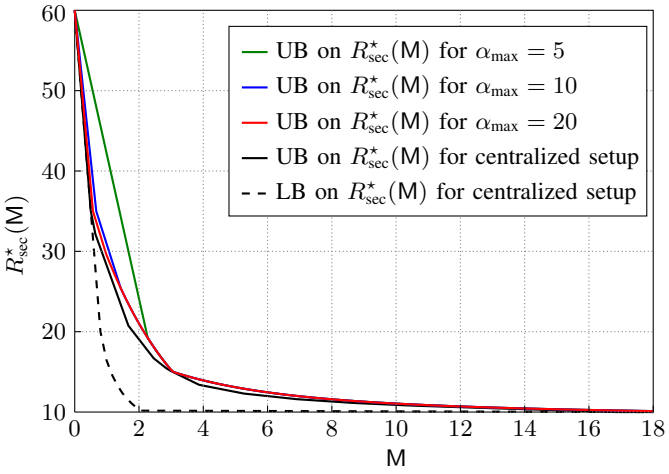

Fig. 2. Upper and lower bounds on $R^\star_{\text{sec}}(\mathsf{M})$ for $\delta_w = 0.6$, $\delta_s = 0.3$, $\delta_z = 0.8$, $\mathsf{K}_w = 10$, $\mathsf{K}_s = 5$, and $\mathsf{D} = 20$.

Notice that $\mathsf{M}^{(\text{Mixed})}(\alpha)$ and $R^{(\text{Mixed})}_{\text{sec}}(\alpha)$ are monotonically increasing and decreasing functions for $\alpha \in [0, \alpha_{\max}]$.

Figure 2 depicts the upper bound in Theorem 1 for different values of $\alpha_{\max}$. The figure also shows upper and lower bounds for the centralized case. The upper bound is obtained by improving [5, Theorem 2] using ideas from [9] and the lower bound is from [5, Theorem 1]. (Being a lower bound on the centralized case, it is also a lower bound on the decentralized case.) For small cache sizes, the performance of our scheme improves with growing $\alpha_{\max}$, i.e., with the length of the key stream S. The reasons is that a large key size implies a large probability of receivers prefetching long independent secret keys that the transmitter can apply as one-time pads to secure communication. For $\alpha_{\max}$ exceeding 10 (i.e., $\mathsf{F}_{\text{key}} \approx 10\mathsf{F}$), the gap between our upper bound and that of the centralized setup is negligible. For moderate and large cache sizes, the gap between the decentralized and the centralized upper bounds is negligible irrespective of $\alpha_{\max}$. We also observe that both upper bounds are close to the lower bound.

Figure 3 compares the upper bound in Theorem 1 with the lower bound from [3]. Note that in [3], a *centralized* prefetching of secret keys is applied before the centralized delivery phase. The required coordination between the transmitter and receivers during the prefetching phase is however incompatible with the decentralized caching paradigm. In our scheme, prefetching of key bits is also decentralized and does not necessitate any coordination with the server. Nevertheless, Figure 3 shows that our scheme outperforms that of [3]. This is due to the fact that our scheme is a *joint cache-channel coding scheme* where the encoder and the decoders simultaneously exploit the cache content and the channel statistics. In contrast, the scheme in [3] is designed for error-free links and we combine it here with a standard BC code for communications to weak receivers. In a subsequent phase, we use a wiretap

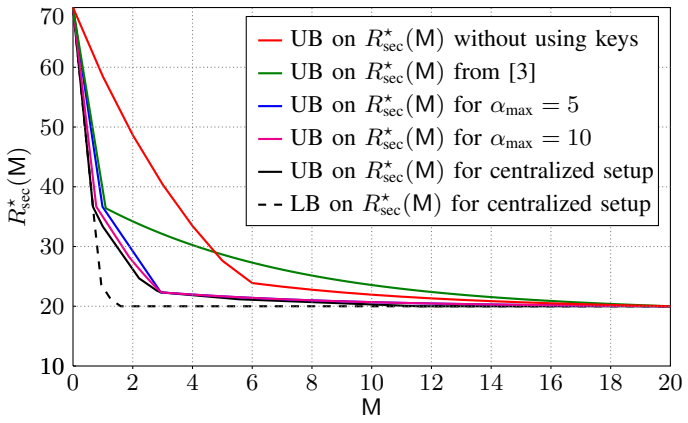

Fig. 3. Upper and lower bounds on $R^\star_{\text{sec}}(\mathsf{M})$ for $\delta_w = 0.7$, $\delta_s = 0.3$, $\delta_z = 0.8$, $\mathsf{K}_w = 5$, $\mathsf{K}_s = 10$, and $\mathsf{D} = 20$.

