Coded Caching in Land Mobile Satellite Systems

Hui Zhao
Communication Systems Dept., EURECOM
Sophia Antipolis, France
Email: hui.zhao@eurecom.fr

Antonio Bazco-Nogueras
IMDEA Networks Institute
Madrid, Spain
Email: antonio.bazco@imdea.org

Petros Elia
Communication Systems Dept., EURECOM
Sophia Antipolis, France
Email: petros.elia@eurecom.fr

Abstract—We investigate the performance of coded caching in land mobile-satellite (LMS) systems, where a satellite station with full access to a content library serves K cache-aided land users. The promising gains that coded caching provides for the error-free shared-link Broadcast Channel are known to suffer in some low signal-to-noise ratio (SNR) scenarios. For this reason, we analyze to what extent the coded caching gains are preserved in LMS systems, which are governed by low-to-moderate SNR due to the long propagation distances. We model the satellite-terrestrial channels through the widely adopted Shadowed-Rician fading, and we show that the coded caching gains are partially preserved even in the low-SNR limit due to the existence of line-of-sight (LOS) components, which allows us to double the goodput of LMS communication at low SNR. These results illustrate the potential of coded caching in LMS systems and motivate further research to design adapted and practical coded caching schemes for LMS systems.

I. Introduction

One of the main challenges in current and future mobile networks is dealing with the ever-increasing demand of video content. Coded caching has arisen as a promising solution to handle this problem, since it has shown that caching modest amounts of information at the receivers can provide remarkable improvements in content delivery time. Coded caching capitalizes on the fact that this traffic considerably fluctuates throughout the day, and do so by storing (or caching) content at the end users beforehand during off-peak hours.

Coded caching was first proposed by Maddah-Ali and Niesen in [1], where they presented a scenario in which a transmitter, which has access to a library of N files, serves K users in an *error-free shared-link* Broadcast Channel (BC) where each user is endowed with a local memory (or cache) with capacity to store a fraction $\gamma \in [0, 1]$ of the library content. In a first phase, these caches are filled during off-peak hours and, in a second phase, each user requests a file from the library. In order to deliver the content, the transmitter generates a common message by applying XOR operations to several messages, each one intended by a different user. Then, it sends this common message over the physical-layer channel to serve $K\gamma + 1$ users, thereby enhancing the content delivery rate by a factor of $K\gamma + 1$ compared to *uncoded caching with simple time-division multiplexing (TDM)*. It is worth noting that in such uncoded caching scheme every user caches a fraction γ of *each file* in the library. During the delivery phase, the transmitter sends the remaining content of the requested files one-by-one [1].

We refer to the coded caching scheme presented in [1] as the *MN scheme* and to the boost factor $K\gamma + 1$ as the *theoretical coded caching gain*. Motivated by these promising gains, the implementation of coded caching in realistic wireless networks has attracted great interest, including works that analyze e.g. dense urban cells [2], cell-free networks [3], MIMO systems [4], [5], among others [6], [7].

Another interesting use case is land mobile-satellite (LMS) systems [8], whose integration in wireless networks is very relevant due to their large service coverage, wide bandwidth range, and high throughput. Indeed, the use of this type of communications is rapidly growing in recent years [9]. LMS systems can provide persistent and low-cost services in the case where the terrestrial network connection is very weak or even unavailable, such as in the case of network disruptions, rural areas, civil aviation, or cruise ships. Furthermore, these systems might serve an enormous amount of users once they handle a bigger portion of the traffic, mainly due to their large coverage area. These characteristics motivate us to analyze the potential of coded caching in LMS systems. Some works that jointly consider satellite and cache-aided terrestrial networks can be found in [10]–[15], which have mostly focused on uncoded caching. However, in practical scenarios we usually have that the user cache size is much smaller than the total size of the library ($\gamma \ll 1$), which makes the gain of uncoded caching over the approach without caching insignificant [1]. In contrast, coded caching provides a multiplicative multicasting gain of $K\gamma + 1$ even though the caches cannot cooperate [1].

Worst-User Bottleneck in Coded Caching:

There exists an important aspect that impacts the performance of coded caching in wireless scenarios, which is referred to as the *worst-user bottleneck*, and that results from the fact that the achievable rate in multicast transmissions is constrained by the user with the worst channel state among the $K\gamma + 1$ simultaneously served users [16]. *This bottleneck is unfortunately exacerbated as the signal-to-noise ratio (SNR) decreases.* Indeed, the gain of XOR-based coded multicasting schemes (e.g., MN) has been shown to entirely vanish in the low-SNR limit over symmetric Rayleigh fading channels [4], [17], [18], where the content delivery rate of coded caching deteriorates down to the rate of uncoded caching with TDM.

The worst-user bottleneck may be particularly important in LMS systems because such systems tend to operate in the low-SNR scenario due to the long distances between the satellite and the terrestrial users. For example, the received

Fig. 1: Schematic representation of the considered setting.

