UC Irvine
Recent Work

Title
The Upload—Broadcast Cost Region of Three-User Linear Computation

Permalink
https://escholarship.org/uc/item/8w39n9nm

Authors
Tan, Shuo
Jafar, Syed A

Publication Date
2026-11-16

Copyright Information
This work is made available under the terms of a Creative Commons Attribution-NonCommercial-NoDerivatives License, available at https://creativecommons.org/licenses/by-nc-nd/4.0/

eScholarship.org Powered by the California Digital Library
University of California

The Upload–Broadcast Cost Region of Three-User Linear Computation

Shuo Tan, *Graduate Student Member, IEEE,* and Syed A. Jafar, *Fellow, IEEE*

***Abstract*—We characterize the upload–broadcast cost region of three-user linear computation through an initially uninformed relay. Each user has side information consisting of linear functions of a finite-field source, sends a separate noninteractive upload, and recovers its linear demand from the relay's common broadcast and its side information. For every feasible task, a single code can simultaneously attain the minimum total upload cost with unconstrained broadcast and the minimum broadcast cost achievable by a fully informed server. The closed cost region is therefore rectangular, even for nonlinear fixed-length zero-error codes of arbitrary blocklength. We give an explicit rank formula for the minimum upload. Over the original field, including the binary field, the minimum linear blocklength for attaining both minima is one if both are integers and two otherwise. As an arbitrary minimum-upload scheme need not preserve the fully informed broadcast optimum, we coordinate upload choices across components using their linear computation broadcast profiles and modify uploads along directed demand cycles when necessary.**

***Index Terms*—Distributed function computation, linear computation broadcast, relay networks, zero-error coding.**

I. Introduction

In distributed computation, the required communication depends on the functions to be computed and the information available at each user. In machine learning and data analytics, jointly processing data held at different devices requires information to be exchanged [1], [2]. Related communication problems also arise in coded caching [3], [4], private information retrieval [5], [6], coded distributed computing [1], data shuffling [7], distributed storage repair [8], and secure aggregation for federated learning [9].

Functional compression exploits the fact that recovering a function can require less communication than recovering the source itself [10], [11], [12]. Network function computation examines how the distribution of information and the opportunities to combine it affect the communication required for a target function [13], [14], [15]. In collaborative computations, overlapping data and previously computed results can create dependencies among the users' demands and side information. A common broadcast can exploit these dependencies to serve several users: each combines the same message with its own side information to recover its demand. Index coding captures this possibility for demands and side information specified by subsets of source messages [16], [17]; formulations with coded side information allow more general linear dependencies [18], [19].

Linear computation broadcast (LCBC) specifies each user's demand and side information as arbitrary linear functions of a common source [20], [21], [22], [23], [24]. In the finite-field setting, a fully informed server observes the source and sends a common broadcast from which every user must recover its demand using its own side information. The objective is to minimize the broadcast cost per computation instance for the prescribed demands and side information. This cost depends on their joint linear structure. The minimum broadcast cost has been characterized for two and three users [20], [21]. In the three-user case, the characterization holds over every finite field for arbitrary linear demands and side information, with a converse valid for nonlinear codes. For more users, generic symmetric capacity results and general achievable schemes are available [22], [23].

These LCBC results take the server's information as given, as does the two-user capacity characterization with a partially informed server [25]. When the relay obtains its information through user uploads, the upload and broadcast costs must be considered jointly.

Related communication problems have been studied in relay networks and coded distributed computing. Ravi and Dey [26] give a multiletter characterization of the zero-error rate region for two users computing an arbitrary common function through a relay, with rates measured by expected variable-length code lengths. Jointly optimal upload and broadcast loads have also been characterized for MapReduce computation through a relay under storage constraints [27] and joint storage and computation constraints [28]. For coded shuffling over a shared broadcast channel with prescribed data placement and output assignments, coding schemes and sufficient conditions for optimality have been established [29]. In these MapReduce models, each user must recover all intermediate values required by its assigned outputs, using locally computed intermediate values as side information.

In this paper, we study this joint communication problem for three users with arbitrary prescribed linear side information and demands. We assume that the users' side information jointly determines all demands. An initially uninformed relay receives a separate, noninteractive upload from each user, formed only from that user's side information. The relay forms a common broadcast from these uploads, and each user recovers its demand using the broadcast and its side information; see Figure 1.

An upload scheme of minimum total cost may nevertheless require more broadcast communication than a fully informed server. Let U^* be the minimum total upload cost when

This work was supported in part by the Office of Naval Research (ONR) under Grant N000142412618, in part by the Army Research Office (ARO) under Grant W911NF2510274, in part by the National Science Foundation (NSF) under Grant CCF-2221379.

The authors are with the Nhu Department of Electrical Engineering and Computer Science, University of California, Irvine, Irvine, CA 92697 USA (e-mail: shuot8@uci.edu; syed@uci.edu).

broadcast is unconstrained, and let D^* be the minimum broadcast cost of a fully informed server, as determined in [21]. Both costs are measured in q-ary units per computation instance. They are separate lower bounds on the costs of every upload–broadcast code. We ask whether a single code can attain both bounds simultaneously. When U^* and D^* are both integers, can this pair be attained by a scalar linear code over the original field $\mathbb{F}_q$?

A. Main Results

For every feasible three-user linear task, we show that a single code attains both bounds. Consequently, Theorem 1 gives the closed upload–broadcast cost region $[U^*, \infty) \times [D^*, \infty)$ under fixed-length zero-error coding, allowing nonlinear codes of arbitrary blocklength. We give an explicit rank formula for U^*. Whenever U^* and D^* are both integers, a scalar $\mathbb{F}_q$-linear upload–broadcast code attains (U^*, D^*) over the original field, including $\mathbb{F}_2$. Otherwise, the minimum $\mathbb{F}_q$-linear blocklength for attaining this pair is two. The binary example below shows that two instances can also be necessary with nonlinear coding.

B. Why Uploads Must Be Coordinated

Consider the following task, where a, b, c, u, v are arbitrary bits and addition is over $\mathbb{F}_2$:

User	Side information	Demand
1	(a, u)	$b + c$
2	(b, v)	$a + c$
3	$(c, u + v)$	$a + b + u$

. (1)

Theorem 3 and (16) give $U^* = 4$ and $D^* = 3/2$; the upload parameters are evaluated in Section III-A4, and the broadcast calculation appears in Appendix C-C. The scalar uploads $M_1 = (a, u),\ M_2 = b,\ M_3 = c$ attain the minimum upload cost, but repeating this tuple requires broadcast cost 2 per instance, even with joint nonlinear broadcasting across repetitions. For L independent uniform instances, let S be any zero-error broadcast of the repeated relay input, with attainable alphabet $\mathcal{S}_L$. By Lemma 8, Users 2 and 3 in the induced task have side information b^L and c^L, respectively. Their decoders recover $a^L + c^L$ and $a^L + b^L + u^L$, so (S, b^L, c^L) determines (a^L, u^L). Hence

$$\begin{aligned} \log_2 |\mathcal{S}_L| &\geq H(S \mid b^L, c^L) \\ &\geq H(a^L, u^L \mid b^L, c^L) = 2L. \end{aligned} \tag{2}$$

Broadcasting $(a+b+c, a+b+c+u)$ on each instance attains equality.

To attain both minima, use different scalar upload tuples on the two instances:

$$M_1 = (a_1, u_1, a_2), \quad M_2 = (b_1, b_2, v_2), \quad M_3 = (c_1, c_2).$$

With $t_\ell = a_\ell + b_\ell + c_\ell$, the relay broadcasts

$$C_1 = t_1 + t_2, \qquad C_2 = t_1 + u_1, \qquad C_3 = t_2 + v_2. \tag{3}$$

User 1 obtains $t_1 = C_2 + u_1$ and then $t_2 = C_1 + t_1$; User 2 obtains $t_2 = C_3 + v_2$ and then $t_1 = C_1 + t_2$. They recover their demands by adding a_ℓ and b_ℓ, respectively. User 3 obtains its demands as $C_2 + c_1$ and $C_3 + c_2 + (u_2 + v_2)$. The cost pair is therefore $(8/2, 3/2) = (U^*, D^*)$.

Finally, a scalar binary broadcast with cost at most $3/2$ has at most $\lfloor 2^{3/2} \rfloor = 2$ attainable messages and would cost at most one bit, contradicting the fully informed lower bound $D^* = 3/2$. Thus two instances are necessary even for nonlinear codes.

Section II states the model and main theorem. Section III proves the minimum-upload formula, and Sections IV and V develop and coordinate the relay choices. Section VI completes the code construction and the proof.

II. Problem Formulation and Main Result

A. Notation

Let $\mathbb{F}_q$ be the finite field of prime-power order q. Unless stated otherwise, dimensions are over $\mathbb{F}_q$ and logarithms have base q. Write $[n] = \{1, \dots, n\}$, $\mathbb{R}_+ = [0, \infty)$, and $a_+ = \max\{a, 0\}$. Vector inequalities hold coordinatewise. Let $\mathbf{e}_j \in \mathbb{F}_q^n$ and $\vec{e}_j \in \mathbb{R}^n$ denote standard basis vectors over the finite field and over the reals, respectively; $\mathbf{1}_s$ is the all-ones vector of length s. We use im, $\ker$, rk, $\langle\cdot\rangle$, and $\oplus$ for image, kernel, rank, linear span, and direct sum, respectively. For a linear map f, $f|_S$ denotes its restriction to a subspace S of its domain. For matrices, $\langle V \rangle$ denotes the column span, $[V, W]$ is column concatenation, and $\mathrm{rk}(V \mid W) = \mathrm{rk}[V, W] - \mathrm{rk}\, W$. For subspaces B, C of a finite-dimensional coefficient space F, define the conditional rank $\mathrm{rk}(C \mid B) = \dim(B + C) - \dim B$. In coordinate constructions, we identify scalar linear forms with their coefficient vectors and use the shorthand $\langle x^\mathsf{T} C \rangle \triangleq \langle C \rangle$ for any coefficient matrix C.