BC code to communicate to the strong receivers, which do not have cache memories. The red curve in Figure 3 depicts the upper bound obtained when in our joint coding scheme the usage of secret keys is replaced by random binning.

In the sequel, we describe the two schemes achieving the secrecy rate-memory pairs in (11) and (12). In the first scheme, only keys are prefetched in the cache memories. In the second scheme, prefetching includes keys and data.

IV. Scheme with Only Keys in the Cache

In this scheme, weak users prefetch only key bits. Assume

$$\alpha_{\max} \geq \frac{\mathsf{K}_w(1-\delta_z)}{1-\delta_w}.$$

Fix a file size F and a small positive $\epsilon > 0$, and define

$$q_\epsilon := 1 - \sqrt[\mathsf{K}_w]{1 - \frac{\mathsf{K}_w(1-\delta_z)}{\alpha_{\max}(1-\delta_w)}} + \epsilon. \tag{15}$$

Fix also a small positive $\epsilon' > 0$ so that

$$\alpha_{\max}\left[1-(1-q_\epsilon)^{\mathsf{K}_w}\right] > \frac{\mathsf{K}_w(1-\delta_z)}{1-\delta_w} + \epsilon'. \tag{16}$$

This is possible because q_ϵ is increasing in ϵ and because $\alpha_{\max}[1-(1-q_0)^{\mathsf{K}_w}] = \frac{\mathsf{K}_w(1-\delta_z)}{1-\delta_w}$.

A. Decentralized Prefetching Phase

For each $b \in \{1,\dots,\mathsf{F}_{\text{key}}\}$, every weak receiver $i \in \mathcal{K}_w$ prefetches the b-th bit of the stream S, i.e., S_b, with probability q_ϵ, independently of all other bits and of all other receivers. For any given subset of receivers $\mathcal{G} \subseteq \mathcal{K}_w$, define the key bits stored exclusively at receivers in $\mathcal{G}$:

$$\mathsf{S}_{\mathcal{G}} := \left\{\mathsf{S}_b \colon \mathsf{S}_b \text{ cached exclusively at receivers in } \mathcal{G}\right\}.$$

No data is cached in this scheme. The cache content at a given weak receiver i can then be written as:

$$V_i = \left\{\mathsf{S}_{\mathcal{G}} \colon \ \mathcal{G} \text{ so that } i \in \mathcal{G}\right\}. \tag{17}$$

By the weak law of large numbers, for sufficiently large key sizes F_{key} and all $\ell \in \{1,\dots,\mathsf{K}_w\}$, approximately the same number of key bits is exclusively cached at any subset of ℓ weak receivers. Let $\mathsf{F}^{(\ell)}_{\text{key}}$ be the expected number of key bits commonly and exclusively cached at any given size-ℓ subset of receivers and $\mathsf{F}^{(0)}_{\text{key}}$ be the expected number of key bits cached at no receiver. For $\ell \in \{1,\dots,\mathsf{K}_w\}$, $\mathsf{F}^{(\ell)}_{\text{key}}$ is given by

$$\mathsf{F}^{(\ell)}_{\text{key}} := \gamma^{(\ell)}_\epsilon \alpha_{\max} \mathsf{F}, \tag{18}$$

where

$$\gamma^{(\ell)}_\epsilon := \ q^\ell_\epsilon \left(1-q_\epsilon\right)^{\mathsf{K}_w-\ell}. \tag{19}$$