SNR at a terrestrial user is typically *below* 10 dB in home TV broadcasts from geosynchronous (GEO) satellites with 36 MHz of bandwidth and several hundred of watts of power [19]. Yet, while this problem may be tragic for the implementation of coded caching in a variety of settings such as Rayleigh fading, we need to explore if the same conclusion holds in LMS systems, where the channel statistics are totally different from Rayleigh fading [20] due to the existence of LOS components. In this paper, we consider Rician-Shadowed fading [20] to model the satellite-terrestrial channel. Our main contributions can be summarized as follows:

- We analyze the goodput (which will be defined in (6)) of the seminal MN scheme for first time in LMS systems. We derive the analytical expression of the maximum achievable goodput and validate these results with numerical evaluations for several practical fading configurations.
- Our analysis shows that the goodput gain of coded caching in practical low-SNR scenarios is maintained in LMS settings when the LOS components are not heavily obstructed, even in the low-SNR limit. This result allows us to *double* the achievable goodput at very low SNR just considering the seminal coded caching scheme, which is tailored to the high-SNR regime. This deviates from the results of previous works [4], [18], which proved that this gain disappears at low SNR for Rayleigh fading, and shows that coded caching has the ability to operate in low-SNR-governed (i.e., below 10 dB) LMS systems.

II. SYSTEM MODEL

We consider a scenario in which a GEO satellite transmitter serves K terrestrial users. Each user requests a file from a content library composed of N files, each of size F bits. We denote the n-th file of the library by W_n, such that the library is given by $\{W_n\}_{n=1}^N$. The broadcasting GEO satellite has full access to the library,[1] and each user has a local memory (or cache) of size MF bits which is filled with content from the library during off-peak hours. The ratio between the local memory size and the total size of the library is denoted by $\gamma \triangleq \frac{M}{N} \in [0, 1]$ and is referred to as the *normalized cache size* of the local memory. This scenario is depicted in Fig. 1. We consider the MN coded caching scheme from [1], which we will

[1]This access can be either because all the content is stored at the satellite or because this satellite is connected to a terrestrial satellite gateway via a feeder link which can often sustain extremely high rates. This satellite gateway has a wired connection to the core network and thus access to the library.

TABLE I: Parameters in four typical fading scenarios

Scenarios	m	b	Ω
Frequent Heavy Shadowing (FHS)	1	0.063	8.97×10^{-4}
Overall Results (OR)	5	0.251	0.278
Average Shadowing (AS)	10	0.126	0.835
Infrequent Light Shadowing (ILS)	20	0.158	1.29

describe in Section II-B, such that as many as $K\gamma + 1$ users (marked with the same color in Fig. 1) are simultaneously served via a satellite BC. In the following, we use $|\cdot|$ to denote the cardinality of a set or the magnitude of a complex number. We define $[n] \triangleq \{1, 2, \dots, n\}$ for any positive integer n, whereas $[a, b]$ denotes the closed interval of real numbers. We also use $\exp(\cdot) \triangleq e^{(\cdot)}$, where e is Euler's number.

A. Channel Model

In order to accurately describe the fluctuation of the signal envelope, we consider the widely adopted Rician-Shadowed fading [20] to model the satellite-terrestrial channel, which can be nicely calibrated to capture both fixed and mobile land terminals, and which can be applied for all types of orbits and for a variety of frequency bands including S-band, L-band, Ku-band, and Ka-band [13], [21], [22]. We recall that the Rician-Shadowed fading model is calibrated by the following set of parameters: the average power of the scatter component, denoted by $2b$, the average power of the LOS component, denoted by Ω, the amplitude of the random shadowing LOS component, denoted by Z, and $m \triangleq \frac{(\mathbb{E}\{Z^2\})^2}{\mathrm{Var}\{Z^2\}}$, which reflects the (average) obstruction of the LOS component (i.e., the blockage of the LOS by buildings, trees, hills, etc.) and where $m = 0$ stands for complete obstruction whereas $m \to \infty$ corresponds to no obstruction [20]. To facilitate the analysis, we assume the same channel statistics (i.e., same m, b and Ω) among the K users, which is reasonable when the users are uniformly distributed within a disk [22]. We consider that the users are located within a disk of radius equal to several kilometers [2]. As the radius is negligible to the height of the GEO satellite, we assume that all the users have the same path-loss. Therefore, we consider *statistically symmetric users*.

GEO satellites are static with respect to an observer from Earth, which makes the Doppler spread negligible for static terrestrial users [21]. Thus, we assume that the coherence time is large, and we do not consider an ergodic channel. Instead, we assume that the channel experiences *quasi-static fading*, which generally comes about in the presence of longer coherence periods and shorter latency constraints, and models low-mobility scenarios which nicely capture coded-caching use-cases where slow users are consuming video content.

While the results hold for any generic Rician-Shadowed fading, in the numerical results we consider four fading scenarios that have been used for performance analysis of LMS systems, and which are obtained from [20, Table III]. Since we will refer to these cases thorough the document, we present them in Table I.

The probability density function (PDF) of the instantaneous channel power gain $|h_k|^2$ at user $k \in [K]$ over Rician-

Shadowed fading channels is given by (cf. [20])

$$f_{|h_k|^2}(x) = \alpha e^{-\beta x} \, {}_1F_1(m; 1; \delta x), \quad x \geq 0, \tag{1}$$

where $\alpha = \left(\frac{2bm}{2bm+\Omega}\right)^m$, $\beta = \frac{1}{2b}$, $\delta = \frac{\Omega}{2b(2bm+\Omega)}$, and where ${}_1F_1(\cdot; \cdot; \cdot)$ denotes the confluent hypergeometric function of the first kind [23]. As a small difference in m reflects a similar (average) obstruction of the LOS component, we consider that m is a positive integer for mathematical tractability, which is a simplifying assumption that is widely adopted in many existing works [13], [21]. Let us denote the normalized transmit power of the satellite as P_t, and the power of the Additive White Gaussian Noise as N_0. Therefore, the instantaneous SNR at user k is $\mathrm{SNR}_k \triangleq \frac{P_t |h_k|^2}{N_0} = \rho |h_k|^2$, where $\rho \triangleq \frac{P_t}{N_0}$.