B. Source, Side Information, and Demands

We follow the notation of three-user LCBC [21]. In one computation instance, the source is an arbitrary $x \in \mathbb{F}_q^d$, where d is the source dimension per instance. User $i \in [3]$ has side information $w_i' = x^\mathsf{T} V_i'$ and demand $w_i = x^\mathsf{T} V_i$. Remove redundant columns as in [21, Sec. VIII-A]: choose a basis V_i' for the side information and demand columns V_i that complete it to a basis of the joint span. Given the side information, the original and retained demands determine each other. This preserves all attainable cost pairs at every blocklength, and gives $V_i' \in \mathbb{F}_q^{d \times m_i'}$, $V_i \in \mathbb{F}_q^{d \times m_i}$, and $\mathrm{rk}[V_i'\ V_i] = m_i' + m_i$. We work with the subspaces

$$A_i \triangleq \langle V_i' \rangle, \qquad U_i \triangleq \langle [V_i'\ V_i] \rangle. \tag{4}$$

Here A_i is User i's side-information coefficient space, and U_i is its *signal space*, the coefficient space spanned jointly by its side information and demand. In [21], U_i denotes the concatenated matrix; here it denotes its column space. Thus $m_i' = \dim A_i$ and $m_i = \mathrm{rk}(V_i \mid V_i') = \dim U_i - \dim A_i$. Given w_i', recovering w_i is equivalent to recovering all linear combinations of the source symbols with coefficient vectors in U_i; the task therefore depends only on (A_i, U_i). Assume throughout that the upload task is feasible: $U_i \subseteq \sum_j A_j$ for every i. Restricting the coefficient space to $F \triangleq \sum_i A_i = \sum_i U_i$ and choosing coordinates, we take $d = \dim F$ without

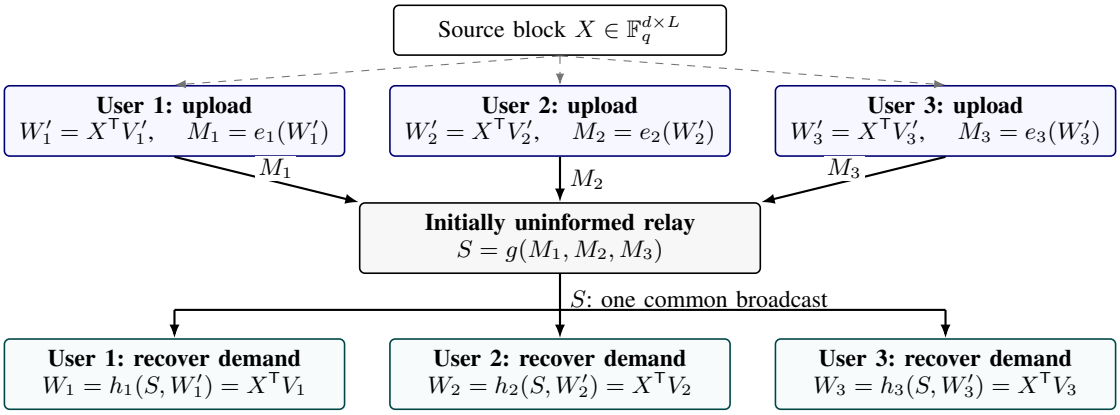


Fig. 1. Three-user upload–broadcast computation. The users send separate noninteractive uploads to a relay with no initial source information. Each user decodes from the same common broadcast and its side information W_i'. Solid lines denote communication; dashed arrows denote prescribed source mappings. Encoding and decoding may be nonlinear. The blocklength L counts computation instances.

loss of generality. Thus d denotes the source dimension of the current task and may decrease under later reductions. Denote the task by $\mathsf{M} = (A_i \subseteq U_i \subseteq F)_{i=1}^3$. Task direct sums take direct sums of the ambient spaces and of each user's side-information and signal spaces; $\mathsf{M}^{\oplus L}$ denotes L copies.

C. Coding Scheme and Cost Region

A block of $L \geq 1$ instances is $X = [x(1), \dots, x(L)] \in \mathbb{F}_q^{d\times L}$, with side information $W_i' = X^{\mathsf{T}}V_i'$ and demands $W_i = X^{\mathsf{T}}V_i$. The rows of $W_i' \in \mathbb{F}_q^{L\times m_i'}$ and $W_i \in \mathbb{F}_q^{L\times m_i}$ index computation instances. The blocklength L, called *batch size* in [21], counts computation instances. A code is *scalar* for $L = 1$ and *vector* for $L > 1$; either may send multiple symbols per instance.

The users send separate noninteractive uploads over noiseless links to an initially uninformed relay, which then sends one common broadcast. Uploads are received individually, not through an additive multiple-access channel. A fixed-length zero-error code consists of finite message alphabets $\mathcal{M}_1, \mathcal{M}_2, \mathcal{M}_3, \mathcal{S}$ and deterministic maps e_i, g, h_i, chosen independently of the source realization, satisfying

$$\begin{aligned} M_i &= e_i(W_i'), \qquad S = g(M_1, M_2, M_3), \\ W_i &= h_i(S, W_i'), \qquad i \in [3], \end{aligned} \tag{5}$$

for every $X \in \mathbb{F}_q^{d\times L}$. Write $M = (M_1, M_2, M_3)$ for the upload tuple. The maps may be nonlinear and couple all L instances; no source distribution is assumed.

Communication is charged by attainable message alphabets. Full column rank of V_i' makes every value of the local side-information block W_i' attainable, so the upload alphabet is $\operatorname{im} e_i$. Let $\mathcal{S}_L \triangleq \{S(X) : X \in \mathbb{F}_q^{d\times L}\}$ be the attainable broadcast alphabet. The normalized costs, in q-ary units per instance, are

$$R_{\mathrm{up}} = \frac{1}{L}\sum_{i=1}^{3} \log_q |\operatorname{im} e_i|, \qquad R_{\mathrm{bc}} = \frac{1}{L}\log_q |\mathcal{S}_L|. \tag{6}$$

Let $\mathcal{R}_L \subseteq \mathbb{R}_+^2$ be the upward-closed set of budgets (U, D) for which a blocklength-L code satisfies $R_{\mathrm{up}} \leq U$ and $R_{\mathrm{bc}} \leq D$. Define the closed cost region

$$\overline{\mathcal{R}} \triangleq \overline{\bigcup_{L\geq 1} \mathcal{R}_L}, \tag{7}$$

with closure in $\mathbb{R}^2$, and the optimal total upload cost with unconstrained broadcast as

$$U^* \triangleq \inf\{U : \ (U, D) \in \overline{\mathcal{R}} \text{ for some } D \geq 0\}. \tag{8}$$

With unconstrained broadcast, upload encoders are feasible if and only if

$$W_i = \widetilde{h}_i\big(W_i', (M_j)_{j\neq i}\big), \qquad i \in [3], \tag{9}$$

for some functions $\widetilde{h}_i$ and every source block (necessity follows by setting $S = g(M)$; sufficiency by broadcasting M).

D. Entropy Interpretation

For converse arguments only, take X uniform on $\mathbb{F}_q^{d\times L}$. The zero-error constraints imply

$$\begin{aligned} H(M_i \mid W_i') = H(W_i \mid S, W_i') = 0, \quad i \in [3], \\ H(S \mid M_1, M_2, M_3) = 0. \end{aligned} \tag{10}$$

The alphabet costs obey

$$LR_{\mathrm{up}} \geq \sum_{i=1}^{3} H(M_i), \qquad LR_{\mathrm{bc}} \geq H(S). \tag{11}$$

These inequalities can be strict: nonlinear messages need not be uniform on their attainable alphabets.

For $B, C \leq F$, write $X^{\mathsf{T}}B$ for the linear combinations of the source symbols specified by a fixed basis of B. The rank–entropy identities are

$$\begin{aligned} H(X^{\mathsf{T}}C \mid X^{\mathsf{T}}B) &= L \operatorname{rk}(C \mid B), \\ H(X^{\mathsf{T}}B) &= L \dim B, \\ I(X^{\mathsf{T}}B; X^{\mathsf{T}}C) &= L \dim(B \cap C). \end{aligned} \tag{12}$$

In particular, $H(W_i \mid W_i') = Lm_i$. The subspace structure remains essential: even the entropies of all subsets of the original side-information and demand variables need not determine the broadcast optimum [21].

E. Fully Informed Broadcast Optimum

Removing uploads and giving the server the source block X yields the three-user linear computation broadcast (LCBC) problem with side information W_i' and demands W_i. Its optimal cost is

$$D^* := \inf_{L,\varphi} \frac{1}{L} \log_q |\operatorname{im} \varphi|, \tag{13}$$

where the infimum is over all $L \geq 1$ and finite-alphabet broadcasts $S = \varphi(X)$ enabling each user i to recover W_i from (S, W_i') for every X. We write $\mathcal{R}_L(\mathsf{M})$, $U^*(\mathsf{M})$, and $D^*(\mathsf{M})$ when the task dependence is needed, and otherwise suppress it.

For distinct $i, j, k \in [3]$, set $U_{ij} = U_i \cap U_j$, $U_{123} = U_1 \cap U_2 \cap U_3$, and $U_{i(j,k)} = U_i \cap (U_j + U_k)$. For each permutation (i, j, k) of $(1, 2, 3)$, define

$$F_{ijk} = \sum_{a=1}^{3} m_a - \operatorname{rk}(U_{ij} \mid A_j) - \operatorname{rk}\big(U_{k(i,j)} \mid A_k\big). \tag{14}$$

For $i, \ell \in [3]$ and $\{i, j, k\} = [3]$, define

$$\begin{aligned} Q_{i\ell} = \sum_{a=1}^{3} m_a - \frac{1}{2}\Big[&\operatorname{rk}(U_{123} \mid A_\ell) + \operatorname{rk}(U_{ij} + U_{ik} \mid A_i) \\ &+ \operatorname{rk}\big(U_{j(i,k)} \mid A_j\big) + \operatorname{rk}\big(U_{k(i,j)} \mid A_k\big)\Big]. \end{aligned} \tag{15}$$

The six F_{ijk} values and nine $Q_{i\ell}$ values form the *LCBC broadcast profile*. We retain the conditioning index ℓ for the entrywise direct-sum aggregation in Section IV-B. By [21, Thm. 1],

$$D^* = \max\left\{ \max_{(i,j,k)} F_{ijk},\ \max_{i,\ell \in [3]} Q_{i\ell} \right\}. \tag{16}$$

Under the alphabet-cost convention in (6), the converse follows from the LCBC entropy bound and (11), while vector linear codes attain the stated value.

A fully informed server can simulate any relay code, so $D \geq D^*$ for every $(U, D) \in \overline{\mathcal{R}}$.

F. Main Result and Proof Overview

Theorem 1 (Optimal region and finite-blocklength achievability)**.** *For every feasible three-user linear task over $\mathbb{F}_q$, the closed cost region under possibly nonlinear fixed-length zero-error coding is*

$$\mathcal{R}_2 = \overline{\mathcal{R}} = [U^*, \infty) \times [D^*, \infty), \tag{17}$$

where U^* is given by Theorem 3 and D^* is the fully informed broadcast optimum in (16)*. Let $L_{\min}^{\mathrm{lin}}$ denote the minimum blocklength of an $\mathbb{F}_q$-linear upload–broadcast code attaining (U^*, D^*). Then*

$$L_{\min}^{\mathrm{lin}} = \begin{cases} 1, & U^*, D^* \in \mathbb{Z}, \\ 2, & \textit{otherwise}. \end{cases} \tag{18}$$

For some binary tasks, attaining both minima requires at least two instances even when nonlinear codes are allowed.

Section VI-C proves the theorem. Larger blocklengths and closure do not enlarge the budget region.

The proof has three parts.

Upload converse. Section III separates the task into omniscience and residual components. With unconstrained broadcast, the upload subproblem for the omniscience component is a three-terminal communication-for-omniscience problem [30], [31]. Its classical rate constraints reduce to the three deficit inequalities in (47). A conditional-entropy bound shows that nonlinear mixing across the two components cannot reduce the sum of their minimum upload costs, yielding U^*.

Relay selection. Section IV reduces broadcasting from a fixed relay input to an LCBC problem and gives sufficient upload-optimal relay choices for each component. The allocation argument uses only the resulting changes in the fifteen entries of the LCBC broadcast profile. Section V coordinates these changes under the full task's broadcast budget. When the initial choices fail, directed demand cycles supply a way to modify the uploads without increasing their cost.

Coding over the original field. Section VI first constructs a scalar code when the minimum upload and the broadcast budget are integers. The binary construction uses additional component choices and a separation property of the induced side-information spaces. Applying the scalar result to two copies of the original task completes the proof for arbitrary half-integer minima.

III. Minimum Upload

We first reduce the task to the direct sum of an omniscience component and a residual component. Their minimum upload costs add. The key converse step establishes this additivity even when nonlinear encoders mix the two components, which are independent under a uniform source.