The number of effectively cached key bits is approximately equal to its expectation with high probability. The described cache placement will thus be admissible with high probability and when F is sufficiently large, if the cache size satisfies

$$\begin{aligned}
\mathsf{MF} &\geq \sum_{\ell=1}^{\mathsf{K}_w} \binom{\mathsf{K}_w-1}{\ell-1} \gamma^{(\ell)}_\epsilon \alpha_{\max}\mathsf{F} + \epsilon \\
&= \alpha_{\max}\mathsf{F} \sum_{\ell=1}^{\mathsf{K}_w} \binom{\mathsf{K}_w-1}{\ell-1} q^\ell_\epsilon \left(1-q_\epsilon\right)^{\mathsf{K}_w-\ell} + \epsilon \\
&= \alpha_{\max}\mathsf{F} \sum_{i=0}^{\mathsf{K}_w-1} \binom{\mathsf{K}_w-1}{i} q^{i+1}_\epsilon \left(1-q_\epsilon\right)^{\mathsf{K}_w-1-i} + \epsilon \\
&= \alpha_{\max} q_\epsilon \mathsf{F} + \epsilon.
\end{aligned} \tag{20}$$

B. Centralized Delivery Phase

Delivery communication is split into two subphases. Subphase 1 consists of

$$n_1 = \frac{\mathsf{K}_w\mathsf{F}}{1-\delta_w} + \frac{\epsilon'}{1-\delta_z}\mathsf{F} \tag{21}$$

channel uses and is dedicated to send information to the K_w weak receivers. Specifically, the transmitter extracts the K_w *independent* keys $\mathsf{S}^\star_1,\dots,\mathsf{S}^\star_{\mathsf{K}_w}$ from S, of length

$$n_{\text{key}} := \left(\frac{1-\delta_z}{1-\delta_w} + \frac{\epsilon'}{\mathsf{K}_w}\right)\mathsf{F}, \tag{22}$$

so that each S_i, $i \in \mathcal{K}_w$, is stored in Receiver i's cache memory. This is possible with high probability when F is sufficiently large because the expected number of totally cached bits is

$$\mathsf{F}_{\text{key}} - \mathsf{F}^{(0)}_{\text{key}} = \alpha_{\max}\left[1-(1-q_\epsilon)^{\mathsf{K}_w}\right]\mathsf{F}, \tag{23}$$

which by (16) and (22) exceeds $\mathsf{K}_w n_{\text{key}}$.

The transmitter then time-shares K_w wiretap codes with secret keys [10] to send message W_{d_1} to Receiver 1 using key $\mathsf{S}^\star_1$, message W_{d_2} to Receiver 2 using key $\mathsf{S}^\star_2$, message W_{d_3} to Receiver 3 using key $\mathsf{S}^\star_3$ and so on.

The probability of decoding error in this Subphase 1 tends to 0 as $\mathsf{F} \ \to \ \infty$, because the communication rate of each message satisfies $\mathsf{F}/(n_1/\mathsf{K}_w) < 1-\delta_w$. Transmission is secured from the eavesdropper because each transmitted

message is secured with a key of rate

$$\frac{n_{\text{key}}}{n_1/\mathsf{K}_w} > (1-\delta_z). \tag{24}$$

Subphase 2 consists of

$$n_2 = \frac{\mathsf{K}_s\mathsf{F}}{\delta_z-\delta_s} + \epsilon'\mathsf{F}. \tag{25}$$

channel uses and is dedicated only to the K_s strong receivers. Specifically, in Subphase 2, the transmitter uses a wiretap BC code [8] to send messages $W_{d_{\mathsf{K}_w+1}}, \dots, W_{d_{\mathsf{K}}}$ to the strong receivers. Communication in this subphase is thus secured from the eavesdropper. Moreover, the probability of decoding error tends to 0 as $\mathsf{F} \to \infty$ because the rate of communication satisfies $\mathsf{F}/(n_2/\mathsf{K}_s) < \delta_z - \delta_s$.

Combine all these observations and sum up $n = n_1 + n_2$. Then, letting $\epsilon, \epsilon' \to 0$, we can conclude achievability of the secrecy rate-memory pair in (11).