By considering [24, Eq. (24)], we simplify the confluent hypergeometric function in (1) for any positive integer m as

$$ {}_1F_1(m; 1; \delta x) = e^{\delta x} \sum_{i=0}^{m-1} \binom{m-1}{i} \frac{(\delta x)^i}{i!}. \tag{2}$$

Hence, upon defining $\zeta(i) \triangleq \binom{m-1}{i} \frac{\delta^i}{i!}$, the PDF of $\mathrm{SNR}_k = \rho |h_k|^2$ can be simplified for any integer $m \geq 1$ as (cf. [21])

$$f_{\mathrm{SNR}_k}(x) = \alpha \sum_{i=0}^{m-1} \frac{\zeta(i)}{\rho^{i+1}} x^i e^{-\frac{\beta-\delta}{\rho}x}, \quad x \geq 0. \tag{3}$$

Let $\Gamma(\cdot, \cdot)$ denote the upper incomplete Gamma function [23]. By integrating the PDF of SNR_k in (3), we can easily obtain the corresponding cumulative distribution function (CDF) as

$$F_{\mathrm{SNR}_k}(x) = 1 - \alpha \sum_{i=0}^{m-1} \frac{\zeta(i)}{(\beta-\delta)^{i+1}} \Gamma\left(i+1, \frac{\beta-\delta}{\rho}x\right). \tag{4}$$

B. MN Coded Caching Scheme

We recall that the coded caching framework in [1] consists of two distinct phases: the placement phase, in which the local cache of the users is filled, and which is carried out during off-peak hours before the actual demand of each user is revealed, and the delivery phase, in which each user requests a file from the library, and the transmitter serves these users while aiming at minimizing the total delivery time required to serve all users.

1) MN Placement phase: During the placement phase, each file W_n is partitioned into $\binom{K}{K\gamma}$ non-overlapping and equal-sized segments (subfiles), such that $W_n \rightarrow \{W_n^{\mathcal{T}} : \mathcal{T} \subseteq [K], |\mathcal{T}| = K\gamma\}$. User k stores all the segments $W_n^{\mathcal{T}}$ such that $k \in \mathcal{T}$, for any $n \in [N]$. The content cached at user k, denoted by $\mathcal{Z}_k$, is hence given by $\mathcal{Z}_k = \{W_n^{\mathcal{T}} : \mathcal{T} \subseteq [K], |\mathcal{T}| = K\gamma, \mathcal{T} \ni k, \forall n \in [N]\}$. It follows that the total content cached at each user amounts to MF bits, which satisfies the local cache size constraint.

2) MN Delivery phase: At the beginning of this phase, each user requests a different file W_{d_k} from the library, where $d_k \in [N]$ denotes the index of the file demanded by user $k \in [K]$. The transmission is divided in different transmission stages. At each transmission stage, the transmitter simultaneously serves a unique set of $K\gamma + 1$ users. Since there are $\binom{K}{K\gamma+1}$ different subsets of $K\gamma + 1$ users in $[K]$, the delivery phase consists of $\binom{K}{K\gamma+1}$ transmission stages, and at each stage the transmitter serves a different subset of users $\mathcal{G} \subseteq [K]$ of $|\mathcal{G}| = K\gamma + 1$ users. Specifically, for the transmission stage intended for a particular set of users $\mathcal{G}$, the transmitted signal is designed as $X_{\mathcal{G}} = \bigoplus_{k \in \mathcal{G}} W_{d_k}^{\mathcal{G} \setminus \{k\}}$, where $\bigoplus$ stands for the bit-wise XOR operator, and the superscript $\mathcal{G} \setminus \{k\}$ implies that the segment transmitted to user k is the one stored at all other users in $\mathcal{G}$. In the physical layer, $X_{\mathcal{G}}$ is mapped into a common multicast[2] message which is then sent to the users in $\mathcal{G}$ via a BC.

After successfully receiving $X_{\mathcal{G}}$, user k can "cache out" the undesired messages in $X_{\mathcal{G}}$ by using its locally-cached content, and thus it obtains the desired subfile $W_{d_k}^{\mathcal{G} \setminus \{k\}}$. This is possible because all subfiles $\{W_n^{\mathcal{G} \setminus \{k'\}}\}_{n=1}^N$, for $k' \in \mathcal{G}$ and $k' \neq k$, have been stored in the cache of user $k \in \mathcal{G}$. After $\binom{K}{K\gamma+1}$ transmission stages, all the users obtain their demanded files. Note that even when $K\gamma$ is not an integer, the MN scheme can still achieve the coded caching gain $K\gamma + 1$ by considering the memory-sharing strategy. We refer to [1] for more details about the MN scheme.