A. Reduced Task and Upload Formula

We decompose the side information into functions known only to one user, functions known to two users, and three-line blocks with local observations $v, r, r - v$. We remove unused directions and identify the block coordinates that remain in the demands after private cancellation. These steps define the rank parameters in the upload formula and separate a component that all users must recover from a residual component with general functional demands. The parameters $d_\bullet$ count dimensions per instance; U_{res} and U_{om} denote the corresponding minimum upload costs in q-ary units per instance. Appendix A proves that the reductions preserve every $\mathcal{R}_L$, including for nonlinear codes, and that the parameters do not depend on the chosen coordinates.

1) Side-information coordinates: Define the active pair dimensions in the original task by

$$d_{\mathrm{pair},i} \triangleq \dim \frac{U_i \cap A_j \cap A_k}{A_1 \cap A_2 \cap A_3}, \qquad \{i, j, k\} = [3]. \tag{19}$$

Thus $d_{\mathrm{pair},i}$ counts the dimensions required by User i and known to both other users, after removing side information common to all three.

Quotient F and all task spaces by $A_1 \cap A_2 \cap A_3$, and retain the notation F, A_i, U_i for the resulting quotient spaces. We

choose a finite-dimensional vector space $\mathcal{T}$, a two-dimensional space E, and coordinates in which

$$F = \bigoplus_{i=1}^{3} \mathcal{P}_i \oplus \bigoplus_{i=1}^{3} \bar{K}_i \oplus (\mathcal{T} \otimes E), \tag{20}$$

$$A_i = \mathcal{P}_i \oplus \bar{K}_j \oplus \bar{K}_k \oplus (\mathcal{T} \otimes \ell_i), \tag{21}$$

where $\{i,j,k\} = [3]$, $\mathcal{P}_i$ is private to User i, and $\bar{K}_i$ is known exactly to the other two users. The space E has basis $\mathbf{v}, \mathbf{r}$ and the lines $\ell_1 = \langle \mathbf{v} \rangle, \ell_2 = \langle \mathbf{r} \rangle, \ell_3 = \langle \mathbf{r} - \mathbf{v} \rangle$. For $T \le \mathcal{T}$, $t \in T$, and $\mathbf{z} \in E$, write $t\mathbf{z} = t \otimes \mathbf{z}$ and $T\mathbf{z} = T \otimes \langle \mathbf{z} \rangle$. For nonzero $t \in \mathcal{T}$, the two-dimensional coefficient space $\langle t \rangle \otimes E$, with local lines $\langle t \rangle \otimes \ell_i$, is a *three-line block*. Its two source coordinates v, r are observed as $v, r, r - v$ by Users 1, 2, and 3, respectively. A basis of $\mathcal{T}$ decomposes $\mathcal{T} \otimes E$ into $\dim \mathcal{T}$ such blocks, occupying $2 \dim \mathcal{T}$ source dimensions.

Fix $(\mathbf{z}_1, \mathbf{z}_2, \mathbf{z}_3) = (\mathbf{r}, \mathbf{v}, \mathbf{v})$ and identify

$$F/A_i \cong \bar{K}_i \oplus \mathcal{P}_j \oplus \mathcal{P}_k \oplus \mathcal{T}$$

through $(\kappa, p_j, p_k, t) \mapsto \kappa + p_j + p_k + t\mathbf{z}_i + A_i$, where $\{i,j,k\} = [3]$. All demand quotients below use these coordinates.

Proposition 2 (Active reduction)**.** *After an invertible change of quotient coordinates preserving each A_i, the pairwise-known block $\bar{K}_i$ decomposes as $\bar{K}_i = K_i^{\mathrm{act}} \oplus K_i^{\mathrm{pas}}$, $\dim K_i^{\mathrm{act}} = d_{\mathrm{pair},i}$, and, for $\{i,j,k\} = [3]$, $U_i/A_i = K_i^{\mathrm{act}} \oplus \mathcal{Q}_i$, $\mathcal{Q}_i \le (\mathcal{P}_j \oplus \mathcal{P}_k) \oplus \mathcal{T}$. The passive blocks K_i^{pas}, together with any private directions unused by all remote demands, can be removed without changing any $\mathcal{R}_L$.*

The proof is given in Appendix A-A.

2) Private parts and pure demands: In the reduced coordinates of Proposition 2, let π_B denote projection onto the coordinate summand B. For $\{i,j,k\} = [3]$, define

$$\mathcal{D}_i = \mathcal{Q}_i \cap \mathcal{T}, \qquad L_i = \pi_{\mathcal{P}_j \oplus \mathcal{P}_k} \mathcal{Q}_i. \tag{22}$$

The vectors in $\mathcal{D}_i$ specify the *pure demands* in $\mathcal{Q}_i$, which have no private part, and L_i collects the private parts. Define the linear map $\phi_i : L_i \to \mathcal{T}/\mathcal{D}_i$ by $\phi_i(p) = t + \mathcal{D}_i$ whenever $p + t \in \mathcal{Q}_i$. This class is unique because any two choices of t differ by an element of $\mathcal{D}_i$.

After deleting private directions unused by remote demands, the retained spaces satisfy

$$\mathcal{P}_j = \sum_{i \ne j} \pi_{\mathcal{P}_j} L_i, \qquad d_{\mathrm{priv}} \triangleq \sum_{j=1}^{3} \dim \mathcal{P}_j. \tag{23}$$

Set

$$C = \mathcal{D}_1 \cap \mathcal{D}_2 \cap \mathcal{D}_3, \qquad d_{\mathrm{com}} = \dim C, \tag{24}$$

$$\mathcal{T}_\varnothing = \mathcal{T}/(\mathcal{D}_1 + \mathcal{D}_2 + \mathcal{D}_3). \tag{25}$$

Every user recovers all linear combinations with coefficient vectors in $C \otimes E$, since $C \otimes E \subseteq U_i$ for every i. In a basis of $\mathcal{T}$ adapted to C, this space consists of d_{com} three-line blocks. The quotient $\mathcal{T}_\varnothing$ removes $\mathcal{D}_1 + \mathcal{D}_2 + \mathcal{D}_3$. The $\mathcal{D}_i$ may involve linear combinations across blocks and need not be coordinate subspaces in a common basis of $\mathcal{T}$.

3) Blocks after private cancellation: Adding locally known block coordinates to private coordinates preserves every A_i. We determine which demand components remain after these changes, modulo $\mathcal{D}_1 + \mathcal{D}_2 + \mathcal{D}_3$.

For $\xi = (\xi_1, \xi_2, \xi_3) \in \bigoplus_i L_i$, write $\xi_{ij} = \pi_{\mathcal{P}_j} \xi_i$. Define linear maps $B : \bigoplus_i L_i \to \bigoplus_i \mathcal{P}_i$ and $G : \bigoplus_i L_i \to \mathcal{T}_\varnothing$ by

$$\begin{aligned} B\xi &= (\xi_{21} + \xi_{31},\ \xi_{12} + \xi_{32},\ \xi_{13} - \xi_{23}), \\ G\xi &= \sum_i \phi_i(\xi_i) \quad \text{in } \mathcal{T}_\varnothing. \end{aligned} \tag{26}$$

Here G uses the canonical projections $\mathcal{T}/\mathcal{D}_i \to \mathcal{T}_\varnothing$. The signs in B follow from the lines ℓ_i and quotient bases fixed above.

These changes replace G by $G + hB$, with arbitrary linear $h : \bigoplus_i \mathcal{P}_i \to \mathcal{T}_\varnothing$. Hence $G_0 \triangleq G|_{\ker B}$ is unchanged. Conversely, any component of G vanishing on $\ker B$ factors through B and can be canceled.

By Lemma 18, a subspace $\widetilde{N}_\varnothing \le \mathcal{T}$ lifting a complement of $\operatorname{im} G_0$ in $\mathcal{T}_\varnothing$ gives a passive summand $\widetilde{N}_\varnothing \otimes E$ after a suitable coordinate change. Deleting this summand preserves every $\mathcal{R}_L$. Choose a basis of the retained subspace of $\mathcal{T}$ adapted to C. The number of three-line blocks remaining after separating $C \otimes E$ is therefore

$$\begin{aligned} d_{\mathrm{tri}} &\triangleq \dim \frac{\mathcal{D}_1 + \mathcal{D}_2 + \mathcal{D}_3}{C} + \operatorname{rk} G_0, \\ \operatorname{rk} G_0 &= \operatorname{rk} \begin{bmatrix} B \\ G \end{bmatrix} - \operatorname{rk} B. \end{aligned} \tag{27}$$

The first term is the dimension of the pure-demand span modulo C. The second counts the additional dimensions in $\mathcal{T}_\varnothing$ that cannot be removed by these changes of private coordinates.

4) Minimum upload and an example:

Theorem 3 (Minimum upload)**.** *For every feasible three-user linear task over $\mathbb{F}_q$, the minimum total upload cost under possibly nonlinear fixed-length zero-error coding is*

$$U^* = U_{\mathrm{res}} + U_{\mathrm{om}}, \qquad U_{\mathrm{res}} = d_{\mathrm{priv}} + d_{\mathrm{tri}}, \tag{28}$$

where $d_{\mathrm{om},i} \triangleq d_{\mathrm{com}} + d_{\mathrm{pair},i}$ is User i's missing dimension in the omniscience component and

$$U_{\mathrm{om}} = \max\left\{ \max_i d_{\mathrm{om},i},\ \frac{1}{2} \sum_i d_{\mathrm{om},i} \right\}. \tag{29}$$

Linear uploads over $\mathbb{F}_q$ attain this value with one instance when U_{om} is integral and with two otherwise.

Sections III-D and III-E prove the converse and give a matching construction. In (28), U_{res} pays for the retained private dimensions and residual three-line blocks, while U_{om} depends on the three omniscience deficits $d_{\mathrm{om},i}$. The matching linear construction sends LU^* symbols over L instances. The blocklength claim in Theorem 3 concerns upload alone; later sections also require an optimal broadcast.

Evaluation for the five-bit task. The task (1) has private symbols a, b, c, one three-line block with side information $u, v, u + v$, and no active pair or pure demands. Thus $d_{\mathrm{priv}} = 3$ and $d_{\mathrm{com}} = d_{\mathrm{pair},i} = 0$. Using $b + c, a + c, a + b$ as bases of L_1, L_2, L_3, respectively, (26) becomes

$$B\xi = (\xi_2 + \xi_3, \xi_1 + \xi_3, \xi_1 + \xi_2), \qquad G\xi = \xi_3.$$

Thus $\ker B = \langle(1,1,1)\rangle$ and $\operatorname{rk} G_0 = 1$, so $d_{\mathrm{tri}} = 1$. Hence $U_{\mathrm{res}} = 4$, $U_{\mathrm{om}} = 0$, and $U^* = 4$.