V. Scheme with Keys and Data in the Cache

In this scheme, weak users prefetch key and data bits. Fix file size F, and choose $\alpha \in [0, \alpha_{\max}]$ and $\epsilon > 0$. Define

$$p_\epsilon := p + \epsilon \tag{26}$$

where p is defined in (13). Moreover, for each $\ell \in \{0,1,\dots,\mathsf{K}_w\}$, define

$$\gamma^{(\ell)} := p^\ell (1-p)^{\mathsf{K}_w-\ell} \tag{27}$$

and

$$\gamma_\epsilon^{(\ell)} := p_\epsilon^\ell (1-p_\epsilon)^{\mathsf{K}_w-\ell}. \tag{28}$$

Choose a sufficiently small $\epsilon' > 0$.

A. Decentralized Prefetching Phase

For each $b \in \{1,\dots,\mathsf{F}_{\text{key}}\}$, every weak receiver $i \in \mathcal{K}_w$ prefetches the b-th bit of the stream S, i.e., S_b, with probability p_ϵ, independently of all other bits and of all other receivers. Similarly, for each $d \in \mathcal{D}$ and each $b \in \{1,\dots,\mathsf{F}\}$, every weak receiver $i \in \mathcal{K}_w$ prefetches the b-th bit of file W_d, i.e., $W_{d,b}$, with probability p, independently of all other bits and of all other receivers.

For any given subset of receivers $\mathcal{G} \subseteq \mathcal{K}_w$, define the key bits stored exclusively at receivers in $\mathcal{G}$:

$$\mathsf{S}_\mathcal{G} := \big\{\mathsf{S}_b \colon \mathsf{S}_b \text{ cached exclusively at receivers in } \mathcal{G}\big\},$$

and for each file W_d, the data bits stored exclusively at receivers in $\mathcal{G}$:

$$W_{d,\mathcal{G}} := \big\{W_{d,b} \colon W_{d,b} \text{ cached exclusively at receivers in } \mathcal{G}\big\}. \tag{29}$$

The cache content at weak receiver i can then be written as:

$$V_i = \bigcup_{\mathcal{G}\colon i\in\mathcal{G}} \Big\{\mathsf{S}_\mathcal{G}, W_{1,\mathcal{G}}, \dots, W_{\mathsf{D},\mathcal{G}}\Big\}, \quad i \in \{1,\dots,\mathsf{K}_w\}. \tag{30}$$

Notice that the expected number of key bits commonly and exclusively prefetched at a given size-ℓ subset of receivers is:

$$\mathsf{F}_{\text{key}}^{(\ell)} := \gamma_\epsilon^{(\ell)}\alpha\mathsf{F}, \qquad \ell \in \{0,1,\dots,\mathsf{K}_w\}. \tag{31}$$

Similarly, the expected number of bits of any file W_d prefetched at a given size-ℓ subset of receivers is

$$\mathsf{F}^{(\ell)} := \gamma^{(\ell)}\mathsf{F}, \qquad \ell \in \{0,1,\dots,\mathsf{K}_w\}. \tag{32}$$

Notice that by the weak law of large numbers, the number of effectively prefetched (data and key) bits is approximately equal to its expectation with high probability.

The described prefetching is thus admissible with high probability, if the cache size

$$\begin{aligned}
\mathsf{MF} &\geq \mathsf{D}\sum_{\ell=1}^{\mathsf{K}_w}\binom{\mathsf{K}_w-1}{\ell-1}\gamma^{(\ell)}\mathsf{F} + \sum_{\ell=1}^{\mathsf{K}_w}\binom{\mathsf{K}_w-1}{\ell-1}\gamma_\epsilon^{(\ell)}\alpha\mathsf{F} + \epsilon \\
&= \mathsf{DF}\sum_{\ell=1}^{\mathsf{K}_w}\binom{\mathsf{K}_w-1}{\ell-1}p^\ell(1-p)^{\mathsf{K}_w-\ell} \\
&\quad + \alpha\mathsf{F}\sum_{\ell=1}^{\mathsf{K}_w}\binom{\mathsf{K}_w-1}{\ell-1}p_\epsilon^\ell(1-p_\epsilon)^{\mathsf{K}_w-\ell} + \epsilon \\
&= p\mathsf{DF} + p_\epsilon\alpha\mathsf{F} + \epsilon.
\end{aligned} \tag{33}$$