III. GOODPUT ANALYSIS

Let us consider that, in each transmission stage of the delivery phase of the MN scheme, the satellite adopts a constant transmission rate (C_{th}) — which will be optimized in Lemma 1 — to serve the corresponding served users in $\mathcal{G}$. In this way, users do not require to feedback their channel state information to the satellite, which can be energy demanding for the users due to the long distances to the satellite. In this case, the main metric of interest is the goodput [25], which is one of the standard metrics for the performance evaluation of LMS systems over fading channels [26]. We assume that the channel state remains constant during a transmission stage and may independently change to other values in the next stage.

We analyze in the following the goodput (i.e., the rate reliably delivered to the users [25]) of the MN scheme. For a given transmission to a user-set $\mathcal{G}$, the satellite delivers a common message to the users in $\mathcal{G}$ over the physical-layer channel. Then, the total instantaneous goodput is defined as

$$C_{\mathrm{out}} \triangleq |\mathcal{G}| C_{th} \mathbb{I}\{\ln(1 + \mathrm{SNR}_{\mathrm{MN}}) \geq C_{th}\} \quad \mathrm{nats/s/Hz}, \tag{5}$$

where $\mathbb{I}\{\cdot\}$ denotes the indicator function, which, for claim $\mathcal{A}$, takes the value $\mathbb{I}\{\mathcal{A}\} = 1$ if $\mathcal{A}$ is true and $\mathbb{I}\{\mathcal{A}\} = 0$ otherwise, and $\mathrm{SNR}_{\mathrm{MN}} \triangleq \min_{k \in \mathcal{G}}\{\mathrm{SNR}_k\}$, since we need to guarantee the successful decoding at the user with the weakest link strength. In (5), the existence of the factor $|\mathcal{G}| = K\gamma + 1$ is due to the fact that $|\mathcal{G}|$ users are served at a time. Consequently, we consider the worst case of outage inasmuch as an outage event happens as long as a single user in $\mathcal{G}$ fails to decode this common message. To analyze the coded caching gain later, and in a similar way as in [4], we take the expectation of C_{out} over channel states, which yields the (average) goodput as

$$\begin{aligned} \overline{C}_{\mathrm{out}} &= |\mathcal{G}| C_{th} \mathbb{E}\left\{\mathbb{I}\{\ln(1 + \mathrm{SNR}_{\mathrm{MN}}) \geq C_{th}\}\right\} \\ &= |\mathcal{G}| C_{th} \int_0^\infty \mathbb{I}\{\ln(1 + x) \geq C_{th}\} f_{\mathrm{SNR}_{\mathrm{MN}}}(x)\mathrm{d}x \\ &= |\mathcal{G}| C_{th} \mathbb{P}\big(\ln(1 + \mathrm{SNR}_{\mathrm{MN}}) \geq C_{th}\big), \end{aligned} \tag{6}$$

[2]There exists a variety of schemes that handle the worst-user bottleneck of multicast transmission or the case with bounded number of subfiles in coded caching. We restrict ourselves to the standard MN scheme to provide simple but meaningful insights on the benefit of coded caching.

where $f_{\mathrm{SNR_{MN}}}(x)$ denotes the PDF of $\mathrm{SNR_{MN}}$. Note that $\overline{C}_{\mathrm{out}}$ is just a statistical mean of C_{out} in (5), which should not be confused by the long-term goodput [4].

We present our first main result, which consists of closed-form expressions for $\overline{C}_{\mathrm{out}}$ and the rate that maximizes $\overline{C}_{\mathrm{out}}$.

Lemma 1. *The average goodput of the MN scheme is given by*

$$\overline{C}_{\mathrm{out}} = |\mathcal{G}|C_{th}\left[\sum_{i=0}^{m-1}\frac{\alpha\zeta(i)}{(\beta-\delta)^{i+1}}\Gamma\left(i+1,\frac{\beta-\delta}{\rho}(e^{C_{th}}-1)\right)\right]^{|\mathcal{G}|}.$$

Further, the optimal $\overline{C}_{\mathrm{out}}$ is $\overline{C}_{\mathrm{out}}^{\star} = \max_{C_{th}\in\mathcal{C}_{th}^*}\overline{C}_{\mathrm{out}}(C_{th})$, where $\mathcal{C}_{th}^*$ is the set of solutions to the equation

$$C_{th}^* = \frac{\sum_{i=0}^{m-1}\frac{\zeta(i)}{(\beta-\delta)^{i+1}}\Gamma\left(i+1,\frac{\beta-\delta}{\rho}(e^{C_{th}^*}-1)\right)}{|\mathcal{G}|e^{C_{th}^*}\sum_{i=0}^{m-1}\frac{\zeta(i)}{\rho^{i+1}}(e^{C_{th}^*}-1)^i\exp\left(-\frac{\beta-\delta}{\rho}(e^{C_{th}^*}-1)\right)}.$$

Proof. We can write (6) as

$$\overline{C}_{\mathrm{out}} = |\mathcal{G}|C_{th}\mathbb{P}\big(\min_{k\in\mathcal{G}}\{\mathrm{SNR}_k\}\geq e^{C_{th}}-1\big)$$
$$= |\mathcal{G}|C_{th}\left[1 - F_{\mathrm{SNR}_k}\left(e^{C_{th}}-1\right)\right]^{|\mathcal{G}|}. \quad (7)$$

Then, we can apply (4) into (7) to obtain the expression of $\overline{C}_{\mathrm{out}}$ in Lemma 1.

