

Correspondence

Abstract—This paper examines the integration of vector coded caching (VCC) into multi-beam satellite communications (SATCOM) systems and demonstrates that even limited receiver-side caching can substantially enhance spectral efficiency. By leveraging cached content to suppress interference, VCC enables the concurrent transmission of multiple precoded signal vectors that would otherwise require separate transmission resources. This leads to a multiplicative improvement in resource utilization in SATCOM. To characterize this performance, we model the satellite-to-ground channel using Rician-shadowed fading and after incorporating practical considerations such as matched-filter precoding, channel state information (CSI) acquisition overhead as well as CSI imperfections at the transmitter, we here derive closed-form expressions for the average sum rate and spectral efficiency gain of VCC in SATCOM. Our analysis, tightly validated through numerical simulations, reveals that VCC can yield spectral efficiency gains of 300% to 550% over traditional multi-user MISO SATCOM with the same resources. These gains—which have nothing to do with multicasting, prefetching gains nor file popularity—highlight VCC as a pure physical-layer solution for future high-throughput SATCOM systems, significantly narrowing the performance gap between satellite and wired networks.

Index Terms—Vector coded caching, satellite communications, spectral efficiency, Rician-shadowed fading.

I. Introduction

The rapid growth of satellite communications (SATCOM) traffic has created an urgent demand for higher spectral efficiency. To address this, *multi-beam satellites* have gained significant attention [1], employing an array of L radio-frequency (RF) feeds to generate multiple beams that reuse the same frequencies. Naturally many challenges remain—like for example that of *inter-beam interference* [2]—leading to a fundamental “*few-feeds-many-users*” mismatch [3] where simply the number of active ground terminals K often far exceeds the number of available feeds L , thus resulting in precoding-enabled multi-beam SATCOM systems that fall short of

This work was supported in part by EURECOM’s industrial members: ORANGE, BMW, SAP, iABG, Norton LifeLock; in part by the ERC-PoC Project LIGHT (Grant 101101031) and the EU H2030 project CONVERGE, in part by the French projects EEMW4FIX (ANR), PERSEUS (PEPR-5G) and FOUNDS (PEPR-5G); and in part by the Huawei France funded Chair towards Future Wireless Networks.

Authors’ addresses: Hui Zhao, Dirk Slock and Petros Elia are with the Communication Systems Department, EURECOM, Sophia Antipolis, France. (E-mail: hui.zhao@eurecom.fr; dirk.slock@eurecom.fr; elia@eurecom.fr) (*Corresponding author: Hui Zhao.*)

Mentions of supplemental materials and animal/human rights statements can be included here.

Color versions of one or more of the figures in this article are available online at <http://ieeexplore.ieee.org>.

the spectral efficiency required to support rapidly growing user demands of large content.

At the same time, this growth of traffic is largely due to Video on Demand (VoD), and thus, as one would expect, caching has been used as a means of alleviating this traffic. However, due to the relatively modest size of receiver-side caches, as compared to the immense libraries of content, traditional prefetching techniques result in very modest gains [4]–[7]. Even the modern techniques of coded caching [4] that cleverly leverage receiver-side caching and coded multicast transmissions to directly handle interference, do not provide substantial spectral efficiency gains in SATCOM networks with large antenna arrays and many beams, mainly because such coded caching solutions were intended for single-antenna settings. Additionally, even advanced *multi-antenna coded caching* schemes [8] offer limited gains over practical multi-beam downlink systems. Consider a downlink system with L transmit antennas/feeds serving K receivers, each with a cache of size equal to a fraction $\gamma \in [0, 1]$ of the entire library, and in the presence of Λ distinct cache states.¹ Most multi-antenna coded caching approaches (e.g., [9]–[12]) achieve a Degrees-of-Freedom (DoF) performance of $L + \Lambda\gamma$, corresponding to a caching gain $\Lambda\gamma \ll L$ that is forced to be very small due to subpacketization (file-size) constraints [13]. As a result, practical DoF performance remains close to L , offering very modest improvement over cacheless multi-user (MU) multiple-input and single-output (MISO) baselines.

Vector coded caching (VCC) fundamentally transforms this scenario by enabling a multiplicative improvement in spectral efficiency compared to traditional MU-MISO [13], [14]. Rather than using the widely adopted XOR-based coded multicasting, VCC instead fuses multiple precoded $L \times 1$ signal vectors into a single superimposed $L \times 1$ signal vector. This superimposed vector, as well as the cache placement, are carefully designed so that the interference—resulting from overloading the system with multiple vectors in one shot—can be ‘cached-out’ at the receiver side. Thus, VCC allows for great efficiencies compared to traditional MU-MISO systems that would require each precoded signal vector to be sent *sequentially*. This new VCC approach overcomes the constraints of conventional multi-antenna coded caching, and—by exploiting both caching and antenna resources—achieves a theoretical DoF of $L(\Lambda\gamma + 1)$, yielding multiplicative DoF gains over downlink MU-MISO systems.

In exploring the actual throughput gains of VCC, recent studies (cf. [15]–[17]) have confirmed the effectiveness of VCC in urban Macro-cell and Micro-cell environments under various realistic considerations, including finite file sizes (i.e., modest values of $\Lambda\gamma$), various fading and pathloss characteristics, and various costs of acquiring channel state information (CSI). In this context, VCC demonstrates a *multiplicative* improvement in spectral

efficiency over independently optimized state-of-art MU-MISO terrestrial networks, validating its effectiveness in complex, high-density deployment scenarios.

Despite the reported performance gains in terrestrial networks [15]–[17], the application of VCC in SATCOM remains largely unexplored. It is worth noting that VCC is fundamentally different from prior SATCOM coded caching works (e.g., [18], [19]), where satellites collectively store the library and users retrieve files by connecting sequentially over an orbital cycle using XOR-based coded multicasting [4], [20] to achieve global coverage.

SATCOM is not only a pertinent setting for applying VCC, but—as it turns out—also enjoys a privileged relationship with VCC. To begin with, SATCOM is heavily constrained by low spectral efficiency—which VCC directly improves—and its most recent advancements are closely tied to VoD [14], [17]. More interestingly, we also see certain synergies. For example, VCC thrives in the presence of many users with statistically symmetric links; two characteristics that are inherent in SATCOM systems, where link uniformity extends over terrestrial coverages with radii of several tens to one hundred kilometers [21]–[23] that encompasses a far larger user population compared to cellular networks. Furthermore, SATCOM systems endure substantially degraded CSI at the transmitter (CSIT) [2], [24], which plays to the strength of VCC which migrates much of the burden of interference cancellation away from the transmitter. At the same time though, SATCOM systems are dominated by reduced received SNR values, which can be particularly problematic for VCC which operates by simultaneously serving multiple precoded signal vectors, thus leaving each vector with diminished power.

Driven by the rapidly increasing demand for higher spectral efficiency in multi-beam SATCOM networks [1]–[3] and the aforementioned synergies between VCC and SATCOM, *our aim is to analyze the extent to which VCC enhances SATCOM spectral efficiency*, while carefully accounting for the fundamentally different characteristics of SATCOM compared to traditional cellular settings. Toward this, we model the satellite-to-ground channel as a Rician-shadowed fading channel [25]–[30], and place our focus on the practical case of matched-filter (MF) precoding [31], all WHILE accounting for CSI acquisition costs and the effect of imperfect CSIT. The primary technical contribution of this work is the *parametrization of VCC delivery performance in SATCOM systems*, achieved through simple closed-form expressions that capture both the overall system spectral efficiency and the multiplicative effective gain over the cacheless MU-MISO baseline. These analytical expressions are shown to be highly accurate when validated against Monte Carlo simulations. In the end, the results reveal that under the aforementioned practical CSIT and SNR considerations, VCC offers a very sizable multiplicative spectral efficiency gain compared to optimized conventional MU-MISO, highlighting its strong potential for future high-demand satellite networks.

¹This aspect of cache states is an esoteric aspect of coded caching, which though turns out to be its Achilles “hill”.

Notations: For a positive integer n , we use $[n]$ to denote the set $\{1, 2, \dots, n\}$. For two sets \mathcal{A} and \mathcal{B} , the notation $\mathcal{A} \setminus \mathcal{B}$ represents the set difference between \mathcal{A} and \mathcal{B} . The operator $|\cdot|$ represents the cardinality of a set or the magnitude of a complex number, depending on the context. $\mathbb{E}\{\cdot\}$ denotes the expectation operator, while $\text{Tr}\{\cdot\}$ represents the trace operator. For a matrix \mathbf{A} , we use \mathbf{A}^T , \mathbf{A}^* , and \mathbf{A}^H to denote the transpose, element-wise complex conjugate, and conjugate transpose of \mathbf{A} , respectively. The notation $\mathbf{0}_L$ represents an $L \times 1$ vector of zeros, and \mathbf{I}_L denotes the $L \times L$ identity matrix. We use \mathcal{CN} to denote the complex Gaussian distribution.

Paper Organization: Section II introduces the system architecture of the considered SATCOM scenario, followed by a detailed description of the VCC strategy along with the adopted precoding approach. Section III focuses on the theoretical analysis of the content delivery performance under the proposed caching framework. Section IV presents simulation results to corroborate the theoretical findings and to demonstrate the performance gain achieved by VCC when compared to systems without caching. Finally, Section V summarizes the main conclusions of this work. For the sake of clarity, several mathematical proofs are deferred to the Appendix.

II. System Model

We consider a SATCOM system in which a satellite equipped with L transmit antennas/feeds serves K cache-aided single-antenna ground users. Each user requests *distinct* content files from a library $\mathcal{F} \triangleq \{W_1, \dots, W_N\}$ containing N equal-sized files. The satellite has full access to the library \mathcal{F} , facilitated by a high-speed feeder link connecting the satellite to a terrestrial satellite gateway, which in turn links to the core network [32]. During an off-peak period (for example, once a month), each user can store (and occasionally update) a fraction $\gamma \in [0, 1]$ of the library content, where γ is the normalized cache size compared to the library size². During all the intermediate transmission phases (peak times), the satellite must deliver individual content to its different users, where each user gets their own content, as is typical in VoD.