B. Centralized Delivery Phase

The delivery phase is divided into $\mathsf{K}_w + 1$ subphases. Subphase ℓ, for $\ell \in \{1,\dots,\mathsf{K}_w\}$, consists of

$$n_\ell := \frac{\binom{\mathsf{K}_w}{\ell}\mathsf{F}^{(\ell-1)}}{1-\delta_w} + \epsilon'\mathsf{F}, \tag{34}$$

channel uses and is further divided into $\binom{\mathsf{K}_w}{\ell}$ equally-long periods. Communication in each of these periods is intended to a different subset of ℓ weak receivers and to all the K_s strong receivers. The last Subphase $\mathsf{K}_w + 1$ is of length

$$n_{\mathsf{K}_w+1} := \frac{\mathsf{K}_s\Big(\mathsf{F}^{(0)} + \sum_{\ell=1}^{\mathsf{K}_w}\binom{\mathsf{K}_w}{\ell}\mathsf{F}^{(\ell,2)}\Big)}{\delta_z-\delta_s} + \epsilon'\mathsf{F}, \tag{35}$$

where for each $\ell \in \{1,\dots,\mathsf{K}_w\}$, we define

$$\mathsf{F}^{(\ell,2)} := \mathsf{F}^\ell - \mathsf{F}^{(\ell,1)}, \tag{36}$$

$$\mathsf{F}^{(\ell,1)} := \min\left\{\mathsf{F}^{(\ell)}, \mathsf{F}^{(\ell-1)}\frac{\delta_w-\delta_s}{\mathsf{K}_s(1-\delta_w)}\right\} - \epsilon'\mathsf{F}. \tag{37}$$

This last subphase is intended only to the strong receivers.

Before transmission starts, the transmitter divides each of the strong receivers' submessages in (29) into 2 more parts:

$$W_{d_j,\mathcal{G}} = \left(W_{d_j,\mathcal{G}}^{(1)}, W_{d_j,\mathcal{G}}^{(2)}\right), \quad \forall j \in \mathcal{K}_s,\ \forall \mathcal{G} \subseteq \mathcal{K}_w, \tag{38}$$

of lengths $\mathsf{F}^{(|\mathcal{G}|,1)}$ and $\mathsf{F}^{(|\mathcal{G}|,2)}$.

The transmitter further extracts from the key stream S the 2^{K_w} *independent* keys $\{\mathsf{S}_\mathcal{G} \colon \forall \mathcal{G} \subseteq \mathsf{K}_w\}$, in a way that for each $\mathcal{G}$ the key $\mathsf{S}_\mathcal{G}$ is of length

$$n_{\text{key},|\mathcal{G}|} := \frac{n_{|\mathcal{G}|}}{\binom{\mathsf{K}_w}{|\mathcal{G}|}} \cdot (1-\delta_z), \tag{39}$$

and receivers in $\mathcal{G}$ can construct it from their cache contents. The described preparations are feasible for sufficiently small ϵ' by the weak law of large numbers, because for $\epsilon = \epsilon' = 0$ we have $n_{\text{key},|\mathcal{G}|} = \mathsf{F}_{\text{key}}^{(|\mathcal{G}|)}$ and because the number of commonly

(not necessarily exclusively) stored key bits at any set of $\mathcal{G}$ receivers is an increasing function of ϵ.

Consider now transmission in Subphase $\ell \in \mathcal{K}_w$ and let $\mathcal{G}_1^{(\ell)}, \ldots, \mathcal{G}_{\binom{\mathsf{K}_w}{\ell}}^{(\ell)}$ denote the $\binom{\mathsf{K}_w}{\ell}$ subsets of $\mathcal{K}_w$ of size ℓ. For each period $t \in \{1, \ldots, \binom{\mathsf{K}_w}{\ell}\}$ of Subphase ℓ, the transmitter creates a piggyback codebook [11]

$$\mathcal{C}_t := \Big\{ \mathbf{x}_t(w_r, w_c) \colon \quad w_r \in \big\{1, \ldots, 2^{\mathsf{F}^{(\ell-1)}}\big\}, \\ w_c \in \big\{1, \ldots, 2^{\mathsf{F}^{(\ell,1)}}\big\} \Big\}, \tag{40}$$

with entries drawn i.i.d. according to a Bernoulli-$1/2$ distribution. The codewords of such a codebook are arranged in an array with rows encoding the message w_r and columns encoding the message w_c, see Figure 4.