To maximize the goodput over C_{th}, we obtain the first derivative of $\overline{C}_{\mathrm{out}}$ with respect to C_{th}. For the sake of readability, let us define $\hat{F}_{\mathrm{SNR}_k}(\cdot) \triangleq 1 - F_{\mathrm{SNR}_k}(\cdot)$, and $\theta \triangleq e^{C_{th}}-1$. Then, we have that

$$\frac{\partial\overline{C}_{\mathrm{out}}}{\partial C_{th}} = |\mathcal{G}|\big(\hat{F}_{\mathrm{SNR}_k}(\theta)\big)^{|\mathcal{G}|-1}\Big(\hat{F}_{\mathrm{SNR}_k}(\theta)$$
$$- |\mathcal{G}|C_{th}f_{\mathrm{SNR}_k}(\theta)e^{C_{th}}\Big). \quad (8)$$

Since $\overline{C}_{\mathrm{out}}$ is a continuous function over $C_{th}\in(0,\infty)$, the maximum value will be in one of the extreme points. Thus, letting this first derivative be equal to 0 allows us to obtain C_{th}^* in Lemma 1 by using the PDF and CDF of SNR_k. $\square$

A. Low SNR Analysis

Next, we analyze the low-SNR regime by allowing the SNR parameter ρ to approach to zero. When $\rho\to 0$, we can consider the fact that $\ln(1+\rho x) = \rho x + o(x)$, where $o(x)$ denotes that $\lim_{\rho\to 0}o(x)/x = 0$, to write that the goodput in (6) satisfies

$$\overline{C}_{\mathrm{out}} = |\mathcal{G}|C_{th}\mathbb{P}(\mathrm{SNR_{MN}}\geq C_{th}+o(\mathrm{SNR_{MN}})) \quad (9)$$
$$\overset{\rho\to 0}{\approx} \widetilde{C}_{\mathrm{out}} \triangleq |\mathcal{G}|C_{th}\big[1 - F_{\mathrm{SNR}_k}(C_{th})\big]^{|\mathcal{G}|}, \quad (10)$$

where $\widetilde{C}_{\mathrm{out}}$ is obtained by omitting the term $o(\mathrm{SNR_{MN}})$ in (9), and where $F_{\mathrm{SNR}_k}(\cdot)$ has been given in (4).

In the following, we present a simplified expression for the values of C_{th} that maximize the goodput at low SNR.

Proposition 1. *The extreme points $C_{th}^* \in \mathcal{C}_{th}^*$ that make the derivative of $\widetilde{C}_{\mathrm{out}}$ equal to zero can be expressed by*

$$C_{th}^* = \frac{\epsilon^*}{|\mathcal{G}|(\beta-\delta)}\rho, \quad as \; \rho\to 0, \quad (11)$$

where $\epsilon^* \in [1,m]$ is a root of the following function:

$$\sum_{i=0}^{m-1}\frac{\zeta(i)}{(\beta-\delta)^{i+1}}\left[\Gamma\left(i+1,\frac{\epsilon^*}{|\mathcal{G}|}\right)-\exp\left(\frac{-\epsilon^*}{|\mathcal{G}|}\right)\frac{(\epsilon^*)^{i+1}}{|\mathcal{G}|^i}\right]=0. \quad (12)$$

TABLE II: Numerical results of ϵ^* in four fading scenarios

ϵ^* \\ Scenarios $|\mathcal{G}|$	FHS	OR	AS	ILS
1	1	1.3783	2.7121	3.4088
2	1	1.4342	3.4655	4.5529
3	1	1.4586	4.0111	5.4215
4	1	1.4723	4.4478	6.1422
5	1	1.4812	4.8149	6.7665
6	1	1.4874	5.1329	7.3215
7	1	1.4919	5.4139	7.8236
8	1	1.4954	5.6658	8.2836
9	1	1.4982	5.8943	8.7090
10	1	1.5005	6.1032	9.1054

Proof. The proof is relegated to the extended version of this work [27] due to space constraints. $\square$

Note that ϵ^* is independent of ρ, and it can have *at most* m different values [27]. Moreover, it follows that $\epsilon^* = 1$ for any scenario in which $m = 1$, e.g. FHS in Table I. Table II lists some numerical values of ϵ^* for the four typical fading scenarios in Table I and for realistic values of the theoretical coded caching gain $|\mathcal{G}|$. For all the cases in Table II, (12) has *only a single* (multiple) positive root. Therefore, if we use the values of ϵ^* in Table II, the derived C_{th}^* in (11) is exactly the global maximum point for $\widetilde{C}_{\mathrm{out}}$.

By using Proposition 1, we can obtain an expression for the optimal $\widetilde{C}_{\mathrm{out}}$, shown in Lemma 2, and which reveals that the optimal $\widetilde{C}_{\mathrm{out}}$ is actually a linear function of ρ.