We adopt a static (quasi-static) model for the satellite–terrestrial channel, where the channel coefficients remain constant within a transmission block and may vary independently across different blocks. To justify this assumption, we consider that land terminals are located within a coverage radius of 100 km, with a circular LEO orbit at an altitude of 600 km above the coverage center. Satellites are spaced 75 km apart, corresponding to an inter-satellite central angle of $\theta \approx 0.674^\circ$. A handover is triggered once the angle between the satellite and the line connecting the Earth’s center to the coverage center exceeds $\theta/2$. This corresponds to a change in the

satellite’s zenith angle from about 0.591° to 1.208° as observed by land terminals located at the opposite edge of the coverage area, indicating an extremely small variation. With precise beam alignment to the serving satellite, the relative geometry between the serving satellite and the covered terminals remains nearly unchanged. Under such conditions and Doppler shift compensation, the LEO-to-ground channel can be regarded as being a static fading channel, similar to that of GEO satellite–terrestrial links.

Under the static channel model, we consider the Rician-shadowed fading channel [25], a common model for satellite-to-ground communication that accounts for both line-of-sight (LOS) and non-line-of-sight (NLOS) components. This model is defined by three key parameters: the power of the scattering component, denoted by 2β ; the power of the LOS component, represented by Ω ; and the parameter m reflecting the extent of LOS obstruction caused by environmental elements like buildings, trees, or terrain. Here, $m = 0$ represents a completely obstructed LOS, while $m \rightarrow \infty$ indicates no obstruction. The Rician-shadowed fading model can accommodate different types of orbits and frequency bands (e.g., Ku-band and Ka-band) [26]–[30]. The channel vector between the satellite and the k -th user ($k \in [K]$) is represented by $\mathbf{h}_k \in \mathbb{C}^{L \times 1}$, which can be expressed as

$$\mathbf{h}_k = Z_k \mathbf{t}_k + \mathbf{h}'_k, \quad (1)$$

where $\mathbf{h}'_k \sim \mathcal{CN}(\mathbf{0}_L, 2\beta \mathbf{I}_L)$ denotes the scattered component; $\mathbf{t}_k \triangleq [\exp(j\theta_k^{(1)}), \dots, \exp(j\theta_k^{(L)})]^T$ collects the LOS phasor from the transmit antennas to user k ; and the real-valued scalar Z_k —following an Nakagami- m distribution with shape parameter m and scale parameter Ω —models the LOS amplitude fluctuation. In the considered LEO-to-ground scenario with slowly moving users, the phase of the LOS component from each transmit antenna to the ground user would, in principle, be deterministic and highly correlated, as it is set by the array geometry and the satellite–user geometry. Nevertheless, there are several reasons to model these LOS phases differently. First, relative satellite motion, oscillator drifts, and hardware-induced phase jitter introduce effective random phase offsets across antennas, which decorrelate LOS contributions over time. Second, deliberate phase dithering is sometimes employed in practice to randomize the LOS component and reduce systematic interference patterns. Accordingly, for analytical tractability we adopt a block-fading abstraction: within a channel coherence block the LOS phases are fixed, whereas across different channel coherence blocks they are modeled as i.i.d. and uniformly distributed over $[0, 2\pi]$.

In a given transmission round, let $\mathbf{x} \in \mathbb{C}^{L \times 1}$, with $\mathbb{E}\{\|\mathbf{x}\|^2\} = P_t$, denote the transmit signal from the satellite that conveys the requested messages for *multiple users*. The received signal at user k is of the form

$$y_k = \mathbf{h}_k^T \mathbf{x} + z_k, \quad (2)$$

where $z_k \sim \mathcal{CN}(0, 1)$ represents the additive white Gaussian noise (AWGN). Hence, P_t here accounts for

²In principle, one would consider the part of the library corresponding to a sizable fraction of the overall VoD traffic—for example 90% of traffic (cf. [14, Example 1]).

Algorithm 1 Vector Coded Caching (VCC) Scheme

Require: Library $\mathcal{F} = \{W_1, \dots, W_N\}$, total number of users K , cache size γ , number of cache states Λ , multiplexing gain Q , requested file indices $\{d_k\}_{k \in [K]}$

Ensure: All users recover their requested files

- 1: **Cache Placement Phase:**
- 2: **for** each file $W_n, n \in [N]$ **do**
- 3: Split W_n into $\binom{\Lambda}{\Lambda\gamma}$ non-overlapping and equal-sized subfiles $\{W_n^{\mathcal{T}} : \mathcal{T} \subseteq [\Lambda], |\mathcal{T}| = \Lambda\gamma\}$
- 4: **end for**
- 5: Partition users into Λ groups, each with $B = K/\Lambda$ users and a dedicated cache state
- 6: **for** each cache state $g \in [\Lambda]$ **do**
- 7: Users in group g cache the identical content
- 8: $\mathcal{Z}_g = \{W_n^{\mathcal{T}} : \mathcal{T} \ni g, \forall n \in [N]\}$
- 9: **end for**

- 9: **Content Delivery Phase:**
- 10: **for** each subset $\Psi \subseteq [\Lambda]$ with $|\Psi| = \Lambda\gamma + 1$ **do**
- 11: **for** each transmission round $r = 1, \dots, B/Q$ **do**
- 12: Select Q users from each cache state $\psi \in \Psi$
- 13: Form symbols $s_{\psi,b}$ from subfiles $W_{d_{\psi,b}}^{\Psi \setminus \{\psi\}}$
- 14: Construct the transmit signal \mathbf{x} in (3)
- 15: **for** each received signal (4) **do**
- 16: Cancel inter-group interference using cached content and composite CSI
- 17: Decode desired symbol $s_{\psi,b}$ by handling intra-group interference via precoding
- 18: **end for**
- 19: **end for**
- 20: **end for**

the effects of the actual transmit power, the large-scale pathloss, antenna gains, and the actual AWGN power.

Next, we demonstrate how VCC enables the efficient use of each user's cached content by detailing the design of the transmit signal vector \mathbf{x} , and showing how users exploit their local cache to decode their intended data while mitigating interference from others.

A. Vector Coded Caching

This subsection outlines the core design of VCC. Similar to conventional cache-aided communications, VCC operates in two phases, namely the *cache placement phase* and the *content delivery phase*, as summarized in **Algorithm 1**. We describe each phase in detail below.

1. Cache Placement Phase

During the placement phase, each file W_n from the library $\mathcal{F} = \{W_1, W_2, \dots, W_N\}$ is divided into $\binom{\Lambda}{\Lambda\gamma}$ non-overlapping and equal-sized subfiles, each labelled by some $\Lambda\gamma$ -tuple \mathcal{T} with $\mathcal{T} \subseteq [\Lambda]$. Users are partitioned into Λ user groups, where each group $g \in [\Lambda]$ contains $B = \frac{K}{\Lambda}$ users, all caching the same subfiles. Specifically,

for users in group g , the cached content is of the form $\mathcal{Z}_g = \{W_n^{\mathcal{T}} : \mathcal{T} \subseteq [\Lambda], |\mathcal{T}| = \Lambda\gamma, \mathcal{T} \ni g, \forall n \in [N]\}$. It is straightforward to verify that this placement scheme adheres to the storage capacity γ at each user³. We refer to users storing the same content as being in the same *cache state*. This setup ensures that users with distinct cache states store some overlapping content, enabling interference management during the delivery phase.

2. Content Delivery Phase

In the delivery phase, there are $\binom{\Lambda}{\Lambda\gamma+1}$ transmission stages. Each stage comprises a unique selection of cache states $\Psi \subseteq [\Lambda]$ with $|\Psi| \triangleq G = \Lambda\gamma + 1$, with Q users sharing the same cache state selected for service in each transmission round. Here, Q can be interpreted as the multiplexing gain provided by transmit antennas. Thus, there are B/Q transmission rounds in each stage. We use $\mathbf{U}_{\psi,b}$ to denote the b -th *active* user in the user-group ψ for any $b \in [Q]$ and $\psi \in [\Lambda]$, and $d_{\psi,b} \in [N]$ denotes the file index requested by user $\mathbf{U}_{\psi,b}$. The signal vector $\mathbf{x} \in \mathbb{C}^{L \times 1}$ in this transmission round is formulated as

$$\mathbf{x} = \alpha \sum_{\psi \in \Psi} \sum_{b \in [Q]} \mathbf{v}_{\psi,b} s_{\psi,b} = \alpha \sum_{\psi \in \Psi} \mathbf{V}_{\psi} \mathbf{s}_{\psi}, \quad (3)$$

where $s_{\psi,b} \in \mathbb{C}$, generated from the subfile $W_{d_{\psi,b}}^{\Psi \setminus \{\psi\}}$, represents the signal symbol intended for user $\mathbf{U}_{\psi,b}$, and $\mathbf{v}_{\psi,b} \in \mathbb{C}^{L \times 1}$ is the precoding vector for $s_{\psi,b}$. In the above, $\mathbf{s}_{\psi} \triangleq [s_{\psi,1}, \dots, s_{\psi,Q}]^T \in \mathbb{C}^{Q \times 1}$ and $\mathbf{V}_{\psi} \triangleq [\mathbf{v}_{\psi,1}, \dots, \mathbf{v}_{\psi,Q}] \in \mathbb{C}^{L \times Q}$ denote the signal vector and the corresponding precoding matrix for the Q cache-sharing users. Moreover, α is responsible for normalizing the average power of \mathbf{x} into P_t .

The received signal at user $\mathbf{U}_{\psi,b}$ is expressed as (4), shown at the top of the next page. According to the cache placement design, user $\mathbf{U}_{\psi,b}$ has cached the subfiles $\{W_n^{\Psi \setminus \{\phi\}} : \phi \in \Psi \setminus \{\psi\}, \forall n \in [N]\}$, which will be used to cancel inter-group interference from other groups $\phi \in \Psi \setminus \{\psi\}$. Furthermore, the intra-group interference among users within the same group can be managed by designing the precoder \mathbf{V}_{ψ} appropriately (cf. Section II-B). This setup provides a theoretical multiplexing gain represented by the DoF GQ , achieving a multiplicative spectral efficiency boost over traditional cacheless MU-MISO systems. After completing $\binom{\Lambda}{\Lambda\gamma+1}$ stages, all users receive their requested files *in full*. We refer to [13], [14] for more details.