5) Omniscience and residual components: Apply Lemma 18 and reuse $\mathcal{T}$ for the retained subspace, of dimension $d_{\mathrm{com}}+d_{\mathrm{tri}}$. Let $W_\varnothing \le \mathcal{T}$ lift $\operatorname{im} G_0$, so that $\dim W_\varnothing = \operatorname{rk} G_0$. Choose $\mathcal{T} = C \oplus \mathcal{T}_{\mathrm{res}}$ with $W_\varnothing \le \mathcal{T}_{\mathrm{res}}$. Since $C \subseteq \mathcal{D}_i$, take linear representatives of ϕ_i in $\mathcal{T}_{\mathrm{res}}$. The coefficient space $C \otimes E$ and the active pair spaces form the omniscience summand

$$\begin{aligned} F_{\mathrm{om}} &= (C \otimes E) \oplus \bigoplus_{i=1}^{3} K_i^{\mathrm{act}}, \\ U_i \cap F_{\mathrm{om}} &= F_{\mathrm{om}}, \qquad i \in [3], \end{aligned} \tag{30}$$

with side-information spaces

$$A_{i,\mathrm{om}} = (C \otimes \ell_i) \oplus K_j^{\mathrm{act}} \oplus K_k^{\mathrm{act}}. \tag{31}$$

The residual coefficient space is

$$F_{\mathrm{res}} = \bigoplus_i \mathcal{P}_i \oplus (\mathcal{T}_{\mathrm{res}} \otimes E), \qquad \dim \mathcal{T}_{\mathrm{res}} = d_{\mathrm{tri}}.$$

For $\beta \in \{\mathrm{om}, \mathrm{res}\}$, write $A_{i,\beta} = A_i \cap F_\beta$, $U_{i,\beta} = U_i \cap F_\beta$, and $\mathsf{M}_\beta = (A_{i,\beta} \subseteq U_{i,\beta} \subseteq F_\beta)_{i=1}^3$. Thus the active reduced task is $\mathsf{M}_{\mathrm{om}} \oplus \mathsf{M}_{\mathrm{res}}$. From now on, M, F, A_i, U_i refer to this task and its spaces unless a different task is specified. The reductions preserve U^*, D^*, and every $\mathcal{R}_L$. The residual has d_{priv} private source coordinates and d_{tri} three-line blocks, hence source dimension $d_{\mathrm{priv}} + 2d_{\mathrm{tri}}$. In particular, $A_i = A_{i,\mathrm{om}} \oplus A_{i,\mathrm{res}}$ and $U_i = F_{\mathrm{om}} \oplus U_{i,\mathrm{res}}$, where $A_{i,\mathrm{res}} = \mathcal{P}_i \oplus (\mathcal{T}_{\mathrm{res}} \otimes \ell_i)$.

For L instances, choose bases adapted to this decomposition and reuse X, W_i' for the reduced source and side information. Here $d = \dim F$ is the dimension of the active reduced source, so $X \in \mathbb{F}_q^{d \times L}$. Using (12), define, for $\beta \in \{\mathrm{om}, \mathrm{res}\}$,

$$\begin{aligned} \Omega = X^{\mathsf{T}} F_{\mathrm{om}}, \qquad W_i'^{,\beta} = X^{\mathsf{T}} A_{i,\beta}, \\ W_i' = (W_i'^{,\mathrm{om}}, W_i'^{,\mathrm{res}}). \end{aligned} \tag{32}$$

Write $\mathbf{K}_i = X^{\mathsf{T}} K_i^{\mathrm{act}}$ and $\mathbf{P}_i = X^{\mathsf{T}} \mathcal{P}_i$. Let $\mathbf{T}_i^\beta$ collect User i's side information from the three-line blocks in component β, with $\mathbf{T}_1^{\mathrm{om}}, \mathbf{T}_2^{\mathrm{om}} \in \mathbb{F}_q^{L \times d_{\mathrm{com}}}$, $\mathbf{T}_1^{\mathrm{res}}, \mathbf{T}_2^{\mathrm{res}} \in \mathbb{F}_q^{L \times d_{\mathrm{tri}}}$, and

$$\mathbf{T}_3^\beta = \mathbf{T}_2^\beta - \mathbf{T}_1^\beta. \tag{33}$$

Under a uniform source, the blocks $\mathbf{T}_1^\beta, \mathbf{T}_2^\beta$, $\beta \in \{\mathrm{om}, \mathrm{res}\}$, and all $\mathbf{K}_i, \mathbf{P}_i$ are mutually independent. For $\{i,j,k\} = [3]$,

$$\begin{aligned} \Omega &= (\mathbf{T}_1^{\mathrm{om}}, \mathbf{T}_2^{\mathrm{om}}, \mathbf{K}_1, \mathbf{K}_2, \mathbf{K}_3), \\ W_i'^{,\mathrm{om}} &= (\mathbf{T}_i^{\mathrm{om}}, \mathbf{K}_j, \mathbf{K}_k), \end{aligned} \tag{34}$$

and

$$W_i'^{,\mathrm{res}} = (\mathbf{P}_i, \mathbf{T}_i^{\mathrm{res}}). \tag{35}$$

From the common broadcast and its side information W_i', User i must recover Ω and $X^{\mathsf{T}} U_{i,\mathrm{res}}$. The residual demand may mix private coordinates and coordinates from the three-line blocks without determining the entire residual source.

Each deleted component splits from the task and has no demand beyond local side information. A reduced code can ignore it. Conversely, fix it to zero in an original code; attainable alphabets cannot grow. Thus the reduction also preserves nonlinear cost regions.

B. Residual Upload Bound

Conditioned on Ω, the next lemma applies functional submodularity to a common function recoverable from the uploads and the side information of either of two users. The two retained conditional mutual information terms will be used to add the omniscience cost.

Lemma 4 (Common-function entropy bound)**.** *Let $W_i'^{,\mathrm{res}} = (Y_i, T_i)$, $i \in [3]$, be finite-alphabet random variables that are pairwise independent given Ω. Let $M_i = e_i(W_i'^{,\mathrm{res}}, \Omega)$ and $M = (M_1, M_2, M_3)$. Suppose*

$$\begin{aligned} H(Y_i \mid M_i, T_i, \Omega) &= 0, \qquad i \in [3], \\ H(T_i \mid T_j, T_k, \Omega) &= 0, \qquad \{i,j,k\} = [3], \end{aligned}$$

and a finite-alphabet variable Θ satisfies

$$\begin{aligned} H(\Theta \mid M, T_1, \Omega) = H(\Theta \mid M, T_2, \Omega) = 0, \\ H(T_1 \mid M, \Theta, T_3, \Omega) = 0. \end{aligned}$$

Define

$$\Gamma = 2H(W_1'^{,\mathrm{res}}, W_2'^{,\mathrm{res}}, W_3'^{,\mathrm{res}} \mid \Omega) - \sum_i H(W_i'^{,\mathrm{res}} \mid \Omega).$$

Then

$$\begin{aligned} \sum_i H(M_i \mid \Omega) \ge \Gamma &+ I(M_2; M_3 \mid M_1, T_1, \Omega) \\ &+ I(M_1; M_3 \mid M_2, T_2, \Omega). \end{aligned} \tag{36}$$

The cyclic bounds hold under the corresponding recovery hypotheses.

Proof: Suppress conditioning on Ω, and set $h = H(W_1'^{,\mathrm{res}}, W_2'^{,\mathrm{res}}, W_3'^{,\mathrm{res}})$ and $a_i = H(W_i'^{,\mathrm{res}})$. Local recovery gives $H(M_i, T_i) = a_i$. Both (M, T_1, T_2) and (M, Θ, T_3) determine the residual source and are functions of it, so both have entropy h. Functional submodularity [21, Lem. 3], applied to (M, T_1) and (M, T_2) with common function (M, Θ), gives

$$\begin{aligned} H(M, T_1) + H(M, T_2) &\ge h + H(M, \Theta) \\ &\ge 2h - H(T_3 \mid M_3) \\ &= 2h - a_3 + H(M_3). \end{aligned} \tag{37}$$

For $i = 1, 2$, let $J_i = I(M_j; M_k \mid M_i, T_i)$, where $\{i,j,k\} = [3]$. Pairwise independence gives

$$H(M, T_i) = a_i + \sum_{j \ne i} H(M_j) - J_i.$$

Substitution yields $\sum_i H(M_i) \ge \Gamma + J_1 + J_2$. Restoring conditioning on Ω proves the stated bound. ∎

C. Adding the Omniscience Cost

The uploads may jointly reveal information about Ω beyond that accounted for by $\sum_i I(M_i; \Omega)$. The next lemma bounds each user's omniscience deficit by two individual mutual information terms and one retained conditional mutual information term. Combining this inequality with the residual bound gives the sum of the two component costs.

Lemma 5 (Upload lower bound with an omniscience component)**.** *Let Ω be independent of the joint tuple*

$(W_1'^{,\mathrm{res}}, W_2'^{,\mathrm{res}}, W_3'^{,\mathrm{res}})$*, whose components are pairwise independent. User* i *has side information* $(W_i'^{,\mathrm{om}}, W_i'^{,\mathrm{res}})$*, where* $W_i'^{,\mathrm{om}}$ *is a function of* Ω*, and sends* $M_i = e_i(W_i'^{,\mathrm{om}}, W_i'^{,\mathrm{res}})$*. For* $\{i,j,k\} = [3]$*, suppose*

$$\begin{aligned} &H(\Omega \mid W_i'^{,\mathrm{om}}, W_i'^{,\mathrm{res}}, M_j, M_k) = 0, \\ &J_i = I(M_j; M_k \mid W_i'^{,\mathrm{res}}, \Omega), \end{aligned} \tag{38}$$

and, for some Γ*,*

$$\sum_\ell H(M_\ell \mid \Omega) \geq \Gamma + J_i + J_j, \qquad i \neq j. \tag{39}$$

With $\delta_i = H(\Omega \mid W_i'^{,\mathrm{om}})$*,*

$$\sum_\ell H(M_\ell) \geq \Gamma + \max\left\{ \max_i \delta_i, \frac{1}{2}\sum_i \delta_i \right\}. \tag{40}$$

Proof: Put $I_i = I(M_i;\Omega)$. For $j \neq i$, the assumptions make $W_i'^{,\mathrm{res}}$ independent of (Ω, M_j). Since $W_i'^{,\mathrm{om}}$ is a function of Ω,

$$\begin{aligned} &I(\Omega; M_j \mid W_i'^{,\mathrm{res}}, W_i'^{,\mathrm{om}}) \\ &\quad = I(\Omega; M_j \mid W_i'^{,\mathrm{om}}) \\ &\quad = I_j - I(M_j; W_i'^{,\mathrm{om}}) \leq I_j. \end{aligned}$$

Recovery and the chain rule yield

$$\begin{aligned} \delta_i &= \sum_{\ell \neq i} I(\Omega; M_\ell \mid W_i'^{,\mathrm{res}}, W_i'^{,\mathrm{om}}) + J_i \\ &\quad - I(M_j; M_k \mid W_i'^{,\mathrm{res}}, W_i'^{,\mathrm{om}}) \\ &\leq I_j + I_k + J_i. \end{aligned} \tag{41}$$

Set $\Delta = \sum_\ell H(M_\ell \mid \Omega) - \Gamma$. Equation (39) gives $\Delta \geq \max_i J_i$ and $2\Delta \geq \sum_i J_i$. Applying these inequalities to (41), individually and summed over i, proves (40), since $\sum_i H(M_i) = \Gamma + \Delta + \sum_i I_i$. ∎

D. Completing the Upload Converse

We verify the hypotheses of the entropy lemmas in two steps. First, each user's upload must distinguish its private coordinates. Second, intersections of the signal spaces U_i provide the required common functions.

Fix $L \geq 1$ and a uniform source for the reduced task in Section III-A. The variables in (32)–(35) denote their L-instance blocks. The residual blocks are pairwise independent and jointly independent of Ω. The rank–entropy dictionary gives

$$H(\Omega \mid W_i'^{,\mathrm{om}}) = L d_{\mathrm{om},i}, \tag{42}$$

$$\begin{aligned} \Gamma &= L\big[2(d_{\mathrm{priv}} + 2d_{\mathrm{tri}}) - (d_{\mathrm{priv}} + 3d_{\mathrm{tri}})\big] \\ &= L U_{\mathrm{res}}. \end{aligned} \tag{43}$$

First, zero-error recovery forces each encoder to distinguish its private coordinates once the remaining local coordinates are fixed:

$$H(\mathbf{P}_i \mid M_i, \mathbf{T}_i^{\mathrm{res}}, \Omega) = 0, \qquad i \in [3]. \tag{44}$$

Indeed, fix $(\mathbf{T}_i^{\mathrm{res}}, \Omega)$. If two private blocks give the same M_i, keep all other source coordinates fixed. The other users' side information and all uploads, hence the broadcast, coincide. Zero-error recovery forces all remote demands to coincide. Their projections span $\mathcal{P}_i$ by (23), so the two private blocks are equal.