Lemma 2. *The maximum goodput $\widetilde{C}_{\mathrm{out}}^*$ in the low-SNR regime can be expressed by*

$$\widetilde{C}_{\mathrm{out}}^* = \frac{\alpha^{|\mathcal{G}|}}{e^{\epsilon^\star}}\left(\frac{\epsilon^\star}{\beta-\delta}\right)^{|\mathcal{G}|+1}L_{m-1}^{|\mathcal{G}|}\left(\frac{-\epsilon^\star\delta}{|\mathcal{G}|(\beta-\delta)}\right)\rho, \quad (13)$$

where $L.(\cdot)$ denotes the Laguerre polynomial [23], and $\epsilon^\star \in [1,m]$ maximizes $\widetilde{C}_{\mathrm{out}}$ among the (at most m) roots of (12).

Proof. For the roots of (12), we have the following identity:

$$\sum_{i=0}^{m-1}\frac{\zeta(i)}{(\beta-\delta)^{i+1}}\Gamma\left(i+1,\frac{\epsilon^\star}{|\mathcal{G}|}\right)$$
$$= \frac{\epsilon^\star}{\beta-\delta}\exp\left(\frac{-\epsilon^\star}{|\mathcal{G}|}\right)\sum_{i=0}^{m-1}\frac{\zeta(i)}{(\beta-\delta)^i}\left(\frac{\epsilon^\star}{|\mathcal{G}|}\right)^i. \quad (14)$$

Applying the identity (14) in $F_{\mathrm{SNR}_k}(x)$ (defined in (4)) and incorporating this to (10) yields (13) after considering that $\sum_{i=0}^{m-1}\binom{m-1}{i}\frac{x^i}{i!} = L_{m-1}(-x)$. $\square$

Note that, if $m = 1$, we have that $\epsilon^\star = 1$ and therefore

$$\widetilde{C}_{\mathrm{out}}^* = \frac{\alpha^{|\mathcal{G}|}}{e(\beta-\delta)^{|\mathcal{G}|+1}}\rho. \quad (15)$$

Furthermore, when $m = 1$ and $\Omega = 0$, we have $\alpha = \beta = \frac{1}{2b}$ and $\delta = 0$, which means that Rician-Shadowed fading is reduced to Rayleigh fading. The optimal goodput becomes

$$\widetilde{C}_{\mathrm{out}}^* = \frac{\alpha^{|\mathcal{G}|}}{e\beta^{|\mathcal{G}|+1}}\rho = \frac{2b\rho}{e}, \quad (16)$$

where $2b\rho$ is the average SNR received at each user. Comparing (16) to the average capacity at low SNR (which converges

to the average SNR in the low-SNR limit) [18], we can see that there is a penalty of $1/e$ in the optimal goodput over Rayleigh fading channels. Moreover, $\widetilde{C}_{\text{out}}^{\star}$ in (16) is independent of $|\mathcal{G}|$, which leads to the same delivery performance as uncoded caching with TDM transmission at low SNR, and which is consistent with [18, Prop. 2].

B. Coded Caching Gain Analysis

To measure the benefit of coded caching over uncoded caching with TDM at finite SNR, we make use of the *effective coded caching gain* [2], [18], which is defined as the ratio of the maximum goodput achieved by the MN scheme over the maximum goodput achieved by uncoded caching with TDM.[3] Note that the latter corresponds to setting $|\mathcal{G}| = 1$ in the MN scheme [1]. Thus, the *effective coded caching gain* (Δ) is

$$\Delta \triangleq \frac{\max_{C_{th} \geq 0} \left\{ \overline{C}_{\text{out}}(|\mathcal{G}|, C_{th}) \right\}}{\max_{C_{th} \geq 0} \left\{ \overline{C}_{\text{out}}(|\mathcal{G}| = 1, C_{th}) \right\}}, \quad (17)$$

where $\overline{C}_{\text{out}}(|\mathcal{G}|, C_{th})$ is given in Lemma 1. The optimization of $\overline{C}_{\text{out}}$ over C_{th} can be solved by using Lemma 1. In the following, we focus on the behavior of this effective gain at low SNR, which is presented in Corollary 1.

Corollary 1. *In the low SNR limit, the effective coded caching gain of the MN scheme is given by*

$$\Delta \overset{\rho \to 0}{=} \left(\frac{\alpha}{\beta - \delta} \right)^{|\mathcal{G}|-1} \frac{(\epsilon^{\star})^{|\mathcal{G}|+1} e^{-\epsilon^{\star}} L_{m-1}^{|\mathcal{G}|} \left(-\frac{\epsilon^{\star} \delta}{|\mathcal{G}|(\beta - \delta)} \right)}{(\epsilon_1^{\star})^2 e^{-\epsilon_1^{\star}} L_{m-1} \left(-\frac{\epsilon_1^{\star} \delta}{\beta - \delta} \right)}, \quad (18)$$

where $\epsilon_1^{\star} \in [1, m]$ maximizes $\widetilde{C}_{\text{out}}^{\star}$ for $|\mathcal{G}| = 1$ among the (at most m) roots of (12).

Proof. The proof follows directly from considering Lemma 2 and the effective coded caching gain definition in (17). □

Note that for $|\mathcal{G}| \leq 10$ (which is a realistic setting [18]), and for the four typical fading scenarios listed in Table I, we can directly take the numerical results of $\epsilon^{\star}$ and $\epsilon_1^{\star}$ in Table II to evaluate the effective coded caching gain in the low-SNR limit, which will be also shown in Fig. 5.