Remark 1. At the receiver side, VCC requires additional processing to suppress interference that cannot be handled by conventional precoding alone. Unlike decoding-based physical-layer (PHY) techniques such as non-orthogonal multiple access (NOMA), this cache-aided interference suppression does not rely on decoding the interfering streams. Instead, the receiver exploits its cached content and the composite CSI to locally regenerate the known

³We do not account for the cost of placement, which only occurs rarely (once a month perhaps), incrementally, and certainly during off-peak hours. This cost does not match the constant cost of content delivery, especially during the crucial peak hours.

$$y_{\psi,b} = \mathbf{h}_{\psi,b}^T \mathbf{x} + z_{\psi,b} = \alpha \mathbf{h}_{\psi,b}^T \mathbf{v}_{\psi,b} s_{\psi,b} + \underbrace{\alpha \mathbf{h}_{\psi,b}^T \sum_{b' \neq b, b' \in [Q]} \mathbf{v}_{\psi,b'} s_{\psi,b'}}_{\text{intra-group interference}} + \underbrace{\alpha \mathbf{h}_{\psi,b}^T \sum_{\phi \neq \psi, \phi \in \Psi} \mathbf{V}_{\phi} \mathbf{s}_{\phi}}_{\text{inter-group interference}} + z_{\psi,b} \quad (4)$$

interfering components and subtract them *jointly* from the received signal. This operation involves only linear processing and constitutes the primary additional PHY complexity introduced by VCC. The associated composite CSI cost will be accounted for in the CSI overhead parameter $\xi_{G,Q}$ (cf. (8)) when evaluating the performance gains over the cacheless baseline.

Remark 2. Unlike content delivery networks (CDNs), which pre-position entire files based on demand prediction, VCC does not bypass transmission by guessing user preferences or pre-loading complete video content. Instead, it uses receiver-side storage to enable PHY multiplexing gains: specifically, to support the decoding of densely packed unicast streams delivered simultaneously through structured precoding. The local storage is thus employed as an enabler of advanced signal separation—allowing multiple uniquely addressed streams to be resolved from a shared transmission. The speedup in spectral efficiency will not arise from pre-downloading content, but rather from a fundamental improvement in the precoding structure.

Next, we describe the precoder for VCC in SATCOM.

B. Matched-Filter (MF) Precoding

We consider MF precoding due to its low complexity and favorable performance in the low-SNR regime typical of SATCOM, where it approaches the spectral efficiency of the more complex MMSE precoding [24], [31], [33].

As seen (cf. (4)), each user must have knowledge of the global CSI in order to cancel inter-group interference. The CSI dissemination process follows, for example, a standard FDD uplink–downlink training framework (cf. [2], [34]) and consists of three phases. In the first phase, the satellite transmits pilot symbols to the served users. In the second, each user feeds back its locally estimated CSI to the satellite gateway, typically via codebook-based quantization that conveys only a discrete index; this limited feedback inevitably leads to quantization errors at the gateway, resulting in imperfect CSIT. Finally, in the third phase, the gateway broadcasts the aggregated global CSI to all GQ users served simultaneously.

Let $\hat{\mathbf{h}}_{\psi,b}$ represent the estimated channel vector for user $\mathbf{U}_{\psi,b}$, obtained via the maximum likelihood (ML) estimator. The estimated channel vector can be expressed as [35] $\hat{\mathbf{h}}_{\psi,b} = \mathbf{h}_{\psi,b} + \tilde{\mathbf{h}}_{\psi,b}$, where $\tilde{\mathbf{h}}_{\psi,b} \in \mathbb{C}^{L \times 1}$ is the estimation error. Each element of $\tilde{\mathbf{h}}_k$ follows an i.i.d. complex Gaussian distribution with zero mean and variance σ_e^2 . Furthermore, $\mathbf{h}_{\psi,b}$ and $\tilde{\mathbf{h}}_{\psi,b}$ are independent.

Under MF precoding, the transmit signal \mathbf{x} from (3) becomes

$$\mathbf{x} = \alpha \sum_{\psi \in \Psi} \hat{\mathbf{H}}_{\psi}^H \mathbf{s}_{\psi}, \quad (5)$$

where $\hat{\mathbf{H}} \triangleq [\hat{\mathbf{h}}_{\psi,1}, \dots, \hat{\mathbf{h}}_{\psi,Q}]^T \in \mathbb{C}^{Q \times L}$ represents the estimated channel matrix corresponding to the channel from the satellite to the selected Q users in the user-group ψ , who—as we recall—share the same cached content.

Remark 3. The computational complexity of $\hat{\mathbf{H}}_{\psi}^H$ is $\mathcal{O}(QL)$, since it only involves a conjugate transpose operation on $\hat{\mathbf{H}}_{\psi}$. In contrast, the zero-forcing (ZF) precoder $\hat{\mathbf{H}}_{\psi}^H (\hat{\mathbf{H}}_{\psi} \hat{\mathbf{H}}_{\psi}^H)^{-1}$ requires the formation of a $Q \times Q$ Gram matrix, a matrix inversion, and an additional matrix multiplication, resulting in an overall computational complexity of $\mathcal{O}(QL + Q^2L + Q^3 + LQ^2)$. The MMSE precoder further introduces a diagonal loading term, leading to a slightly higher computational complexity than that of ZF. Therefore, ZF/MMSE precoding is significantly more computationally demanding than MF precoding.

After removing inter-group interference via locally cached content and perfect downlink channel training, the received signal at user $\mathbf{U}_{\psi,b}$ who aims to decode its intended symbol $s_{\psi,b}$, is given by

$$y'_{\psi,b} = \alpha \mathbf{h}_{\psi,b}^T \hat{\mathbf{h}}_{\psi,b}^* s_{\psi,b} + \alpha \sum_{b' \in [Q] \setminus \{b\}} \mathbf{h}_{\psi,b}^T \hat{\mathbf{h}}_{\psi,b'}^* s_{\psi,b'} + z_{\psi,b}.$$

With common Gaussian signaling, the signal-to-interference-plus-noise ratio (SINR) for decoding $s_{\psi,b}$ at user $\mathbf{U}_{\psi,b}$ takes the form

$$\text{SINR}_{\psi,b} = \frac{\alpha^2 |\mathbf{h}_{\psi,b}^T \hat{\mathbf{h}}_{\psi,b}^*|^2}{1 + \alpha^2 \sum_{b' \in [Q] \setminus \{b\}} |\mathbf{h}_{\psi,b}^T \hat{\mathbf{h}}_{\psi,b'}^*|^2}. \quad (6)$$

The effective rate for $\mathbf{U}_{\psi,b}$ can be expressed as

$$R_{\psi,b} = \xi_{G,Q} \log_2 (1 + \text{SINR}_{\psi,b}), \quad (7)$$

where the term $\xi_{G,Q} \triangleq 1 - \frac{GQ\Theta}{T}$ accounts for the CSI acquisition overhead, T denotes the channel coherence block length (in symbols), and Θ is the total pilot length per user and per block. The effective sum rate for simultaneously serving GQ users in VCC is

$$R_{\text{sum}}(G, Q) = \xi_{G,Q} \sum_{\psi \in \Psi} \sum_{b \in [Q]} \log_2 (1 + \text{SINR}_{\psi,b}). \quad (8)$$

We note that setting $G = 1$ corresponds to the cacheless (traditional MU-MISO) counterpart (cf. (3) and (4)). To quantify the performance gain over conventional cacheless MU-MISO systems, we define the *effective gain* as⁴

$$\mathcal{G} = \frac{\max_{Q \in \{2, \dots, Q_{\max}\}} \mathbb{E}\{R_{\text{sum}}(G, Q)\}}{\max_{Q' \in \{2, \dots, Q'_{\max}\}} \mathbb{E}\{R_{\text{sum}}(1, Q')\}}, \quad (9)$$

⁴This paper examines the role of precoding in cache-aided MU-MISO systems. The case $Q = Q' = 1$, in which both VCC and the cacheless system operate without precoding, is not considered. In practice, a spot beam's spatial resolution limits how many users can share the same time–frequency resources; typically only a few (≤ 10) can be multiplexed simultaneously, so $Q_{\max}, Q'_{\max} \leq 10$, even for $L \gg 10$.

where the expectation is taken over the channel realizations across the statistically symmetric users. This allows both cache-aided and cacheless schemes to achieve their *optimized tradeoff* between multiplexing and beamforming gains. The effective gain \mathcal{G} thus captures the multiplicative boost in spectral efficiency over optimized conventional MU-MISO systems at finite SNRs. As one might expect, achieving the full theoretical gain of G corresponds to idealized high SNR conditions [13].

Remark 4. As indicated in (9), this work does not concern it self with the local caching gain⁵, which refers to the reduction in delivery load achieved by storing parts of the requested content at the user side in advance. We repeat that this means that the gains that we record, are certainly not an outcome of the fact that a fraction of the data need not be communicated because it is cached. Instead, the gains here are from comparing the delivery rates of the data that is in fact communicated during delivery.

III. Main Results

In this section, we derive closed-form expressions for the average sum rate and the effective gain of VCC in SATCOM systems under imperfect CSIT, with the wireless channel modeled using Rician-shadowed fading. For analytical tractability, we assume that all users experience identical channel parameters (i.e., the same m , β , and Ω), which preserves the generality of the insights when focusing on average system performance [22], [26], [28].

We first have the following result for the statistics of the estimated channel $\hat{\mathbf{h}}_{\psi,b}$.

Proposition 1. In the ML estimation of the channel vector $\mathbf{h}_{\psi,b}$ over Rician-shadowed fading channels, each element of $\hat{\mathbf{h}}_{\psi,b}$ shares the same statistical parameters m and Ω as the corresponding element of $\mathbf{h}_{\psi,b}$, differing only in the average scattering component power, which becomes $2\beta + \sigma_e^2$.