Lemma 6 (Common functions from signal-space intersections)**.** *For the active residual task* $\mathsf{M}_{\mathrm{res}}$*,*

$$\mathcal{T}_{\mathrm{res}} \otimes E \subseteq U_{3,\mathrm{res}} + (U_{1,\mathrm{res}} \cap U_{2,\mathrm{res}}), \tag{45}$$

with cyclic counterparts. The common function $\Theta_{12} = X^{\mathsf{T}}(U_{1,\mathrm{res}} \cap U_{2,\mathrm{res}})$*, and its cyclic counterparts, satisfy the hypotheses of Lemma 4. Consequently, with* Γ *as in* (43) *and* $J_i = I(M_j; M_k \mid W_i'^{,\mathrm{res}}, \Omega)$*, where* $\{i,j,k\} = [3]$*,*

$$\sum_\ell H(M_\ell \mid \Omega) \geq \Gamma + J_i + J_j, \qquad i \neq j. \tag{46}$$

Appendix B-A proves (45) and checks the recovery conditions with $Y_i = \mathbf{P}_i$, $T_i = \mathbf{T}_i^{\mathrm{res}}$. By (44), $(M_i, \mathbf{T}_i^{\mathrm{res}})$ and $W_i'^{,\mathrm{res}}$ determine each other given Ω, so the retained mutual informations are exactly J_i.

Converse in Theorem 3: From its side information and the other users' uploads, User i can reconstruct the broadcast and recover Ω, establishing (38). Apply Lemma 5 with (43), (46), and $\delta_i = L d_{\mathrm{om},i}$ to obtain

$$\sum_i H(M_i) \geq L[U_{\mathrm{res}} + U_{\mathrm{om}}].$$

Since $H(M_i) \leq \log_q |\operatorname{im} e_i|$, this lower-bounds the upload cost of every nonlinear fixed-length zero-error code at every blocklength, and hence throughout the closed cost region. ∎

E. Matching Upload Construction

We code the two components separately. For the residual component, choose a basis from the local generators in Lemma 19 (Appendix B-B). This gives scalar uploads of total length U_{res} symbols.

With unconstrained broadcast, the relay may forward the entire upload tuple. Any two users' side information in the omniscience component determines Ω, and User i's missing dimension is $d_{\mathrm{om},i}$. The classical rate constraints for communication for omniscience [30, Eq. (10) and Prop. 1] therefore reduce to the three deficit inequalities below, together with nonnegativity. The following lemma attains the resulting minimum sum rate by noninteractive linear uploads with zero error over the original field, using at most two instances.

Lemma 7 (Noninteractive omniscience upload)**.** *For the omniscience component in* (30)–(31)*, the minimum total cost of separate noninteractive linear uploads enabling every user to recover* Ω *is*

$$\begin{aligned} U_{\mathrm{om}} = \min_{r_1,r_2,r_3 \geq 0} \ &\sum_i r_i \\ \text{s.t.} \quad &r_j + r_k \geq d_{\mathrm{om},i}, \qquad \{i,j,k\} = [3]. \end{aligned} \tag{47}$$

The optimum is attained over $\mathbb{F}_q$ *with one instance when integral and with two otherwise.*

Appendix B-C proves the sufficiency of the deficit bounds by giving the stated finite-blocklength construction.

Achievability in Theorem 3: Set $L_0 = 1$ when $U_{\rm om}$ is integral and $L_0 = 2$ otherwise. Combine L_0 copies of the scalar residual upload in Lemma 19 with the omniscience upload in Lemma 7 on the same instances. Their direct sum is valid because the residual coordinate change fixes $F_{\rm om}$. Broadcasting the entire upload tuple allows every user to decode. The total upload length is $L_0[U_{\rm res} + U_{\rm om}]$, which gives the claimed normalized cost over $\mathbb{F}_q$. ∎

This completes the upload problem with unconstrained broadcast. Broadcasting the entire upload tuple need not attain the fully informed broadcast optimum D^*. We next choose the relay input while keeping the total upload cost at U^*, so that the induced broadcast task also attains D^*.

IV. From Local Uploads to a Joint Broadcast

Upload-optimal linear schemes may make different linear functions available at the relay. We track how this choice affects each entry of the LCBC broadcast profile and reduce the design to a finite set of component choices. Appendix C supplies their structural justification.

A. Feasible Relay Inputs

We first work with one instance. A subspace $P \le F$ specifies the relay input $X^{\mathsf T}P$, namely the linear functions made available by the uploads. A *feasible relay subspace* satisfies

$$P = \sum_{i=1}^{3} (P \cap A_i), \qquad U_i \subseteq A_i + P. \tag{48}$$

The first condition permits a basis of P to be selected from the local spaces $P \cap A_i$, giving uploads of total cost $\dim P$. The second ensures demand recovery from the relay input and each user's side information. For L instances, replace (F, A_i, U_i) by $(F^{\oplus L}, A_i^{\oplus L}, U_i^{\oplus L})$ and divide block costs by L.

Lemma 8 (Reduction to the relay space)**.** *Let $P \le F$, $A_i \le U_i \le F$, and $U_i \subseteq A_i + P$ for $i \in [3]$. At every blocklength, broadcasting from $X^{\mathsf T}P$ has the same achievable zero-error broadcast alphabet sizes as a fully informed task on P with*

$$\widehat{A}_i = A_i \cap P, \qquad \widehat{U}_i = U_i \cap P. \tag{49}$$

Denote this induced task by $\mathsf{M}_P = (\widehat{A}_i \subseteq \widehat{U}_i \subseteq P)_{i=1}^3$. Both tasks allow nonlinear fixed-length broadcasts.

Proof: Fix L. Since $U_i = A_i + \widehat{U}_i$, every induced-task code solves the original task with the same broadcast.

Conversely, write $A_i = \widehat{A}_i \oplus C_i$. The induced-task decoder runs the original decoder with side information $X^{\mathsf T}\widehat{A}_i$ and zero C_i-coordinates. Since $C_i \cap P = 0$, every source block X has an auxiliary block $\widetilde{X}_i$ such that $\widetilde{X}_i^{\mathsf T}P = X^{\mathsf T}P$, $\widetilde{X}_i^{\mathsf T}C_i = 0$. This block gives the same broadcast and the supplied side information, so zero-error decoding recovers $\widetilde{X}_i^{\mathsf T}\widehat{U}_i = X^{\mathsf T}\widehat{U}_i$. The auxiliary blocks may differ across users. Since $X^{\mathsf T}P$ ranges over all source blocks of the induced task and the encoder is unchanged, both tasks have the same achievable alphabets. ∎

The induced task preserves the demand dimensions $\mathrm{rk}\left(\widehat{U}_i \mid \widehat{A}_i\right) = m_i$; its optimum follows from (14)–(16).

B. Tracking the LCBC Broadcast Profile

We order the entries of the LCBC broadcast profile from (14) and (15) as

$$\begin{aligned} \boldsymbol{F} &= (F_{123}, F_{132}, F_{213}, F_{231}, F_{312}, F_{321}), \\ \boldsymbol{Q} &= (Q_{11}, Q_{12}, Q_{13}, Q_{21}, Q_{22}, Q_{23}, Q_{31}, Q_{32}, Q_{33}). \end{aligned} \tag{50}$$

Here $\boldsymbol{F} \in \mathbb{R}^6$ and $\boldsymbol{Q} \in \mathbb{R}^9$ collect the first- and second-family entries, respectively. The shorthand $F_j = (\boldsymbol{F})_j$ follows the order in (50), and $\boldsymbol{e}_j$ is the jth standard basis vector of $\mathbb{R}^6$. For a task M, write $(B_1(\mathsf{M}), \dots, B_{15}(\mathsf{M})) = (\boldsymbol{F}(\mathsf{M}), \boldsymbol{Q}(\mathsf{M}))$. For a task direct sum, subspace sums and intersections decompose componentwise, so every conditional rank in (14) and (15) adds. Consequently,

$$B_\nu\left(\bigoplus_a \mathsf{M}_a\right) = \sum_a B_\nu(\mathsf{M}_a), \qquad \nu \in [15]. \tag{51}$$

The broadcast optimum is obtained by taking the maximum only after these entrywise sums:

$$D^*\left(\bigoplus_a \mathsf{M}_a\right) = \max_{\nu \in [15]} \sum_a B_\nu(\mathsf{M}_a) \le \sum_a D^*(\mathsf{M}_a). \tag{52}$$

The inequality can be strict. Consider two independent one-symbol tasks: in one, User 1 knows x and User 2 demands it; in the other, User 2 knows y and User 1 demands it. User 3 has no demand or side information. Each task separately needs one broadcast symbol, whereas their direct sum needs only $x + y$. The relay choices must therefore be coordinated across components to meet the broadcast budget for the full task. For a feasible component relay subspace P, let $(\boldsymbol{F}_P, \boldsymbol{Q}_P)$ be the profile of the induced task M_P in Lemma 8. For upload-optimal P, define its *first-family increment* over the fully informed component profile as

$$\boldsymbol{e}(P) \triangleq \boldsymbol{F}_P - \boldsymbol{F}_{\rm full}. \tag{53}$$

A component is *profile-preserving* if some upload-optimal relay subspace preserves all fifteen entries. For upload-optimal componentwise choices $P = \bigoplus_a P_a$, the induced task also decomposes as $\mathsf{M}_P = \bigoplus_a (\mathsf{M}_a)_{P_a}$. Hence, whenever each choice preserves the second-family entries,

$$\begin{aligned} \boldsymbol{F}_P &= \boldsymbol{F}(\mathsf{M}) + \sum_a \boldsymbol{e}(P_a), \\ \boldsymbol{Q}_P &= \boldsymbol{Q}(\mathsf{M}), \qquad \mathsf{M} = \bigoplus_a \mathsf{M}_a. \end{aligned} \tag{54}$$

For a budget $D \ge D^*(\mathsf{M})$, it therefore suffices to keep each first-family increment sum at most $D - F_j(\mathsf{M})$. The second-family entries already lie within the budget.

For integral U^* and an integer budget $D \ge D^*$, we seek a scalar relay subspace $P \le F$ satisfying

$$\begin{aligned} P = \sum_{i=1}^{3} (P \cap A_i), \qquad U_i \subseteq A_i + P, \\ \dim P = U^*, \qquad B_\nu(\mathsf{M}_P) \le D, \quad \nu \in [15]. \end{aligned} \tag{55}$$

The profile constraints bound the induced LCBC optimum by D. Sections V and VI construct this relay subspace and its scalar broadcast over the original field. The general two-instance result will follow by applying the scalar construction to $\mathsf{M}^{\oplus 2}$.

C. Decomposing the Relay Choice

We first construct the component relays over a field K with $|K| \geq 3$. For binary tasks, Appendix F-C gives additional choices that fit integer broadcast budgets directly over $\mathbb{F}_2$.