For $m = 1$, such as in FHS, we have $\epsilon^{\star} = \epsilon_1^{\star} = 1$, and Δ in Corollary 1 can be simplified as $\Delta = \left(\frac{\alpha}{\beta - \delta} \right)^{|\mathcal{G}|-1} = 1$, which means that the MN scheme loses all gains in the low-SNR limit when $m = 1$ (i.e., with heavy LOS obstruction).

IV. NUMERICAL RESULTS

In this section, we perform some numerical analysis to validate the correctness of the derived expressions and present some interesting comparisons. Every simulated result is averaged over 10^5 channel realizations. We consider the four channel fading scenarios listed in Table I.

First, we evaluate the goodput as a function of C_{th} in Fig. 2. We can observe that the goodput has a single maximum, which

---

[3]Unlike some existing uncoded caching schemes where the cache hit/miss probability is typically investigated by considering the statistics of the users' requests, in the MN coded caching framework every user caches some content of each file in the library. During the delivery phase, the transmitter sends the remaining content of the requested files that is not cached in the users [1].


Fig. 2: Goodput versus C_{th} for $|\mathcal{G}| = 4$ and $\rho = -10$ dB.


Fig. 3: Optimal goodput versus ρ for $|\mathcal{G}| = 5$.

can be found by applying Lemma 1. It is also worth noting that Lemma 2 (based on the simplified expression of Proposition 1) provides a very accurate result at low SNR.

Fig. 3 plots the maximum goodput (optimized over C_{th}) versus ρ in the ILS, AS, OR and FHS scenarios. As mentioned in Lemma 2, the optimal goodput linearly grows with ρ at low SNR. For the same system settings as in Fig. 3, Fig. 4 shows the effective coded caching gain, and we can observe that the effective gain converges to a lower-bound as ρ decreases. However, the lower-bound depends on the fading scenario. Specifically, ILS has the largest effective coded caching gain in the low-SNR limit, followed by the figures for AS and OR, while the effective gain in FHS is the lowest (and equals 1). The reason is that ILS enjoys the lightest obstruction of the LOS component.

For the ILS scenario, the MN scheme almost doubles the goodput with respect to uncoded caching with TDM transmission even in the low-SNR limit. Note that this differs from the results on Rayleigh fading, where it has been shown that MN achieves no gain in the low-SNR limit [17], [18]. Moreover, except for the worst case (i.e., FHS), the MN scheme shows a significant boost effect over uncoded caching also for $\rho \leq 10$ dB, which, by considering that $\mathbb{E}\{|h_k|^2\} = 2b + \Omega$, corresponds to an average SNR $= \rho \mathbb{E}\{|h_k|^2\}$ of less than 10 dB in the AS case and than 1.3 dB in the FHS case

Fig. 5 shows the effective coded caching gain in the low-SNR limit versus the theoretical coded caching gain $|\mathcal{G}|$ (which is only achieved in the high-SNR limit). We can see how the effective gain increases for every scenario except for FHS.

Fig. 4: Effective coded caching gain versus ρ for $|\mathcal{G}| = 5$.

Fig. 5: Effective coded caching gain versus $|\mathcal{G}|$ as $\rho \to 0$.

Thus, in those cases, simultaneously serving more users from the single-antenna transmitter enhances the overall delivery performance. The steeper gain boost for the ILS and AS cases shows that the MN scheme (and coded caching) is advantageous mainly when the LOS component is not heavily obstructed.

V. Conclusion and Future Work

This work analyzes the goodput and the effective gain of coded caching in a LMS system where Rician-Shadowed fading is adopted to model the satellite-terrestrial channel. Our analysis shows that even the simplest coded caching scheme can double the goodput of LMS systems operating at very low SNR, despite the fact that such scheme is tailored for the high-SNR regime and suffers from the worst-user bottleneck. This implies that coded caching has the ability to operate in such (low-SNR governed) LMS systems. Satellite integration in beyond-5G networks, as well as satellite networks, are expected to be a relevant feature in the incoming communications developments. These results motivate further research on cache-aided LMS systems, e.g., by investigating the gain of more complex and adapted coded caching schemes to compensate the worst-user bottleneck for cases where the LOS component is heavily obstructed, but also by developing new schemes that adapt to the particular properties of LMS and other satellite systems.

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] H. Zhao, A. Bazco-Nogueras, and P. Elia, "Wireless coded caching with shared caches can overcome the near-far bottleneck," in *Proc. IEEE Int. Symp. Inf. Theory (ISIT)*, Jul. 2021, pp. 350–355.

[3] M. Bayat, R. K. Mungara, and G. Caire, "Achieving spatial scalability for coded caching via coded multipoint multicasting," *IEEE Trans. Wireless Commun.*, vol. 18, no. 1, pp. 227–240, Jan. 2019.

[4] K.-H. Ngo, S. Yang, and M. Kobayashi, "Scalable content delivery with coded caching in multi-antenna fading channels," *IEEE Trans. Wireless Commun.*, vol. 17, no. 1, pp. 548–562, Jan. 2018.

[5] H. Zhao, E. Lampiris, G. Caire, and P. Elia, "Multi-antenna coded caching analysis in finite SNR and finite subpacketization," in *Proc. Int. ITG Workshop on Smart Antennas (WSA)*, Nov. 2021, pp. 433–438.