Proof: Let $h_{\psi,b}^{(\ell)}$ denote the ℓ -th element of $\mathbf{h}_{\psi,b}$. For the channel gain $h_{\psi,b}^{(\ell)}$ over Rician-shadowed fading channels, we can write it as (cf. [25]) $h_{\psi,b}^{(\ell)} = h_{\text{LOS}} + h_{\text{NLOS}}$ where h_{LOS} represents the random fluctuation of the LOS component. In the above, h_{NLOS} models the NLOS component, which follows a complex Gaussian distribution with zero-mean and variance 2β . For the ℓ -th element of $\hat{\mathbf{h}}_{\psi,b}$, denoted by $\hat{h}_{\psi,b}^{(\ell)}$, we can write it as [35] $\hat{h}_{\psi,b}^{(\ell)} = h_{\psi,b}^{(\ell)} + \tilde{h}_{\psi,b}^{(\ell)} = h_{\text{LOS}} + h_{\text{NLOS}} + \tilde{h}_{\psi,b}^{(\ell)}$, where $\tilde{h}_{\psi,b}^{(\ell)} \sim \mathcal{CN}(0, \sigma_e^2)$ denotes the ML estimation error. Note that $\tilde{h}_{\text{NLOS}} \triangleq h_{\text{NLOS}} + \tilde{h}_{\psi,b}^{(\ell)}$ follows a complex Gaussian distribution with zero-mean and variance $2\beta + \sigma_e^2$. Thus, $\hat{h}_{\psi,b}^{(\ell)} = h_{\text{LOS}} + \tilde{h}_{\text{NLOS}}$ forms a new Rician-shadowed channel gain as described in Proposition 1. ■

Next, we have the closed-form expression for α^2 under MF precoding.

⁵We remind the reader (cf. [4]) that in practice, such gains are essentially trivial since the storage capacity of end-user devices [4] is very small compared to the entire content library.

Proposition 2. The squared power control factor α^2 under MF precoding with imperfect CSIT is of the form

$$\alpha^2 = \frac{P_t}{GQL(2\beta + \sigma_e^2 + \Omega)}. \quad (10)$$

Moreover, by setting $G = 1$ and $Q = Q'$ in α^2 , we have the squared power control factor α_0^2 in the cacheless MU-MISO counterpart.

Proof: The proof is relegated to Appendix A. ■

We also have the following.

Proposition 3. For

$$\Xi_1 \triangleq \mathbb{E}\{||\mathbf{h}_{\psi,b}||^4\}, \quad \Xi_2 \triangleq \mathbb{E}\{|\mathbf{h}_{\psi,b}^T \hat{\mathbf{h}}_{\psi,b}^*|^2\}, \quad (11)$$

then

$$\begin{aligned} \Xi_1 &= L \left[\left(4\beta^2 + 4\beta\Omega + \frac{\Omega^2}{m} \right) + (2\beta + \Omega)^2 \right] \\ &\quad + L(L-1) \left[\left(1 + \frac{1}{m} \right) \Omega^2 + 4\beta\Omega + 4\beta^2 \right] \end{aligned} \quad (12)$$

$$\Xi_2 = L(2\beta + \Omega)(2\beta + \sigma_e^2 + \Omega). \quad (13)$$

Proof: The proof is relegated to Appendix B. ■

In the following, we present the average sum rate and the corresponding effective gain of VCC under MF precoding with imperfect CSIT. We recall that the power control parameters α^2 and α_0^2 have been derived in Proposition 2, while Ξ_1 and Ξ_2 are presented in Proposition 3.

Theorem 1. Under MF precoding and in the presence of imperfect CSIT, the average sum rate $\bar{R}_{\text{sum}} \triangleq \mathbb{E}\{R_{\text{sum}}\}$ can be approximated as

$$\bar{R}_{\text{sum}} \approx \xi_{G,Q} GQ \log_2 \left(1 + \frac{\alpha^2(\Xi_1 + \sigma_e^2 L(2\beta + \Omega))}{1 + \alpha^2(Q-1)\Xi_2} \right), \quad (14)$$

and the effective gain can be approximated as

\mathcal{G}

$$\begin{aligned} &\approx \frac{\max_{Q \in \{2, \dots, Q_{\max}\}} GQ \xi_{G,Q} \log_2 \left(1 + \frac{\alpha^2(\Xi_1 + \sigma_e^2 L(2\beta + \Omega))}{1 + \alpha^2(Q-1)\Xi_2} \right)}{\max_{Q' \in \{2, \dots, Q'_{\max}\}} Q' \xi_{1,Q'} \log_2 \left(1 + \frac{\alpha_0^2(\Xi_1 + \sigma_e^2 L(2\beta + \Omega))}{1 + \alpha_0^2(Q'-1)\Xi_2} \right)}. \end{aligned} \quad (15)$$

Proof: The proof is relegated to Appendix C. ■

Remark 5. Our analytical method differs from conventional approaches (e.g., [29], [30]) that also adopt Rician-shadowed fading for modeling terrestrial-satellite channels. In those studies, the typical procedure involves first expressing the performance metrics (e.g., average capacity) as infinite integrals derived from the probability density function (PDF) of the channel gain, which are then evaluated or approximated to obtain closed-form results. In contrast, our approach avoids direct integration over the PDF and instead computes a few simple expectations. This not only improves tractability and yields a compact expression for parameterizing the delivery performance of VCC in SATCOM, but also ensures that the approximation becomes increasingly tight as L and Q grow.

TABLE I: Parameters in three typical shadowing scenarios [25]

Shadowing Scenarios	m	β	Ω
Frequent Heavy Shadowing (FHS)	0.739	0.063	8.97×10^{-4}
Average Shadowing (AS)	10.1	0.126	0.835
Infrequent Light Shadowing (ILS)	19.4	0.158	1.29

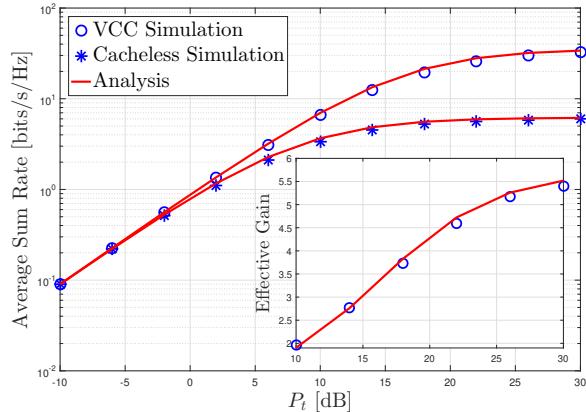


Fig. 1: Average sum rate and effective gain versus P_t for $L = 8$ and $Q_{\max} = Q'_{\max} = 8$ in FHS, where $\text{SNR}_{\text{ave}} = P_t - 9.0$ [dB].

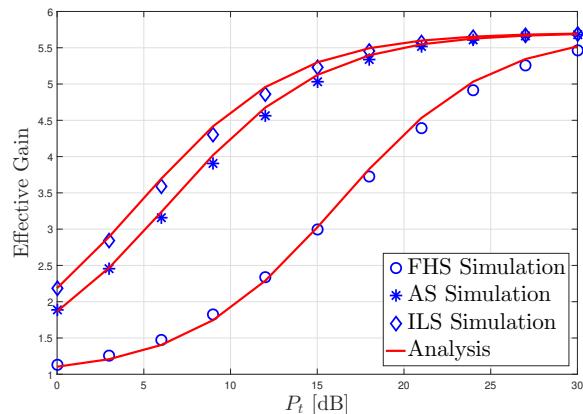


Fig. 2: Effective gain versus P_t for $L = 8$ and $Q_{\max} = Q'_{\max} = 8$, where the values of SNR_{ave} in FHS, AS and ILS are respectively $P_t - 9.0$ [dB], $P_t + 0.4$ [dB], and $P_t + 2.1$ [dB].

Remark 6. In the average sum rate expression \bar{R}_{sum} in (14), the intra-group interference term is

$$\alpha^2(Q-1)\Xi_2 = \frac{P_t(Q-1)}{GQ(2\beta+\Omega)},$$

which is independent of the CSIT error, and which underscores the robustness of MF precoding.

IV. Numerical Performance Evaluation and Validation

In this section, we numerically evaluate and plot the performance of VCC in SATCOM systems. Unless stated otherwise, the variance of the estimation error σ_e^2 is assumed to be $\frac{1}{8}$ of the AWGN power. We consider low-mobility land terminals, and therefore assume a coherence length of $T = 10^4$, which can correspond to a coherence

time of 10 ms and a coherence bandwidth of 1 MHz.⁶ Moreover, we also assume $G = 6$ and $\Theta = 12$. The choice of $G = \Lambda\gamma + 1 = 6$ reflects the achievable theoretical gain under practical file size limitations, as discussed in [14]. This value can arise from multiple combinations of the number of cache states Λ and the cache size γ , such as $\Lambda = 50$ with $\gamma = \frac{1}{10}$ or $\Lambda = 80$ with $\gamma = \frac{1}{16}$.

A. Numerical Results Under Static Channel

In this subsection, we consider the static channel model as analyzed in this paper. We evaluate the system performance under three different shadowing scenarios, with the specific Rician-shadowed channel parameters listed in Table I. According to [15, Eq. (5.3)], the average downlink SNR over a Rician-shadowed channel is $\text{SNR}_{\text{ave}} = P_t(2\beta + \Omega)$, measured from a single-antenna satellite to a user terminal. For reference, SNR_{ave} is also indicated in Figs. 1–5.

We proceed to illustrate in Figs. 1–5, the spectral efficiency gains brought about by introducing VCC in a downlink VoD-delivery satellite system. This gain (y-axis) is a *multiplicative* gain on spectral efficiency, and it incorporates the differing CSIT overhead and CSIT estimation errors of the two compared systems. We also recall that the recorded gains are over downlink systems that enjoy the same resources (same power, same number of antennas), and which are *independently* optimized with respect to the multiplexing-vs-beamforming tradeoff. Finally, let's also recall from Remark 4 that the gains do not arise from pre-caching, but that indeed these are gains on the speed of VoD delivery of data that must actually be delivered after the demands are declared.

Fig. 1 illustrates the average sum rate and the effective gain versus P_t in the most challenging shadowing environment, frequent heavy shadowing (FHS). We note that FHS is not limited to GEO satellite–terrestrial links and can also arise in LEO scenarios [28], [37], [38]. As expected, the average sum rate increases with P_t , and when we reach $P_t = 15$ dB, VCC yields a gain—over a system with the same power and antenna resources, but without VCC—of $\times 3$ in spectral efficiency (this is also often referred to as a 200% boost in spectral efficiency). What we also note is that the analytical outcomes from Theorem 1 (solid lines) are tightly validated by simulations.