A task is *indecomposable* if no nontrivial ambient direct sum respects all six spaces A_i, U_i. For a feasible task, its *side-information excess* is $\chi \triangleq \sum_i \dim A_i - \dim F \geq 0$, with equality exactly when $F = A_1 \oplus A_2 \oplus A_3$.

Proposition 9 (Residual decomposition)**.** *Let $\mathsf{M}_{\mathrm{res}} = \bigoplus_a \mathsf{M}_a$ be an indecomposable decomposition over $\mathbb{F}_q$, with ambient spaces F_a and $A_{i,a} = A_i \cap F_a$, $U_{i,a} = U_i \cap F_a$, and $\mathsf{M}_a = (A_{i,a} \subseteq U_{i,a} \subseteq F_a)_{i=1}^3$. Each component is an active residual task, and minimum upload is additive:*

$$\begin{aligned} U^*(\mathsf{M}_{\mathrm{res}}) &= 2\dim F_{\mathrm{res}} - \sum_i \dim A_{i,\mathrm{res}} \\ &= \sum_a U^*(\mathsf{M}_a) = U_{\mathrm{res}}, \end{aligned} \tag{56}$$

where $U^*(\mathsf{M}_a) = 2\dim F_a - \sum_i \dim A_{i,a}$.

Appendix A-B proves this statement. The classification in Appendix C-A then identifies the positive-excess components over their defining finite field.

The zero-excess case follows from the upload converse. Here $F = A_1 \oplus A_2 \oplus A_3$, and (44) forces each encoder to distinguish all $q^{L \dim A_i}$ local side-information blocks after inactive directions have been removed. Uploading the linear combinations specified by a basis of each A_i is therefore optimal, and $P = F$ preserves the fully informed profile. Only positive-excess components require a choice of relay input.

D. A Small Component That Requires Coordination

Over every finite field, R_1 has independent coordinates X_1, X_2, X_3, V, R, with $T = R - V$. Users 1, 2, and 3 know (X_1, V), (X_2, R), and (X_3, T), and demand $X_2 + X_3$, $X_1 + X_3$, and $X_1 + X_2 + V$, respectively. Over $\mathbb{F}_2$, this is (1). The upload optimum is four, attained by the feasible relay subspaces $P_V = \langle X_1, V, X_2, X_3\rangle$, $P_R = \langle X_1, X_2, R, X_3\rangle$, and $P_T = \langle X_1 + V, X_2, X_3, T\rangle$, whose generators are local. Their first-family increments are $\boldsymbol{e}_4 + \boldsymbol{e}_6$, $\boldsymbol{e}_2 + \boldsymbol{e}_5$, and $\boldsymbol{e}_1 + \boldsymbol{e}_3$, respectively.

For every $n \geq 1$, these three increment vectors are attained for R_n by relay subspaces that preserve all nine second-family entries. Its fully informed profile is $\boldsymbol{F}(\mathsf{R}_n) = n\mathbf{1}_6$, $2\boldsymbol{Q}(\mathsf{R}_n) = 3n\mathbf{1}_9$. Appendix C-C gives the model and rank verification. Each choice uses one entry in each cyclic orientation, $\{F_1, F_4, F_5\}$ and $\{F_2, F_3, F_6\}$, so it must fit the broadcast budget of the full task.

E. Sufficient Choices and Their Broadcast Cost

Table I lists sufficient relay choices attaining the minimum component upload cost. Each preserves all nine second-family entries and changes the first-family entries by the displayed increment. An increment $\boldsymbol{e}_j$ raises F_j by one; whether this changes D^* depends on the other profile entries.

Appendix C defines the component families in the table. For $T_1^{(r)}$, $O^{(r)}$, and the family $\mathsf{R} = \{\mathsf{R}_n : n \geq 1\}$, the increments are independent of size n; the boundary tasks Z_{ab} are two-dimensional. Allocation therefore uses only component multiplicities, the stated rank bounds, and the full task's broadcast profile.

For any three-user task, define the *pair margins* from its fully informed profile; their subscripts index first-family entries:

$$\begin{gathered} \kappa_{16} \triangleq 2Q_{22} - F_1 - F_6, \quad \kappa_{24} \triangleq 2Q_{33} - F_2 - F_4, \\ \kappa_{35} \triangleq 2Q_{11} - F_3 - F_5. \end{gathered} \tag{57}$$

For each listed component, all nine $Q_{i\ell}$ coincide.

These margins are always nonnegative:

$$\kappa_{16},\ \kappa_{24},\ \kappa_{35} \geq 0. \tag{58}$$

For example, conditional-rank submodularity gives

$$\begin{aligned} &\mathrm{rk}(U_{12} \mid A_2) + \mathrm{rk}(U_{23} \mid A_2) \\ &\quad \geq \mathrm{rk}(U_{12} + U_{23} \mid A_2) + \mathrm{rk}(U_{123} \mid A_2)\,. \end{aligned}$$

Substitution in (14) and (15) yields $\kappa_{16} \geq 0$; cyclic relabeling gives the other two. These margins translate the second-family entries into capacity constraints for allocating first-family increments within the full task's broadcast budget.

In Table I, $\mathcal{E}_t$ lists the increments for type t, and $\underline{\kappa}_{ab}(t)$ is the listed pair-margin lower bound. For every component γ of listed type $t(\gamma)$,

$$\kappa_{ab}(\gamma) \geq \underline{\kappa}_{ab}(t(\gamma)), \quad (a,b) \in \{(1,6),(2,4),(3,5)\}. \tag{59}$$

Only pair-capacity constraints use these lower bounds; all other calculations use the component's full profile.

Proposition 10 (Sufficient residual relay choices)**.** *Over any finite field K with $|K| \geq 3$, every residual indecomposable outside Table I, including every zero-excess summand, is profile-preserving. Each listed family admits all stated increment choices by upload-optimal relay subspaces over K, preserving all nine $Q_{i\ell}$.*

Appendix C-D proves the proposition using standard quiver results to identify components and explicit local uploads to verify their profiles.

a) The omniscience component: When U_{om} is integral, Lemma 7 supplies scalar uploads. Recall that $d_{\mathrm{om},i} = d_{\mathrm{com}} + d_{\mathrm{pair},i}$ is User i's missing dimension in the omniscience component. Its fully informed profile follows directly from (14) and (15), since $U_i = F_{\mathrm{om}}$:

$$\begin{aligned} &\boldsymbol{F}_{\mathrm{om}} = (d_{\mathrm{om},1}, d_{\mathrm{om},1}, d_{\mathrm{om},2}, d_{\mathrm{om},2}, d_{\mathrm{om},3}, d_{\mathrm{om},3}), \\ &2Q_{i\ell} = \sum_j d_{\mathrm{om},j} - d_{\mathrm{om},\ell}. \end{aligned} \tag{60}$$

Every feasible relay subspace on L instances induces an omniscience task with deficits $Ld_{\mathrm{om},i}$, hence the same normalized profile. Its pair margins and the quantities in (67)–(68) vanish. The omniscience profile still contributes to the total broadcast budget, and its entries need not coincide when the active pair dimensions $d_{\mathrm{pair},i}$ differ.

TABLE I

SUFFICIENT FIRST-FAMILY INCREMENTS IN THE ORDER (50). MARGINS ARE EVALUATED AT $n = 1$ FOR INDEXED FAMILIES AND $n = 0$ FOR BOUNDARY TYPES Z_{ab}. IN $T_1^{(r)}$ AND $O^{(r)}$, USER r HAS THE ENLARGED SIGNAL SPACE U_r OR SIDE-INFORMATION SPACE A_r, RESPECTIVELY.

Component type t	Pair-margin lower bounds $(\underline{\kappa}_{16}, \underline{\kappa}_{24}, \underline{\kappa}_{35})$	Sufficient increments $\mathcal{E}_t$
$T_1^{(1)}$	$(1,1,2)$	$\{\boldsymbol{e}_5, \boldsymbol{e}_3\}$
$T_1^{(2)}$	$(2,1,1)$	$\{\boldsymbol{e}_6, \boldsymbol{e}_1\}$
$T_1^{(3)}$	$(1,2,1)$	$\{\boldsymbol{e}_4, \boldsymbol{e}_2\}$
$O^{(1)}$	$(1,1,0)$	$\{\boldsymbol{e}_4+\boldsymbol{e}_6, \boldsymbol{e}_2, \boldsymbol{e}_1\}$
$O^{(2)}$	$(0,1,1)$	$\{\boldsymbol{e}_4, \boldsymbol{e}_3, \boldsymbol{e}_2+\boldsymbol{e}_5\}$
$O^{(3)}$	$(1,0,1)$	$\{\boldsymbol{e}_6, \boldsymbol{e}_5, \boldsymbol{e}_1+\boldsymbol{e}_3\}$
R	$(1,1,1)$	$\{\boldsymbol{e}_1+\boldsymbol{e}_3, \boldsymbol{e}_2+\boldsymbol{e}_5, \boldsymbol{e}_4+\boldsymbol{e}_6\}$
Z_{16}	$(1,0,0)$	$\{\boldsymbol{e}_1, \boldsymbol{e}_6\}$
Z_{24}	$(0,1,0)$	$\{\boldsymbol{e}_2, \boldsymbol{e}_4\}$
Z_{35}	$(0,0,1)$	$\{\boldsymbol{e}_3, \boldsymbol{e}_5\}$

V. RELAY ALLOCATION AND DEMAND-CYCLE REPAIR

The relay subspaces in Proposition 10 preserve all nine second-family entries. It remains to choose their first-family increment vectors so that all six first-family entries fit within a common broadcast budget. Directed demand cycles allow us to modify relay choices when the componentwise allocation fails.

A. Why Componentwise Relay Choices May Fail

For the complete active reduced task M, fix a broadcast budget $D \geq D^*(\mathsf{M})$, measured in q-ary units per instance. The slack available for each first-family increment is

$$s_j = D - F_j(\mathsf{M}), \qquad j \in [6]. \tag{61}$$

These slacks are nonnegative. In the order (50), the index sets $H_+ = \{1,4,5\}$ and $H_- = \{2,3,6\}$ correspond to the two cyclic orientations. Write $s_\sigma = \sum_{j \in H_\sigma} s_j$, $\sigma \in \{+,-\}$, for their total slacks. Let μ_{R} denote the number of summands in the family R. Each listed relay choice for the family R consumes one unit of slack in each orientation.

A *positive directed demand cycle* consists of three independent symbols x, y, z, with side information and demands

$$\text{User 1: } z \to y, \quad \text{User 2: } y \to x, \quad \text{User 3: } x \to z, \tag{62}$$

where $a \to b$ means that the user knows a and demands b. The reverse cycle has the same side information z, y, x and demands x, z, y, respectively.

Consider R_1 together with one positive cycle. Its fully informed profile, computed in (79), is $\boldsymbol{F} = (3,2,2,3,3,2)$, and all nine $Q_{i\ell} = 3$. At $D = D^*(\mathsf{M}) = 3$, $\mu_{\mathsf{R}} = 1$ and $\boldsymbol{s} = (0,1,1,0,0,1)$. Every R_1 relay choice in Table I uses a coordinate in H_+ with zero slack, so none fits the budget. The construction in Lemma 13 resolves this obstruction by mixing uploads across the R_1 summand and the cycle. Its increment is $\boldsymbol{e}_2$; it preserves all second-family entries and the minimum upload cost $U^*(\mathsf{R}_1) + 3 = 7$.