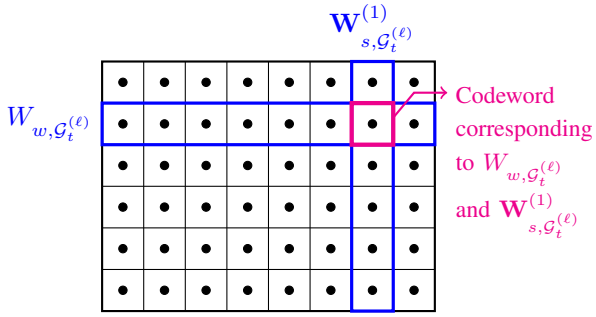

Fig. 4. Structure of piggyback codebook with codewords arranged in an array. Here the rows encode $W_{w,\mathcal{G}_t^{(\ell)}}$ and the columns $\mathbf{W}^{(1)}_{s,\mathcal{G}_t^{(\ell)}}$.

In period t of subphase ℓ, the transmitter conveys[1]

$$W_{w,\mathcal{G}_t^{(\ell)}} := \Big(\Big(\bigoplus_{i \in \mathcal{G}_t^{(\ell)}} W_{d_i,\mathcal{G}_t^{(\ell)} \setminus \{i\}} \Big) + \mathsf{S}_{\mathcal{G}_t^{(\ell)}} \Big) \mod 2^{\mathsf{F}^{(\ell-1)}} \tag{41}$$

to all weak receivers in $\mathcal{G}_t^{(\ell)}$, and the message

$$\mathbf{W}^{(1)}_{s,\mathcal{G}_t^{(\ell)}} := \Big(W^{(1)}_{d_{\mathsf{K}_w+1},\mathcal{G}_t^{(\ell)}}, \ldots, W^{(1)}_{d_{\mathsf{K}},\mathcal{G}_t^{(\ell)}} \Big) \tag{42}$$

to all K_s strong receivers. To this end, it uses the piggyback codebook in Figure 4 and sends the codeword $\mathbf{x}_t\big(W_{w,\mathcal{G}_t^{(\ell)}}, \mathbf{W}^{(1)}_{s,\mathcal{G}_t^{(\ell)}}\big)$ in row $W_{w,\mathcal{G}_t^{(\ell)}}$ and column $\mathbf{W}^{(1)}_{s,\mathcal{G}_t^{(\ell)}}$.

Each weak receiver $i \in \mathcal{G}_t^{(\ell)}$ has stored $\mathbf{W}^{(1)}_{s,\mathcal{G}_t^{(\ell)}}$ in its cache memory and can decode based on the restricted codebook $\mathcal{C}_{t,\mathcal{G}_t^{(\ell)}}\big(\mathbf{W}^{(1)}_{s,\mathcal{G}_t^{(\ell)}}\big)$ consisting only of the codewords in the column indicated by $\mathbf{W}^{(1)}_{s,\mathcal{G}_t^{(\ell)}}$:

$$\mathcal{C}_{t,\mathcal{G}_t^{(\ell)}}\big(\mathbf{W}^{(1)}_{s,\mathcal{G}_t^{(\ell)}}\big) := \Big\{ \mathbf{x}_1\big(w_r, \mathbf{W}^{(1)}_{s,\mathcal{G}_t^{(\ell)}}\big) \colon \\ w_r \in \Big\{1, \ldots, 2^{\mathsf{F}^{(\ell-1)}}\Big\} \Big\}.$$

Its decoding performance is thus the same as if this message $\mathbf{W}^{(1)}_{s,\mathcal{G}_t^{(\ell)}}$ had not been sent at all, and it decodes correctly with probability tending to 1 as $\mathsf{F} \to \infty$, because

$$\frac{\mathsf{F}^{(\ell-1)}}{n_\ell / \binom{\mathsf{K}_w}{\ell}} < 1 - \delta_w. \tag{43}$$