Fig. 2 shows the effective gain versus P_t across various shadowing conditions: FHS, average shadowing

⁶As a practical reference, for GEO satellite–ground channels, the coherence time can reach approximately 500 ms after excluding system overheads such as gateway and user terminal processing delays [36]. For a 600 km-altitude LEO orbit, the satellite–ground round-trip time (RTT) ranges from about 4 ms to roughly twice this value, depending on the elevation angle. Given dense LEO satellite deployment and precise antenna alignment, the LEO-to-ground channel can exhibit a coherence time of 10 ms even when accounting for the RTT and system overheads. Moreover, a 1 MHz coherence bandwidth is reasonable for LEO-to-ground links under conditions with a strong LOS component (e.g., medium and light shadowing).

(AS), and infrequently light shadowing (ILS). For the three shadowing cases examined, VCC delivers sizable spectral efficiency gains relative to the cacheless system in the practical P_t regime. For instance, according to the link budget in Table II ($P_t \approx 18.1$ dB, corresponding to an EIRP⁷ of 45 dBW), the spectral efficiency is boosted by more than 400% (i.e., $\times 5$) in both AS and ILS cases. Even under the most adverse FHS condition, VCC can still *quadruple* the throughput compared to its cacheless counterpart.

EXAMPLE 1 *To clarify the above points, let us revisit the example of Fig. 2, focusing on the curve corresponding to the AS case. Consider a satellite equipped with $L = 8$ transmit antennas, and let $P_t = 9$ dB. For this setting, the plot shows a multiplicative gain of $\times 4$.*

What does this mean exactly? Suppose an operator, with the above resources ($L = 8$ transmit antennas and $P_t = 9$ dB), optimizes its system under the MF precoding assumption by balancing multiplexing and beamforming gains. After accounting for CSIT overhead, the resulting spectral efficiency is R_1 bits/s/Hz. Now, applying the VCC approach with the same resources L, P_t , and after accounting for the new CSIT overhead, yields a spectral efficiency R_{vcc} . Our results show that $R_{\text{vcc}} \approx 4 \times R_1$. To further clarify, recall Remark 4, which effectively states that the recorded gain persists even if the baseline downlink system (without VCC) is equipped with the same receiver-side cache as VCC. In that case, the baseline would indeed avoid transmitting already cached content, yet the multiplicative advantage of VCC over the baseline remains unchanged.

Fig. 3 plots how the effective gain varies with the number of transmit antennas L . This aspect is particularly relevant in satellite systems, where large-scale antenna arrays have strong potential for deployment on LEO satellites [2], [39]. The main observation here is that having more transmit antennas can in fact amplify the gain from VCC over the baseline system (with the same increased L). This observation—that higher L yields higher multiplicative gains—marks a difference from the case of terrestrial networks over Rayleigh fading channels (cf. [14, Fig. 3]), and this difference can be attributed to the fundamentally different nature of the channels and the spatial resolution constraint in SATCOM.

Fig. 4 depicts the sensitivity to CSIT imperfection. Increasing the CSIT-estimation error slightly reduces the effective gain; the reduction is marginal across the entire P_t range, highlighting that, with MF precoding, the inter-user interference term in the SINR is nearly insensitive to CSIT errors (see Remark 6). In contrast, the ZF-based results in [17] show an *increase* in effective gain

TABLE II: Link budget evaluation for a single-antenna LEO satellite at 600 km altitude [28], [40], with EIRP = 45 dBW, user terminal G/T = 5 dB/K (small VSAT), carrier frequency 20 GHz, bandwidth 36 MHz, zenith angle (90°), including gaseous absorption (0.9 dB) and rain attenuation (10 dB). The resulting P_t is approximately 18.1 dB.

Shadowing Scenarios	FHS	AS	ILS
SNR _{ave} [dB]	9.1	18.5	20.2
Effective Gain in Fig. 2	≥ 4	≥ 5	≥ 5.5

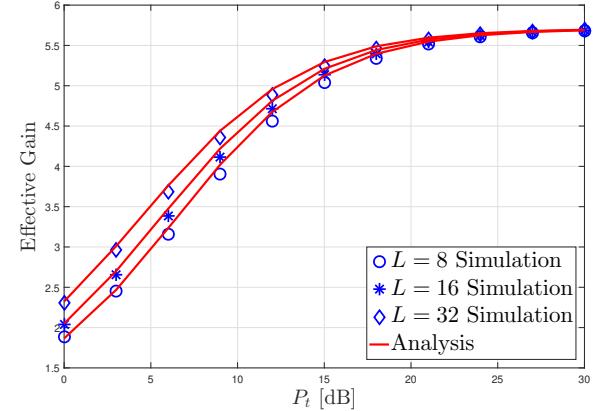


Fig. 3: Effective gain versus P_t in AS, where $Q_{\text{max}} = Q'_{\text{max}} = 8$ and $\text{SNR}_{\text{ave}} = P_t + 0.4$ [dB]. For $P_t \approx 18.1$ dB in AS, the $\text{SNR}_{\text{ave}} \approx 18.5$ dB, corresponding to gain ≥ 5 .

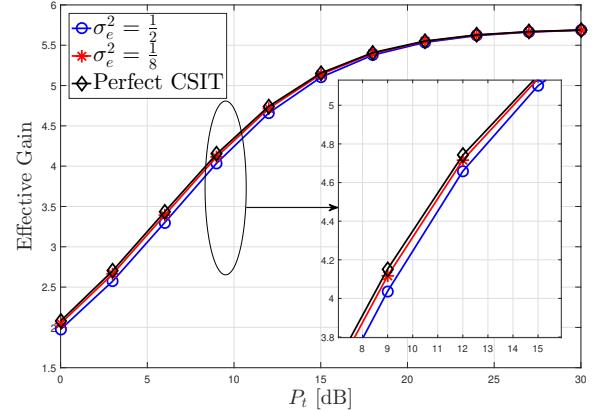


Fig. 4: Effective gain versus P_t in AS, where $L = 16$, $Q_{\text{max}} = Q'_{\text{max}} = 8$, and $\text{SNR}_{\text{ave}} = P_t + 0.4$ [dB]. For $P_t \approx 18.1$ dB in AS, the $\text{SNR}_{\text{ave}} \approx 18.5$ dB, corresponding to gain ≥ 5 .

as the CSIT error grows, because VCC's cross-group interference cancellation is CSIT-free and therefore comparatively insensitive to CSIT errors, causing the relative-gain metric to rise with the error level.

The size of the coherence block T has a direct impact on the performance of VCC. With a fixed CSI acquisition cost, a larger T reduces the relative overhead and thereby increases the effective throughput, which translates into a higher effective gain. Conversely, smaller T values amplify the impact of CSI cost and degrade the gain. Fig. 5 illustrates the effective gain achieved for different values of T . As T decreases, the gain drops markedly, while the dependence on T is less pronounced at low P_t . When T is large (e.g., 10^4), the choice of $Q_{\text{max}}, Q'_{\text{max}}$

⁷EIRP (Effective Isotropic Radiated Power) is defined as transmit power (in dBW) plus antenna gain (dBi) minus feeder losses (dB). It represents the apparent radiated power in the direction of maximum gain, as if from an ideal isotropic antenna. For instance, a satellite with 10 W transmit power (10 dBW) and 30 dBi antenna gain yields an EIRP of 40 dBW.

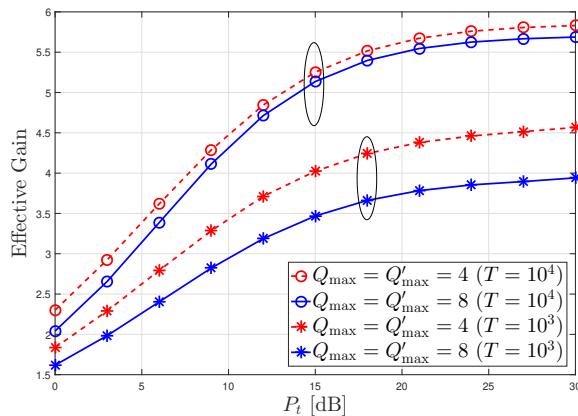


Fig. 5: Effective gain versus P_t in AS, where $L = 16$ and $\text{SNR}_{\text{ave}} = P_t + 0.4$ [dB]. For $P_t \approx 18.1$ dB in AS, the $\text{SNR}_{\text{ave}} \approx 18.5$ dB, corresponding to gains between 3.5 and 5.5.

has little influence; for small T (e.g., 10^3), increasing the *spatial-multiplexing cap*—i.e., the maximum number of simultaneous spatial streams afforded by multiple antennas—actually reduces the effective gain, while a smaller cap increases it. It is also worth noting that even with $T = 10^3$, the gains remain substantial: at $P_t = 18.1$ dB, both $Q_{\text{max}}, Q'_{\text{max}} \in \{4, 8\}$ yield an effective gain above 3.5, i.e., more than a *three-and-a-half-fold* throughput improvement over conventional MU-MIMO SATCOM.

VCC's 3×5 spectral efficiency gain not only outperforms conventional MU-MISO, but also shows that similar performance can be reached with fewer antennas, RF chains, or less accurate CSIT—opening new research directions in designing leaner, more efficient SATCOM architectures.

B. Extension to Dynamic Channel

In this subsection, we consider a dynamic channel model for SATCOM in order to demonstrate that the multiplicative spectral-efficiency gains achieved by VCC over the cacheless MU-MISO baseline are preserved under more practical channel conditions. We consider users that are *randomly and uniformly* distributed over a terrestrial coverage area with a radius D . The expectation in the effective gain definition in (9) is further averaged over the random realizations of the user spatial distribution. We do not explicitly model Doppler effects arising from the relative motion between the serving satellite and terrestrial terminals. In practical deployments, such effects can be accurately estimated and effectively compensated, since satellite orbital parameters and relative user locations are typically known or can be reliably tracked [28], [41].