B. The Aggregate Selection Criterion

Let $\boldsymbol{\mu} = (\mu_t)_t$ collect the multiplicities of the component types listed in the ten rows of Table I, in table order, summed over $n \geq 1$ for indexed families. In particular, $\mu_{\mathsf{R}} = \sum_{n \geq 1} \mu_{\mathsf{R},n}$, where $\mu_{\mathsf{R},n}$ counts R_n. For a nonnegative slack vector $\boldsymbol{s}$, an *aggregate selection* assigns weights $\lambda_{t,\boldsymbol{e}}$ to the choices $\boldsymbol{e} \in \mathcal{E}_t$ such that

$$\begin{aligned} &\lambda_{t,\boldsymbol{e}} \geq 0, \qquad \sum_{\boldsymbol{e} \in \mathcal{E}_t} \lambda_{t,\boldsymbol{e}} = \mu_t, \\ &\sum_t \sum_{\boldsymbol{e} \in \mathcal{E}_t} \lambda_{t,\boldsymbol{e}} (\boldsymbol{e})_j \leq s_j, \qquad\qquad j \in [6]. \end{aligned} \tag{63}$$

The weight $\lambda_{t,\boldsymbol{e}}$ specifies how many copies of family t, per instance, use the scalar relay choice $\boldsymbol{e}$. Integer weights select whole component copies and therefore give scalar uploads, followed by one joint broadcast.

For the complete-task slack vector in (61), $D \geq Q_{22}(\mathsf{M})$ gives

$$s_1 + s_6 = 2D - F_1(\mathsf{M}) - F_6(\mathsf{M}) \geq \kappa_{16}(\mathsf{M}),$$

and similarly for the other two pairs. Entrywise profile additivity in (51) and (58)–(59) imply

$$s_a + s_b \geq \sum_t \mu_t \underline{\kappa}_{ab}(t), \quad (a,b) \in \{(1,6),(2,4),(3,5)\}. \tag{64}$$

Profile-preserving summands have nonnegative margins, so omitting their contributions preserves the lower bounds in (64); the omniscience margins are zero. Under these pair constraints, feasibility reduces to the two orientation constraints below.

Proposition 11 (Aggregate selection criterion)**.** *For $(\boldsymbol{\mu}, \boldsymbol{s}) \in \mathbb{R}_+^{10} \times \mathbb{R}_+^6$ satisfying* (64)*, the aggregate selection condition* (63) *is feasible if and only if*

$$s_+ \geq \mu_{\mathsf{R}}, \qquad s_- \geq \mu_{\mathsf{R}}. \tag{65}$$

If the multiplicities and slacks are integers, the allocation weights may also be chosen as integers.

Proof sketch: The two constraints are necessary because each relay choice for R consumes one unit in each orientation. For sufficiency, match the relay choices for R first. The pair

constraints and (65) provide enough capacity for this match; the remaining pair capacities then cover the other rows of Table I. Appendix D-A gives a constructive allocation using only addition, subtraction, minima, and maxima. These operations preserve integrality. ∎

C. Turning a Budget Deficit into a Demand Cycle

A failed orientation constraint identifies enough independent directed demands to form the cycles needed for repair. We extract these cycles from the complement of the summands in R.

For distinct i, j, k, let $\mathsf{E}_{i\to j}$ be one independent source symbol known only to User i and demanded by User j, with subspaces

$$A_i = U_i = U_j = \langle v \rangle, \qquad A_j = A_k = U_k = 0. \tag{66}$$

Extracting this task means splitting it off on independent source coordinates. Let $\mu_{i\to j}$ be the maximum number of such copies. The following additive quantities connect the positive-cycle edges to the LCBC profile entries:

$$\begin{aligned} p_{3\to 2} = F_1 + F_4 - 2Q_{23}, \quad p_{2\to 1} = F_1 + F_5 - 2Q_{12}, \\ p_{1\to 3} = F_4 + F_5 - 2Q_{31}, \end{aligned} \tag{67}$$

and for the reverse-cycle edges,

$$\begin{aligned} p_{2\to 3} = F_2 + F_6 - 2Q_{32}, \quad p_{3\to 1} = F_2 + F_3 - 2Q_{13}, \\ p_{1\to 2} = F_3 + F_6 - 2Q_{21}. \end{aligned} \tag{68}$$

All F- and Q-entries above are evaluated on the fully informed task. To extract cycles outside the summands in R, write

$$\mathsf{M} = \mathsf{M}_{\mathsf{R}} \oplus \mathsf{M}_{\mathrm{rem}}, \qquad \mathsf{M}_{\mathsf{R}} = \bigoplus_{n\ge 1} \mathsf{R}_n^{\oplus \mu_{\mathsf{R},n}}. \tag{69}$$

The complement $\mathsf{M}_{\mathrm{rem}}$ includes the omniscience component. The orientation deficits are

$$\varepsilon_\sigma = (\mu_{\mathsf{R}} - s_\sigma)_+, \qquad \sigma \in \{+,-\}. \tag{70}$$

Lemma 12 (Demand cycles from orientation deficits)**.** *For distinct i, j, k, every three-user task satisfies*

$$p_{k\to j} := F_{ijk} + F_{jki} - 2Q_{jk} \le \mu_{k\to j}. \tag{71}$$

Copies of all six directed edge types can be extracted simultaneously on disjoint source coordinates, attaining the stated multiplicities. Consequently, for each edge e of orientation $\sigma \in \{+,-\}$,

$$\mu_e(\mathsf{M}_{\mathrm{rem}}) \ge \varepsilon_\sigma. \tag{72}$$

Thus $\mathsf{M}_{\mathrm{rem}}$ supplies enough independent cycles to cover either deficit without changing a R_n summand. These statements hold over every field.

Proof: For the extraction step, project the directions in $A_k \cap U_j$ that lie outside $A_j + U_{ij}$ onto a complementary summand. This projection splits off copies of $\mathsf{E}_{k\to j}$ and preserves all six task subspaces. The resulting multiplicity is additive and vanishes on the other five edge types, so all six extractions can be performed successively. The exact quotient formula and the rank identity proving (71) are given in Appendix D-B.

To obtain cycles, note that every p_e uses two first-family entries from its orientation and a term $-2Q$. Since $Q \le D$ and $s_j \ge 0$, we have $p_e(\mathsf{M}) \ge -s_\sigma$. Also $p_e(\mathsf{R}_n) = -n$ by (107). Additivity therefore gives

$$\begin{aligned} \mu_e(\mathsf{M}_{\mathrm{rem}}) \ge p_e(\mathsf{M}_{\mathrm{rem}}) &= p_e(\mathsf{M}) + \sum_{n\ge 1} n\mu_{\mathsf{R},n} \\ &\ge \mu_{\mathsf{R}} - s_\sigma. \end{aligned} \tag{73}$$

The edge multiplicities are nonnegative. Simultaneous extraction allows the three edges in an orientation to be grouped into as many independent cycles as their smallest multiplicity, proving the claim. ∎

D. Directed-Cycle Upload Modification

Pair a cycle from Lemma 12 with a R_n summand. Mixing one uploaded coordinate with a cycle coordinate preserves the minimum upload cost and leaves a single increment in the opposite orientation.

Lemma 13 (Demand-cycle repair)**.** *For every $n \ge 1$ and every finite field, the direct sum of R_n and the cycle* (62) *admits an upload-optimal feasible relay subspace of dimension $3n+4$. In the profile order* (50)*, its increment is*

$$\boldsymbol{F}_P - \boldsymbol{F}_{\mathrm{full}} = \boldsymbol{e}_2, \qquad \boldsymbol{Q}_P = \boldsymbol{Q}_{\mathrm{full}}. \tag{74}$$

For the reverse cycle, such a relay subspace has increment $\boldsymbol{e}_4$ and again preserves every second-family entry. Both use scalar linear uploads over the original field.

Proof: Recall the coordinates of R_n from Appendix C-C. Here n is the component size, whereas the uploads below use blocklength one. The private vectors are $\mathbf{X}_i \in \mathbb{F}_q^n$. The three-line block has source coordinates V, R, with $T = R - V$. Users 1, 2, and 3 have side information $(\mathbf{X}_1, V)$, $(\mathbf{X}_2, R)$, and $(\mathbf{X}_3, T)$, respectively. The demands are $\mathbf{X}_2 + \mathbf{X}_3$, $\mathbf{X}_1 + \mathbf{X}_3$, and $\mathbf{X}_1 + A\mathbf{X}_2 + hV$, where $A = I_n + N$, N is one nilpotent Jordan block, and $h \in \mathbb{F}_q^n \setminus \operatorname{im} N$. Together with the cycle coordinates x, y, z, use the uploads

	$+$	$-$
M_1	$(\mathbf{X}_1, z - V)$	$(\mathbf{X}_1, V, z)$
M_2	$(\mathbf{X}_2, R, y)$	$(\mathbf{X}_2, y - R)$
M_3	$(\mathbf{X}_3, x)$	$(\mathbf{X}_3, x)$.

(75)

Every uploaded symbol is available from its user's side information, and the corresponding coefficient vectors are linearly independent. The total number of uploaded symbols, $3n+4 = (3n+1)+3$, equals the minimum upload cost by (56). Given the full relay input, Users 1 and 2 recover their demands directly. User 3 obtains $V = R - T$ and then $z = (z - V) + V$ in the positive orientation; it obtains $R = V + T$ and then $y = (y - R) + R$ in the reverse orientation. It also recovers its vector demand.

We now show that only one conditional rank of a pairwise signal-space intersection changes. For a coefficient row ξ, the relay hyperplanes are $P^\sigma = \ker \ell^\sigma$, where $\ell^+(\xi) = \xi_V + \xi_z$ and $\ell^-(\xi) = \xi_R + \xi_y$. The induced signal and side-information

spaces are $\widehat{U}_i = U_i \cap P^\sigma$ and $\widehat{A}_i = A_i \cap P^\sigma$, respectively. Feasibility gives

$$\dim \widehat{U}_i - \dim \widehat{A}_i = \dim U_i - \dim A_i = n+1. \tag{76}$$

Appendix E shows that, before and after restriction, the triple intersection is zero and each incident pair sum spans its user's demand quotient. Thus (108) gives

$$F_{ijk} = 2(n+1) - \mathrm{rk}(U_{ij} \mid A_j), \qquad 2Q_{i\ell} = 3(n+1), \tag{77}$$

with hats on the spaces after restriction. In particular, all nine second-family entries are preserved.

For a hyperplane $P = \ker \ell$, rank-nullity gives

$$\mathrm{rk}(B \mid A) - \mathrm{rk}(B \cap P \mid A \cap P) = \mathrm{rk}(\ell|_B) - \mathrm{rk}(\ell|_{B\cap A}). \tag{78}$$

Among the pairwise signal-space intersections, only U_{13} is restricted by P^+. The functional ℓ^+ is nonzero on $U_{13} \cap A_1 = \langle z \rangle$, whereas $U_{13} \cap A_3 = 0$. Thus only $\mathrm{rk}(U_{13} \mid A_3)$ drops by one. For P^-, only U_{23} among these intersections is restricted, with $U_{23} \cap A_2 = \langle y \rangle$ and $U_{23} \cap A_3 = 0$, so only $\mathrm{rk}(U_{23} \mid A_3)$ drops by one. Equation (77) then gives increments $\boldsymbol{e}_2$ and $\boldsymbol{e}_4$, respectively.