[6] T. X. Vu, S. Chatzinotas, and B. Ottersten, "Edge-caching wireless networks: Performance analysis and optimization," *IEEE Trans. Wireless Commun.*, vol. 17, no. 4, pp. 2827–2839, Apr. 2018.

[7] S. Jin *et al.*, "A new order-optimal decentralized coded caching scheme with good performance in the finite file size regime," *IEEE Trans. Commun.*, vol. 67, no. 8, pp. 5297–5310, Aug. 2019.

[8] X. Lin, S. Rommer, S. Euler, E. A. Yavuz, and R. S. Karlsson, "5G from space: An overview of 3GPP non-terrestrial networks," *IEEE Commun. Stand. Mag.*, accepted, doi: 10.1109/MCOMSTD.011.2100038.

[9] P. Wang *et al*, "Convergence of satellite and terrestrial networks: A comprehensive survey," *IEEE Access*, vol. 8, pp. 5550–5588, 2020.

[10] H. Wu, J. Li, H. Lu, and P. Hong, "A two-layer caching model for content delivery services in satellite-terrestrial networks," in *Proc. IEEE Glob. Commun. Conf. (GLOBECOM)*, Dec. 2016, pp. 1–6.

[11] A. Kalantari *et al.*, "Cache-assisted hybrid satellite-terrestrial backhauling for 5G cellular networks," in *Proc. IEEE Glob. Commun. Conf. (GLOBECOM)*, Dec. 2017, pp. 1–6.

[12] T. X. Vu *et al.*, "Efficient 5G edge caching over satellite," in *Proc. 36th Int. Commun. Satellite Syst. Conf. (ICSSC)*, Oct. 2018, pp. 1–5.

[13] K. An, Y. Li, X. Yan, and T. Liang, "On the performance of cache-enabled hybrid satellite-terrestrial relay networks," *IEEE Wireless Commun. Lett.*, vol. 8, no. 5, pp. 1506–1509, Oct. 2019.

[14] X. Zhang *et al.*, "On the performance of hybrid satellite-terrestrial content delivery networks with non-orthogonal multiple access," *IEEE Wireless Commun. Lett.*, vol. 10, no. 3, pp. 454–458, Mar. 2021.

[15] X. Wang, H. Li, T. Lan, and Q. Wu, "Overlay coded multicast for edge caching in 5G-satellite integrated networks," in *Proc. IEEE Wireless Commun. and Netw. Conf. (WCNC)*, May 2020, pp. 1–7.

[16] N. Jindal and Z. Luo, "Capacity limits of multiple antenna multicast," in *Proc. IEEE Int. Symp. Inf. Theory (ISIT)*, 2006, pp. 1841–1845.

[17] H. Zhao, A. Bazco-Nogueras, and P. Elia, "Resolving the worst-user bottleneck of coded caching: Exploiting finite file sizes," in *Proc. IEEE Inf. Theory Workshop (ITW)*, Apr. 2021, pp. 1–5.

[18] H. Zhao, A. Bazco-Nogueras and P. Elia, "Wireless coded caching can overcome the worst-user bottleneck by exploiting finite file sizes," *IEEE Trans. Wireless Commun.*, accepted, doi: 10.1109/TWC.2022.3140895.

[19] "Space communication calculations," Australian Space Academy, 2019, accessed on: 21/10/2021. [Online]. Available: http://www.spaceacademy.net.au/spacelink/spcomcalc.htm

[20] A. Abdi, W. Lau, M.-S. Alouini, and M. Kaveh, "A new simple model for land mobile satellite channels: first- and second-order statistics," *IEEE Trans. Wireless Commun.*, vol. 2, no. 3, pp. 519–528, May 2003.

[21] P. K. Sharma *et al.*, "Performance analysis of overlay spectrum sharing in hybrid satellite-terrestrial systems with secondary network selection," *IEEE Trans. Wireless Commun.*, vol. 16, no. 10, pp. 6586–6601, Oct. 2017.

[22] G. Pan, J. Ye, Y. Tian, and M.-S. Alouini, "On HARQ schemes in satellite-terrestrial transmissions," *IEEE Trans. Wireless Commun.*, vol. 19, no. 12, pp. 7998–8010, Dec. 2020.

[23] I. S. Gradshteyn and I. M. Ryzhik, *Table of integrals, series, and products*, 7th ed. Academic press, 2007.

[24] A. Erdelyi, "Transformation of a certain series of products of confluent hypergeometric functions. Applications to Laguerre and Charlier polynomials," *Compos. Math.*, vol. 7, pp. 340–352, 1940.

[25] T. Wu and V. K. Lau, "Robust rate, power and precoder adaptation for slow fading MIMO channels with noisy limited feedback," *IEEE Trans. Wireless Commun.*, vol. 7, no. 6, pp. 2360–2367, Jun. 2008.

[26] K. P. Liolis, A. D. Panagopoulos, and P.-D. Arapoglou, "An analytical unifying approach for outage capacity achieved in SIMO and MISO broadband satellite channel configurations," in *3rd European Conf. on Antennas and Propag.*, Mar. 2009, pp. 2911–2915.

[27] H. Zhao, A. Bazco-Nogueras, and P. Elia, "Coded caching in land mobile satellite systems," 2021. [Online]. Available: https://box.networks.imdea.org/s/0fptig1dnMsTz8M