Strong receivers have no cache memories and decode both messages $W_{w,\mathcal{G}_t^{(\ell)}}$ and $\mathbf{W}^{(1)}_{s,\mathcal{G}_t^{(\ell)}}$ based on the entire codebook $\mathcal{C}_t$. This decoding is correct with probability tending to 1 as $\mathsf{F} \to \infty$, because

$$\frac{\mathsf{F}^{(\ell-1)} + \mathsf{K}_s \mathsf{F}^{(\ell,1)}}{n_\ell / \binom{\mathsf{K}_w}{\ell}} \le \frac{\mathsf{F}^{(\ell-1)}}{n_\ell / \binom{\mathsf{K}_w}{\ell}} \cdot \frac{1-\delta_s}{1-\delta_w} < 1 - \delta_s, \tag{44}$$

where the inequalities hold by (37) and (43). Communication in this subphase is secure because the modulo-operation in (41) acts as a random binning for both messages.

In subphase $\mathsf{K}_w + 1$, the transmitter uses a wiretap BC code to send the missing parts of their messages to the strong receivers. The analysis for this phase is standard and omitted.

Combining all these considerations and letting $\epsilon, \epsilon' \to 0$, establishes achievability of the tradeoff in (12).

VI. Summary

We have derived an upper bound on the decentralized secrecy rate-memory tradeoff of a K-receiver wiretap BC under a joint secrecy constraint. In this setup, K_w receivers are weak and have cache memories and K_s receivers are strong and have no cache memories. We propose a coding scheme where prefetching of key bits and data bits is decentralized. For small cache sizes, the performance of our scheme improves with increasing length of the random key sequence stored at the transmitter. For moderate and large cache sizes, it performs close to the fundamental limit, regardless of the size of this key sequence.

[1] Here, $\bigoplus$ denotes a bitwise XOR operation and mod the modulo operator.

References

[1] M. A. Maddah-Ali and U. Niesen, "Fundamental limits of caching," *IEEE Trans. Inf. Theory*, vol. 60, no. 5, pp. 2856–2867, May 2014.

[2] M. A. Maddah-Ali and U. Niesen, "Decentralized coded caching attains order-optimal memory-rate tradeoff," *IEEE/ACM Trans. on Networking*, vol. 23, no. 4, pp. 1029–1040, Aug. 2015.

[3] A. Sengupta, R. Tandon, and T. C. Clancy, "Fundamental limits of caching with secure delivery," *IEEE Trans. on Inf. Forensics and Sec.*, vol. 10, no. 2, pp. 355–370, Feb. 2015.

[4] S. Saeedi Bidokhti, M. Wigger and A. Yener, "Benefits of cache assignment on degraded broadcast channels." ArXiv: 1702.08044.

[5] S. Kamel, M. Sarkiss, M. Wigger, and G. Rekaya-Ben Othman "Secrecy capacity-memory tradeoff of erasure broadcast channels." ArXiv:1801.00606, Jan. 2018.

[6] S. Kamel, M. Sarkiss and M. Wigger, "Secure joint cache-channel coding over erasure broadcast channels," *Proc. of IEEE WCNC*, San Francisco, CA, Mar. 2017.

[7] S. Kamel, M. Sarkiss and M. Wigger, "Decentralized joint cache-channel coding over erasure broadcast channels," *Proc. of IEEE MENACOMM*, Jounieh, Lebanon, Apr. 2018.

[8] E. Ekrem and S. Ulukus, "Multi-receiver wiretap channel with public and confidential messages," *IEEE Trans. Inf. Theory*, vol. 59, no. 4, pp. 2165–2177, Apr. 2013.

[9] M. M. Amiri and D. Gündüz, "Cache-aided content delivery over erasure broadcast channels," *IEEE Trans. on Comm.*, vol. 66, no. 1, pp. 370–381, Jan. 2018.

[10] H. Yamamoto, "Rate-distortion theory for the Shannon cipher system," *IEEE Trans. Inf. Theory*, vol. 43, no. 3, pp. 827–835, May 1997.

[11] S. Saeedi Bidokhti, R. Timo, and M. Wigger, "Noisy broadcast networks with receiver caching." to appear *IEEE Trans. Inf. Theory*, 2018.