Owing to the spatial distribution of users, the satellite–terrestrial channels experience rapid state variations between LOS and NLOS conditions. Following [37], [38], the dynamic transitions between LOS and NLOS conditions are modeled using a *two-state Markov process*. Specifically, under the dynamic channel model, the chan-

nel vector $\mathbf{h}_k \in \mathbb{C}^{L \times 1}$ of the k -th user is given by

$$\mathbf{h}_k = \begin{cases} \mathbf{h}_k^{\text{LOS}}, & \text{with probability } \mathcal{P}_k, \\ \mathbf{h}_k^{\text{NLOS}}, & \text{with probability } 1 - \mathcal{P}_k, \end{cases} \quad (16)$$

where \mathcal{P}_k denotes the probability of LOS propagation determined by the elevation angle ζ_k between the k -th user and the serving satellite. As suggested in [37], [42], this probability can be simply expressed as

$$\mathcal{P}_k = \exp(-\eta \cot \zeta_k), \quad (17)$$

where η is an environment-dependent parameter. The LOS and NLOS channel components, $\mathbf{h}_k^{\text{LOS}}$ and $\mathbf{h}_k^{\text{NLOS}}$ in (16), are assumed to be statistically independent and both follow the Rician-shadowed fading model proposed in [25]. Specifically, we adopt the ILS model to characterize $\mathbf{h}_k^{\text{LOS}}$ and the FHS model to characterize $\mathbf{h}_k^{\text{NLOS}}$. The corresponding channel parameters are summarized in Table I. We note that the analytical expression in (17) for modeling the LOS probability has been shown to be in good agreement with the 3GPP model [42].

Fig. 6 illustrates the effective gain achieved by VCC under static and dynamic channel models for a terrestrial coverage area with radius $D = 10$ km, corresponding to a medium-sized city. The satellite altitude H is set to 600 km, and the serving LEO satellite is located at the zenith position above the center of the terrestrial coverage area. Under this geometry,⁸ the elevation angle ζ_k of terrestrial terminals ranges from 89° to 90° . In the static case, the channel is modeled using the ILS scenario in Table I, whereas the dynamic channel follows the model in (16), where the parameter η in (17) for \mathcal{P}_k is set to 0.35, corresponding to an urban environment [37], [42]. It is observed that the effective gain under dynamic channels closely matches that of the static case across the entire P_t range, with only negligible performance degradation. This confirms that the multiplicative spectral-efficiency gains provided by VCC over the cacheless MU-MISO baseline are well preserved under dynamic satellite–terrestrial channels. It also provides justification for adopting a static and symmetric channel in the channel modeling and performance analysis of this work.

V. Conclusion and Discussions

In this paper, we analyzed the performance of VCC in state-of-art (multi-beam) SATCOM systems, focusing on scenarios characterized by low-to-moderate SNR conditions and various shadowing environments. We derived an analytical model that accurately derives the average sum rate and the effective gain, validated by simulations under different system configurations and channel conditions. Our results demonstrate that VCC provides

⁸The pathloss variation across users is negligible. For instance, the maximum pathloss occurs at the cell-edge users located at a distance of $D = 10$ km from the coverage center, whose distance to the serving satellite is $\sqrt{H^2 + D^2} \approx 600.08$ km. Consequently, for simplicity, we assume identical pathloss for all users.

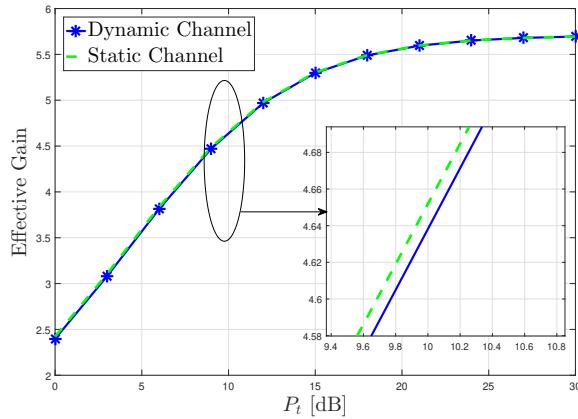


Fig. 6: Effective gain versus P_t in various channel models, where $L = 16$, $Q_{\max} = Q'_{\max} = 8$, $D = 10$ km and $\eta = 0.35$ (urban area).

substantial spectral efficiency gains even in challenging environments. In the LEO setting experiencing average and light shadowing, the gains are more pronounced and VCC can boost the spectral efficiency by a stunning factor of 3.0–5.5, even in more challenging dynamic channels. This work further positions VCC not just as a tool for extending SATCOM coverage to remote areas, but as a promising enabler for narrowing the significant spectral efficiency gap between satellite and fiber-optic communication systems for VoD.

A key factor affecting the performance gap between VCC and the traditional baseline is the CSI acquisition cost, which grows with the need to serve approximately G times more users. Our system model adopts a deliberately conservative training scheme: it omits advanced techniques like pilot reuse or coordination and assumes a fixed pilot length regardless of SNR. While this ensures robust, and perhaps conservative, performance estimates, it also indicates room for further gains through optimized training strategies—a direction we leave for future work. Moreover, incorporating time-varying channel effects induced by satellite–user relative motion, as well as exploiting distance-based spatial separation to further suppress inter-user interference in coded caching–aided SATCOM delivery, constitute interesting directions for future research.

We advocate that VCC is a novel method for accelerating satellite video delivery, doing so by rethinking the precoding structure at the PHY—not by guessing user behavior or caching entire videos in advance. Our approach is a unicast approach that retains standard VoD semantics: content is requested interactively, and our system delivers it more efficiently using a redesigned transmission method. Specifically, we modify how transmit symbols are constructed at the baseband precoding level—with no changes to the RF front-end, modulation scheme, or antenna hardware. The result is a $3 \times 5 \times$ spectral efficiency gain, enabled by layering multiple unicast streams over the same spectrum and decoding them using structured receiver logic. The modest storage

at each user terminal is not used to bypass transmission but instead supports the decoding process.

Hui Zhao
Dirk Slock, Life Fellow, IEEE
Petros Elia, Member, IEEE
Communication Systems Department, EURECOM, Sophia Antipolis, France

Appendix A: Proof of Proposition 2

Considering the transmit signal under MF precoding in (5) and the average transmit power constraint P_t , we can write the squared power control factor α^2 as

$$\alpha^2 = \frac{P_t}{\sum_{\phi \in \Psi} \sum_{\psi \in \Psi} \text{Tr}\{\mathbb{E}\{\hat{\mathbf{H}}_{\phi} \hat{\mathbf{H}}_{\psi}^H\}\}}. \quad (18)$$

For $\phi = \psi$, we have that

$$\begin{aligned} \sum_{\phi \in \Psi} \text{Tr}\{\mathbb{E}\{\hat{\mathbf{H}}_{\phi} \hat{\mathbf{H}}_{\phi}^H\}\} &= \sum_{\phi \in \Psi} \sum_{b \in [Q]} \mathbb{E}\{\|\hat{\mathbf{h}}_{\phi,b}\|^2\} \\ &= \sum_{\phi \in \Psi} \sum_{b \in [Q]} L(2\beta + \sigma_e^2 + \Omega) \\ &= GQL(2\beta + \sigma_e^2 + \Omega), \end{aligned} \quad (19)$$

where we consider the fact that $\mathbb{E}\{\|\hat{\mathbf{h}}_{\phi,b}\|^2\} = L(2\beta + \sigma_e^2 + \Omega)$ (cf. [15, Prop. 5.1]). Next, we consider the case of $\psi \neq \phi$, which yields that

$$\begin{aligned} \sum_{\phi \in \Psi} \sum_{\psi \in \Psi \setminus \{\phi\}} \text{Tr}\{\mathbb{E}\{\hat{\mathbf{H}}_{\phi} \hat{\mathbf{H}}_{\psi}^H\}\} \\ = \sum_{\phi \in \Psi} \sum_{\psi \in \Psi \setminus \{\phi\}} \text{Tr}\{\mathbb{E}\{\hat{\mathbf{H}}_{\phi}\} \mathbb{E}\{\hat{\mathbf{H}}_{\psi}^H\}\} \stackrel{(a)}{=} 0, \end{aligned} \quad (20)$$

where (a) follows from the fact that, for a scalar channel gain $h_{\psi,b}^{(\ell)}$ in Rician-shadowed fading channels, the expectations of both the scatter components and estimation errors (each being complex Gaussian with zero-mean) are zero. Thus, only the LOS component with Nakagami- m amplitude needs to be modeled, whose mean is

$$\mathbb{E}\{h_{\psi,b}^{(\ell)}\} = \frac{\Gamma(m + \frac{1}{2})}{\Gamma(m)} \sqrt{\frac{\Omega}{m}} \mathbb{E}\left\{\exp\left(j\theta_{\psi,b}^{(\ell)}\right)\right\}, \quad (21)$$

where $\theta_{\psi,b}^{(\ell)}$ is the phase of the LOS component. As $\theta_{\psi,b}^{(\ell)}$ is uniformly distributed across channel blocks, we have $\mathbb{E}\{\exp(j\theta_{\psi,b}^{(\ell)})\} = 0$. Consequently, the LOS phasor averages out, which leads to

$$\mathbb{E}\{h_{\psi,b}^{(\ell)}\} = 0. \quad (22)$$

Finally, substituting (19) and (20) into (18), we can derive the closed-form expression for α^2 in Proposition 2.

Appendix B: Proof of Proposition 3

We first derive the expression for Ξ_1 as given in (12). Based on its definition in (11), Ξ_1 can be reformulated as

$$\begin{aligned} \Xi_1 &= \mathbb{E}\{\|\mathbf{h}_{\psi,b}\|^4\} = \mathbb{E}\{\|\mathbf{h}_{\psi,b}\|^2 \|\mathbf{h}_{\psi,b}\|^2\} \\ &= \mathbb{E}\left\{\left(\sum_{\ell=1}^L |h_{\psi,b}^{(\ell)}|^2\right) \left(\sum_{\vartheta=1}^L |h_{\psi,b}^{(\vartheta)}|^2\right)\right\}, \end{aligned} \quad (23)$$

which can be further written as

$$\Xi_1 = \sum_{\ell=1}^L \mathbb{E} \left\{ |h_{\psi,b}^{(\ell)}|^4 \right\} + \sum_{\ell=1}^L \sum_{\vartheta=1, \vartheta \neq \ell}^L \mathbb{E} \left\{ |h_{\psi,b}^{(\ell)}|^2 |h_{\psi,b}^{(\vartheta)}|^2 \right\}. \quad (24)$$

For $\mathbb{E} \left\{ |h_{\psi,b}^{(\ell)}|^4 \right\}$ in (24), we have that

$$\begin{aligned} \mathbb{E} \left\{ |h_{\psi,b}^{(\ell)}|^4 \right\} &= \text{Var} \left\{ |h_{\psi,b}^{(\ell)}|^2 \right\} + \left(\mathbb{E} \left\{ |h_{\psi,b}^{(\ell)}|^2 \right\} \right)^2 \\ &\stackrel{(a)}{=} \left(4\beta^2 + 4\beta\Omega + \frac{\Omega^2}{m} \right) + (2\beta + \Omega)^2, \end{aligned} \quad (25)$$

where step (a) follows directly from [15, Prop. 5.1].

For $\mathbb{E} \left\{ |h_{\psi,b}^{(\ell)}|^2 |h_{\psi,b}^{(\vartheta)}|^2 \right\}$ in (24), we have that

$$\begin{aligned} &\mathbb{E} \left\{ |h_{\psi,b}^{(\ell)}|^2 |h_{\psi,b}^{(\vartheta)}|^2 \right\} \\ &= \mathbb{E} \left\{ |Z_{\psi,b} \exp(\jmath\theta_{\psi,b}^{(\ell)}) + \mathbf{h}'_{\psi,b}(\ell)|^2 \right. \\ &\quad \times \left. |Z_{\psi,b} \exp(\jmath\theta_{\psi,b}^{(\vartheta)}) + \mathbf{h}'_{\psi,b}(\vartheta)|^2 \right\} \\ &= \mathbb{E} \left\{ Z_{\psi,b}^4 + Z_{\psi,b}^2 |\mathbf{h}'_{\psi,b}(\ell)|^2 \right. \\ &\quad + Z_{\psi,b}^2 |\mathbf{h}'_{\psi,b}(\vartheta)|^2 + |\mathbf{h}'_{\psi,b}(\ell)|^2 |\mathbf{h}'_{\psi,b}(\vartheta)|^2 \left. \right\} \\ &= \left(1 + \frac{1}{m} \right) \Omega^2 + 4\beta\Omega + 4\beta^2, \end{aligned} \quad (26)$$

where $\mathbf{h}'_{\psi,b} \sim \mathcal{CN}(\mathbf{0}_L, 2\beta\mathbf{I}_L)$ denotes the NLOS channel vector at user $\mathbf{U}_{\psi,b}$, and $\mathbf{h}'_{\psi,b}(\ell)$ denotes the ℓ -th element. Applying (25) and (26) to (24) leads directly to (12).

Referring to the definition of Ξ_2 in (11), we first obtain

$$\begin{aligned} \Xi_2 &= \mathbb{E} \left\{ |\mathbf{h}_{\psi,b}^T \hat{\mathbf{h}}_{\psi,b'}^*|^2 \right\} = \mathbb{E} \left\{ \hat{\mathbf{h}}_{\psi,b'}^T \mathbf{h}_{\psi,b}^* \mathbf{h}_{\psi,b}^T \hat{\mathbf{h}}_{\psi,b'}^* \right\} \\ &= \mathbb{E} \left\{ \left(\sum_{\ell=1}^L \hat{h}_{\psi,b'}^{(\ell)} \left(h_{\psi,b}^{(\ell)} \right)^* \right) \left(\sum_{\vartheta=1}^L h_{\psi,b}^{(\vartheta)} \left(\hat{h}_{\psi,b'}^{(\vartheta)} \right)^* \right) \right\} \\ &= \sum_{\ell=1}^L \mathbb{E} \left\{ |h_{\psi,b}^{(\ell)}|^2 \right\} \mathbb{E} \left\{ |\hat{h}_{\psi,b'}^{(\ell)}|^2 \right\} \\ &\quad + \sum_{\ell=1}^L \sum_{\vartheta=1, \vartheta \neq \ell}^L \mathbb{E} \left\{ \hat{h}_{\psi,b'}^{(\ell)} \left(\hat{h}_{\psi,b'}^{(\vartheta)} \right)^* \right\} \mathbb{E} \left\{ h_{\psi,b}^{(\vartheta)} \left(h_{\psi,b}^{(\ell)} \right)^* \right\}. \end{aligned} \quad (27)$$

For $\mathbb{E} \left\{ h_{\psi,b}^{(\vartheta)} \left(h_{\psi,b}^{(\ell)} \right)^* \right\}$ in (27), we have

$$\begin{aligned} &\mathbb{E} \left\{ h_{\psi,b}^{(\vartheta)} \left(h_{\psi,b}^{(\ell)} \right)^* \right\} \\ &= \mathbb{E} \left\{ \left(Z_{\psi,b} \exp(\jmath\theta_{\psi,b}^{(\vartheta)}) + \mathbf{h}'_{\psi,b}(\vartheta) \right) \right. \\ &\quad \times \left. \left(Z_{\psi,b} \exp(-\jmath\theta_{\psi,b}^{(\ell)}) + (\mathbf{h}'_{\psi,b}(\ell))^* \right) \right\} \\ &\stackrel{(a)}{=} 0, \end{aligned} \quad (28)$$

where the step (a) follows from the fact that $\theta_{\psi,b}^{(\ell)}$ and $\theta_{\psi,b}^{(\vartheta)}$ i.i.d. uniformly distributed over $(0, 2\pi]$ and $\mathbf{h}'_{\psi,b} \sim \mathcal{CN}(\mathbf{0}_L, 2\beta\mathbf{I}_L)$, and which leads to

$$\mathbb{E} \left\{ \hat{h}_{\psi,b'}^{(\ell)} \left(\hat{h}_{\psi,b'}^{(\vartheta)} \right)^* \right\} \mathbb{E} \left\{ h_{\psi,b}^{(\vartheta)} \left(h_{\psi,b}^{(\ell)} \right)^* \right\} = 0. \quad (29)$$

Invoking [15, Prop. 5.1], we have $\mathbb{E} \left\{ |h_{\psi,b}^{(\ell)}|^2 \right\} = 2\beta + \Omega$ and $\mathbb{E} \left\{ |\hat{h}_{\psi,b'}^{(\ell)}|^2 \right\} = 2\beta + \sigma_e^2 + \Omega$. Combining these with (29), we finally obtain (13).

Appendix C: Proof of Theorem 1

By considering [43, Lem. 1], the average (effective) rate for user $\mathbf{U}_{\psi,b}$ can be tightly approximated as

$$\begin{aligned} \bar{R}_{\psi,b} &= \xi_{G,Q} \mathbb{E} \left\{ \log_2(1 + \text{SINR}_{\psi,b}) \right\} \\ &\approx \xi_{G,Q} \log_2 \left(1 + \frac{\alpha^2 \mathbb{E} \left\{ |\mathbf{h}_{\psi,b}^T \hat{\mathbf{h}}_{\psi,b}^*|^2 \right\}}{1 + \alpha^2 \sum_{b' \in [Q] \setminus \{b\}} \mathbb{E} \left\{ |\mathbf{h}_{\psi,b}^T \hat{\mathbf{h}}_{\psi,b'}^*|^2 \right\}} \right). \end{aligned} \quad (30)$$

In (30), it is easy to see that $\mathbb{E} \left\{ |\mathbf{h}_{\psi,b}^T \hat{\mathbf{h}}_{\psi,b'}^*|^2 \right\} = \Xi_2$ which has been derived in Proposition 3. For the term $\mathbb{E} \left\{ |\mathbf{h}_{\psi,b}^T \hat{\mathbf{h}}_{\psi,b}^*|^2 \right\}$ in (30), we have that

$$\begin{aligned} \mathbb{E} \left\{ |\mathbf{h}_{\psi,b}^T \hat{\mathbf{h}}_{\psi,b}^*|^2 \right\} &= \mathbb{E} \left\{ \left\| \mathbf{h}_{\psi,b} \right\|^2 + \mathbf{h}_{\psi,b}^T \tilde{\mathbf{h}}_{\psi,b}^* \right\}^2 \\ &\stackrel{(a)}{=} \mathbb{E} \left\{ \left\| \mathbf{h}_{\psi,b} \right\|^4 \right\} + \mathbb{E} \left\{ \tilde{\mathbf{h}}_{\psi,b}^T \mathbf{h}_{\psi,b}^* \mathbf{h}_{\psi,b}^T \tilde{\mathbf{h}}_{\psi,b}^* \right\} \\ &= \Xi_1 + \mathbb{E} \left\{ \text{Tr} \left\{ \mathbf{h}_{\psi,b}^* \mathbf{h}_{\psi,b}^T \tilde{\mathbf{h}}_{\psi,b}^* \tilde{\mathbf{h}}_{\psi,b}^T \right\} \right\} \\ &\stackrel{(b)}{=} \Xi_1 + \text{Tr} \left\{ \mathbb{E} \left\{ \mathbf{h}_{\psi,b}^* \mathbf{h}_{\psi,b}^T \right\} \mathbb{E} \left\{ \tilde{\mathbf{h}}_{\psi,b}^* \tilde{\mathbf{h}}_{\psi,b}^T \right\} \right\} \\ &= \Xi_1 + \sigma_e^2 \mathbb{E} \left\{ \left\| \mathbf{h}_{\psi,b} \right\|^2 \right\} \stackrel{(c)}{=} \Xi_1 + \sigma_e^2 L(2\beta + \Omega), \end{aligned} \quad (31)$$

where (a) follows from the fact that $\mathbb{E} \left\{ \left\| \mathbf{h}_{\psi,b} \right\|^2 \mathbf{h}_{\psi,b} \tilde{\mathbf{h}}_{\psi,b}^* \right\} = \mathbb{E} \left\{ \left\| \mathbf{h}_{\psi,b} \right\|^2 \mathbf{h}_{\psi,b} \right\} \mathbb{E} \left\{ \tilde{\mathbf{h}}_{\psi,b}^* \right\} = 0$ due to $\tilde{\mathbf{h}}_{\psi,b} \sim \mathcal{CN}(\mathbf{0}_L, \sigma_e^2 \mathbf{I}_L)$, (b) follows from the interchange between the expectation and trace operators, and (c) follows from the fact that $\mathbb{E} \left\{ \left\| \mathbf{h}_{\psi,b} \right\|^2 \right\} = L(2\beta + \Omega)$ (cf. [15, Prop. 5.1]). Substituting (31) into (30), we can derive the approximation for $\bar{R}_{\psi,b}$. Considering that $\bar{R}_{\text{sum}} = \sum_{\psi \in \Psi} \sum_{b \in [Q]} \bar{R}_{\psi,b}$, we finally have (14). Considering the effective gain defined in (9) and using the approximation for \bar{R}_{sum} , we can easily derive (15).

REFERENCES

- [1] M. Khammassi, A. Kammoun, and M.-S. Alouini, “Precoding for high-throughput satellite communication systems: A survey,” *IEEE Commun. Surv. Tutor.*, vol. 26, no. 1, pp. 80–118, 2024.
- [2] B. S. Mysore, E. Lagunas, S. Chatzinotas, and B. Ottersten, “Precoding for satellite communications: Why, how and what next?” *IEEE Commun. Lett.*, vol. 25, no. 8, pp. 2453–2457, Aug. 2021.
- [3] C. Li, H. Zhu, J. Zhang, Y. Zhu, J. Cheng, D. Bian, and G. Li, “User selection in ZF precoding multi-user satellite MIMO downlink with QoS constraints,” *IEEE Trans. Commun.*, vol. 71, no. 11, pp. 6402–6415, Nov. 2023.
- [4] M. A. Maddah-Ali and U. Niesen, “Fundamental limits of caching,” *IEEE Trans. Inf. Theory*, vol. 60, no. 5, pp. 2856–2867, May 2014.
- [5] H. Joudeh, E. Lampiris, P. Elia, and G. Caire, “Fundamental limits of wireless caching under mixed cacheable and uncachable traffic,” *IEEE Trans. Inf. Theory*, vol. 67, no. 7, pp. 4747–4767, Jul. 2021.
- [6] F. Engelmann and P. Elia, “A content-delivery protocol, exploiting the privacy benefits of coded caching,” in *Proc. Int. Symp. Modeling and Optimization in Mobile, Ad Hoc, and Wireless Networks (WiOpt)*, 2017.
- [7] E. Lampiris, J. Zhang, and P. Elia, “Cache-aided cooperation with no CSIT,” in *Proc. IEEE Int. Symp. Inf. Theory (ISIT)*, 2017, pp. 2960–2964.
- [8] S. P. Shariatpanahi *et al.*, “Multi-server coded caching,” *IEEE Trans. Inf. Theory*, vol. 62, no. 12, pp. 7253–7271, Dec. 2016.

[9] E. Lampiris and P. Elia, "Full coded caching gains for cache-less users," *IEEE Trans. Inf. Theory*, vol. 66, no. 12, pp. 7635–7651, Dec. 2020.

[10] M. J. Salehi *et al.*, "Low-complexity high-performance cyclic caching for large MISO systems," *IEEE Trans. Wireless Commun.*, vol. 21, no. 5, pp. 3263–3278, May 2022.

[11] A. Tölli, S. P. Shariatpanahi, J. Kaleva, and B. H. Khalaj, "Multi-antenna interference management for coded caching," *IEEE Trans. Wireless Commun.*, vol. 19, no. 3, pp. 2091–2106, Mar. 2020.

[12] X. Xu and M. Tao, "Modeling, analysis, and optimization of caching in multi-antenna small-cell networks," *IEEE Trans. Wireless Commun.*, vol. 18, no. 11, pp. 5454–5469, Nov. 2019.

[13] E. Lampiris and P. Elia, "Adding transmitters dramatically boosts coded-caching gains for finite file sizes," *IEEE J. Sel. Areas Commun.*, vol. 36, no. 6, pp. 1176–1188, Jun. 2018.

[14] H. Zhao, A. Bazco-Nogueras, and P. Elia, "Vector coded caching multiplicatively boosts the throughput of realistic downlink systems," *IEEE Trans. Wireless Commun.*, vol. 22, no. 4, pp. 2683–2698, Apr. 2023.

[15] H. Zhao, "High performance cache-aided downlink systems: Novel algorithms and analysis," Ph.D. dissertation, Dept. Comput. Sci. Telecommun. Electron., Sorbonne University, Paris, France, 2022.

[16] H. Zhao and P. Elia, "Vector coded caching substantially boosts MU-MIMO: Pathloss, CSI and power-allocation considerations," in *Proc. 26th Int. ITG WSA and 13th SCC*, Feb. 2023.

[17] H. Zhao and P. Elia, "Vector coded caching multiplicatively boosts MU-MIMO systems under practical considerations," *IEEE Trans. Wireless Commun.*, vol. 25, pp. 11350–11366, 2026.

[18] X. Xie, K. Huang, J. Zhang, S. Gu, and Q. Zhang, "Coded caching in satellite networks," *IEEE Trans. Commun.*, vol. 73, no. 5, pp. 3227–3239, May 2025.

[19] K. Huang, J. Huang, J. Zhang, S. Gu, and Q. Zhang, "Coded caching in satellite networks with non-uniform user distribution," *IEEE Trans. Veh. Technol.*, vol. 73, no. 5, pp. 7447–7452, May 2024.

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

[21] P. Ivaniš, J. Milojković, V. Blagojević, and S. Brkić, "Capacity analysis of hybrid satellite–terrestrial systems with selection relaying," *Entropy*, vol. 26, no. 5, 2024.

[22] G. D. Chondrogiannis *et al.*, "HARQ-assisted satellite-terrestrial communications over generalized- K fading," *IEEE Commun. Lett.*, vol. 28, no. 8, pp. 1899–1903, Aug. 2024.

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

[24] M. He, H. Zhao, X. Miao, S. Wang and G. Pan, "Secure rate-splitting multiple access transmissions in LMS systems," *IEEE Commun. Lett.*, vol. 28, no. 1, pp. 19–23, Jan. 2024.

[25] 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.

[26] V. Singh *et al.*, "STAR-RIS-enhanced NOMA-aided overlay multiuser cognitive satellite-terrestrial networks with discrete phase shifts design," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 61, no. 5, pp. 15 065–15 076, Oct. 2025.

[27] Q. T. Ngo, K. T. Phan, W. Xiang, A. Mahmood, and J. Slay, "Two-tier cache-aided full-duplex hybrid satellite–terrestrial communication networks," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 58, no. 3, pp. 1753–1765, Jun. 2022.

[28] E. Kim, I. P. Roberts, and J. G. Andrews, "Downlink analysis and evaluation of multi-beam LEO satellite communication in shadowed Rician channels," *IEEE Trans. Veh. Technol.*, vol. 73, no. 2, pp. 2061–2075, Feb. 2024.

[29] X. Miao *et al.*, "Secure satellite–vehicle communications with randomly distributed vehicles on different roads," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 60, no. 1, pp. 291–303, Feb. 2024.

[30] K. Dolas and M. R. Bhatnagar, "On performance of IRS-assisted hybrid satellite–terrestrial cooperative communication," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 59, no. 2, pp. 2020–2028, Apr. 2023.

[31] F. Rusek *et al.*, "Scaling up MIMO: Opportunities and challenges with very large arrays," *IEEE Signal Process. Mag.*, vol. 30, no. 1, pp. 40–60, Jan. 2013.

[32] H. Zhao, A. Bazco-Nogueras, and P. Elia, "Coded caching in land mobile satellite systems," in *Proc. IEEE Int. Conf. Commun. (ICC)*, May 2022.

[33] H. Zhao, D. Slock, and P. Elia, "Exact SINR analysis of matched-filter precoder," in *Proc. 25th IEEE Int. Workshop Signal Process. Adv. Wireless Commun. (SPAWC)*, Sep. 2024.

[34] G. Caire, N. Jindal, M. Kobayashi, and N. Ravindran, "Multiuser MIMO achievable rates with downlink training and channel state feedback," *IEEE Trans. Inf. Theory*, vol. 56, no. 6, pp. 2845–2866, Jun. 2010.

[35] M. K. Arti, "Performance evaluation of maximal ratio combining in shadowed-Rician fading land mobile satellite channels with estimated channel gains," *IET Commun.*, vol. 9, no. 16, pp. 2013–2022, 2015.

[36] M. Ángel Vázquez *et al.*, "Precoding, scheduling, and link adaptation in mobile interactive multibeam satellite systems," *IEEE J. Sel. Areas Commun.*, vol. 36, no. 5, pp. 971–980, May 2018.

[37] F. Qu, S. Lu, G. Xu, J. Shi, A. Nauman, and D. Niyato, "Coverage probability analysis and optimization for stochastic geometry-based multi-satellite cooperative systems under dynamic channels," *IEEE Wireless Commun. Lett.*, vol. 15, pp. 91–95, 2026.

[38] S. Yang, Y. Zhu, O. A. Dobre, G. K. Karagiannidis, and Z. Ding, "Performance analysis for NOMA-assisted LEO communications: A two-dimensional stochastic geometric approach," *IEEE Trans. Wireless Commun.*, vol. 24, no. 5, pp. 3822–3836, May 2025.

[39] L. You, K.-X. Li, J. Wang, X. Gao, X.-G. Xia, and B. Ottersten, "Massive MIMO transmission for LEO satellite communications," *IEEE J. Sel. Areas Commun.*, vol. 38, no. 8, pp. 1851–1865, Aug. 2020.

[40] 3rd Generation Partnership Project (3GPP), "Study on new radio (NR) to support non-terrestrial networks (NTN)," 3GPP, Sophia Antipolis, France, Tech. Rep. 38.811, Jul. 2020, v15.3.0.

[41] T. Li, H. Zhao, R. Zhang, and G. Pan, "Ergodic capacity analysis for terrestrial-LEO-GEO relay systems with stochastic orbit modeling," *IEEE Trans. Veh. Technol.*, vol. 74, no. 12, pp. 19 737–19 742, Dec. 2025.

[42] A. Al-Hourani and I. Guvenc, "On modeling satellite-to-ground path-loss in urban environments," *IEEE Commun. Lett.*, vol. 25, no. 3, pp. 696–700, Mar. 2021.

[43] Q. Zhang, S. Jin, K.-K. Wong, H. Zhu, and M. Matthaiou, "Power scaling of uplink massive MIMO systems with arbitrary-rank channel means," *IEEE J. Sel. Topics Signal Process.*, vol. 8, no. 5, pp. 966–981, Oct. 2014.