The full pair ranks also give

$$\begin{aligned} \boldsymbol{F}^+_{\mathrm{full}} &= (n+2, n+1, n+1, n+2, n+2, n+1), \\ \boldsymbol{F}^-_{\mathrm{full}} &= (n+1, n+2, n+2, n+1, n+1, n+2). \end{aligned} \tag{79}$$

These profiles follow from (77); see Appendix E. All constructions and rank identities hold over the original field, including characteristic two. ∎

E. Completing the Aggregate Relay Selection

We now combine cycle repair with Proposition 11 while keeping the original broadcast budget. Summing the three pair constraints in (64) gives

$$\sum_{j=1}^{6} s_j \geq 3\mu_{\mathsf{R}}. \tag{80}$$

Proposition 14 (Scalar relay selection at an integer budget)**.** *For a task in active reduced form over $\mathbb{F}_q$ with $U^* \in \mathbb{Z}$, every integer $D \geq D^*$ admits a relay subspace satisfying* (55)*. For $q = 2$, it can be chosen so that, after a permutation of the users,*

$$P = \widehat{A}_1 + \widehat{A}_2 + \widehat{A}_3, \qquad \widehat{A}_1 \cap (\widehat{A}_2 + \widehat{A}_3) = 0, \tag{81}$$

where $\widehat{A}_i = P \cap A_i$.

Proof: First suppose $q \geq 3$. The component multiplicities and the slacks $s_j = D - F_j(\mathsf{M})$ are integers. If both orientation constraints in (65) hold, Proposition 11 gives an integer allocation of the scalar relay choices.

Suppose $s_+ < \mu_{\mathsf{R}}$, and set $\varepsilon_+ = \mu_{\mathsf{R}} - s_+$. By Lemma 12, the complement of the R summands contains at least ε_+ independent positive cycles. Since ε_+ is an integer, select that many R_n summands, pair them with the cycles, and apply Lemma 13. Treat each pair as one combined component. The extracted edges are zero-excess, profile-preserving indecomposables. By Krull–Schmidt, their extraction leaves the positive-excess component multiplicities unchanged. The repair therefore reserves $\varepsilon_+ \boldsymbol{e}_2$ and leaves

$$\boldsymbol{s}' = \boldsymbol{s} - \varepsilon_+ \boldsymbol{e}_2, \qquad \mu'_{\mathsf{R}} = \mu_{\mathsf{R}} - \varepsilon_+ = s_+. \tag{82}$$

The $(2,4)$ pair constraint gives

$$s_2 \geq \mu_{\mathsf{R}} - s_4 \geq \mu_{\mathsf{R}} - s_+ = \varepsilon_+, \tag{83}$$

so $\boldsymbol{s}' \geq 0$. Every pair-capacity right side decreases by ε_+; only the $(2,4)$ left side decreases, by the same amount. The pair constraints remain valid. Moreover, $s'_+ = \mu'_{\mathsf{R}}$, while (80) gives

$$s'_- = s_- - \varepsilon_+ \geq 2\mu_{\mathsf{R}} \geq \mu'_{\mathsf{R}}. \tag{84}$$

Thus Proposition 11 allocates the remaining components. If $s_- < \mu_{\mathsf{R}}$, use the reverse-cycle repair, reserve $\varepsilon_- \boldsymbol{e}_4$, and interchange the two orientations. The two orientation constraints cannot fail simultaneously by (80).

All selections and repairs use integer numbers of components and preserve the second-family entries. Their uploads have total dimension U_{res}. Since $U^* = U_{\mathrm{res}} + U_{\mathrm{om}}$ and $U_{\mathrm{res}} \in \mathbb{Z}$, the hypothesis gives $U_{\mathrm{om}} \in \mathbb{Z}$. Add the scalar omniscience upload of Lemma 7, which preserves its component profile. This proves (55) for $q \geq 3$.

For $q = 2$, Lemma 25 supplements the same integer allocation with binary component choices and uses their pair margins to accommodate the resulting increments. It gives (55) over $\mathbb{F}_2$. Appendix F-D proves (81) for the selected relay subspace. ∎

VI. Scalar Codes and Two-Instance Achievability

We now construct a broadcast from the selected relay input over the original field. The binary case uses the separation property (81).

A. Scalar Broadcast Construction

For the next two propositions, let $\mathsf{M} = (A_i \subseteq U_i \subseteq F)_{i=1}^3$ be a fully informed LCBC task, with optimum $D^* = D^*(\mathsf{M})$.

Proposition 15 (Scalar LCBC over fields of size at least three)**.** *For $q \geq 3$, every fully informed three-user LCBC task over $\mathbb{F}_q$ admits a scalar linear code of cost $\lceil D^* \rceil$ over the same field, without requiring $F = A_1 + A_2 + A_3$.*

The proof is in Appendix F-G. The binary LCBC example in [21, Example 4 and App. A] shows that the field-size condition is needed for general LCBC tasks.

Proposition 16 (Binary scalar LCBC with separated side-information spaces)**.** *Suppose a binary three-user LCBC task satisfies $F = A_1 + A_2 + A_3$ and $A_1 \cap (A_2 + A_3) = 0$ after a permutation of the users. It admits a scalar linear code over $\mathbb{F}_2$ of cost $\lceil D^* \rceil$.*

The proof is in Appendix F-I.

B. Scalar Upload–Broadcast Coding

Corollary 17 (Scalar coding at integral minimum upload)**.** *For every feasible three-user linear task over $\mathbb{F}_q$, if $U^* \in \mathbb{Z}$, then a scalar linear code attains $(U^*, \lceil D^* \rceil)$.*

Proof: Apply the reductions of Section III-A, which preserve every $\mathcal{R}_L$, and then Proposition 14 with $D = \lceil D^* \rceil$. The selected relay subspace has upload cost U^* and induced optimum $D^*(\mathsf{M}_P) \le D$. For $q = 2$, it also satisfies (81). Thus Proposition 15, or Proposition 16 when $q = 2$, gives a scalar broadcast of cost $\lceil D^*(\mathsf{M}_P) \rceil \le D$. By Lemma 8, this solves the original task. The costs equal the stated values: a smaller upload cost contradicts U^*, and a smaller integer broadcast cost is below D^*. ∎

C. Proof of the Main Theorem

Proof of Theorem 1: The upload converse and the fully informed broadcast lower bound give $U \ge U^*$ and $D \ge D^*$ throughout $\overline{\mathcal{R}}$, including for nonlinear codes.

By Theorem 3 and (16), $U^*, D^* \in \frac{1}{2}\mathbb{Z}$. The upload rank formula and (51) imply

$$U^*(\mathsf{M}^{\oplus 2}) = 2U^*(\mathsf{M}), \qquad D^*(\mathsf{M}^{\oplus 2}) = 2D^*(\mathsf{M}).$$

Since these doubled minima are integers, Corollary 17 gives a scalar $\mathbb{F}_q$-linear code for $\mathsf{M}^{\oplus 2}$ with costs $(2U^*, 2D^*)$. Viewed as a two-instance code for M, it attains the normalized costs (U^*, D^*). Together with upward closure and the converse, this yields

$$[U^*, \infty) \times [D^*, \infty) \subseteq \mathcal{R}_2 \subseteq \overline{\mathcal{R}} \subseteq [U^*, \infty) \times [D^*, \infty).$$

When $U^*, D^* \in \mathbb{Z}$, Corollary 17 applied directly to M attains the corner with blocklength one. Otherwise, no scalar $\mathbb{F}_q$-linear code attains the corner: each attainable message alphabet is a linear image of cardinality q^r for an integer r, so both costs are integers. Together with the two-instance construction, this proves (18). The five-bit task in Section I-B establishes tightness of the uniform two-instance bound even for nonlinear codes. ∎

VII. Conclusion

For every feasible three-user linear task with separate noninteractive uploads to an initially uninformed relay, we characterized a rectangular closed upload–broadcast cost region, even under nonlinear fixed-length zero-error coding. The minimum total upload cost with unconstrained broadcast, given by an explicit rank formula, is simultaneously attainable with the fully informed broadcast optimum. Over the original field, attaining both minima requires minimum linear blocklength one when both are integers and two otherwise. The uniform two-instance bound is tight even for nonlinear codes.

Our converse establishes upload-cost additivity between the omniscience and residual components even under nonlinear mixing. Our construction selects among upload-optimal relay inputs through their induced LCBC profiles, using directed demand cycles to modify uploads without increasing their cost. Thus, broadcast optimality depends on which information is uploaded, not merely how much.

Extensions to more users and individual upload constraints raise the question of which structural conditions preserve simultaneous optimality and when an upload–broadcast tradeoff arises.

Acknowledgments

Under the authors' guidance, OpenAI's ChatGPT (GPT-5.6 and GPT-6) assisted with drafting and refining the Section I, including the literature review and related-work discussion, as well as the preparation of Section II, Section III, derivations in Sections IV, V, and VI, and technical derivations and proof exposition in the appendices. This assistance included substantial computational support for quiver decomposition searches and combinatorial exploration of water-filling allocations.

The authors developed the core ideas and strategy: first identify upload-optimal schemes and their design freedom, then select among them to attain the fully informed broadcast optimum. The authors also simplified, and refined the mathematical arguments and proofs. All main results have been formally verified in Lean; the Lean proofs and algorithm code are available at https://github.com/TSEason2022/The-Upload-Broadcast-Cost-Region-of-Three-User-Linear-Computation.

Appendix A
Normal Form, Passive Reduction, and Additivity

A. Construction and Invariance of the Upload Parameters

For the construction, separate the pure-demand and private-cancellation contributions to d_{tri}:

$$d_{\mathrm{pure}} := \dim \frac{\mathcal{D}_1 + \mathcal{D}_2 + \mathcal{D}_3}{C},$$
$$d_{\mathrm{mix}} := \operatorname{rk} G_0, \qquad d_{\mathrm{tri}} = d_{\mathrm{pure}} + d_{\mathrm{mix}}.$$

Only their sum enters the upload formula; the separate counts specify the construction in Appendix B-B.

1) Three-subspace decomposition: Apply the three-subspace decomposition [21, Lem. 2] to (A_1, A_2, A_3) after removing their common intersection. The lemma's basis labels now refer to side information: $\mathcal{P}_i = \langle B_{ic} \rangle$ and $\bar{K}_i = \langle B_{jk} \rangle$, $\{i, j, k\} = [3]$. Its aligned basis matrices $B_1 = B_{1(2,3)}$, $B_2 = B_{2(1,3)}$, and $B_3 = B_{3(1,2)}$ satisfy $B_1 + B_2 = B_3$. Identify the column spaces of $-B_1, B_2$ with $\mathcal{T}\mathbf{v}, \mathcal{T}\mathbf{r}$, using the aligned bases. Then B_3 represents $\mathcal{T}(\mathbf{r} - \mathbf{v})$, giving (20)–(21) over the original field. The same lemma applies to the three subspaces $\mathcal{D}_i \le \mathcal{T}$ in the upload proof.

Proof of Proposition 2: Write $\mathcal{W}_i = U_i / A_i$ for the demand quotient after conditioning on User i's side information. In the fixed quotient coordinates, $F/A_i = \bar{K}_i \oplus \mathcal{Y}_i$, where $\mathcal{Y}_i = (\mathcal{P}_j \oplus \mathcal{P}_k) \oplus \mathcal{T}$. By the modular law, $K_i^{\mathrm{act}} = \mathcal{W}_i \cap \bar{K}_i$ has dimension $d_{\mathrm{pair},i}$. Choose a complement K_i^{pas} in $\bar{K}_i$ and put $\mathcal{Q}_i = \pi_{\mathcal{Y}_i} \mathcal{W}_i$. Then

$$\mathcal{W}_i = K_i^{\mathrm{act}} \oplus \{q + \psi_i(q) : q \in \mathcal{Q}_i\}$$

for some linear $\psi_i : \mathcal{Q}_i \to K_i^{\mathrm{pas}}$. Extend ψ_i to $\mathcal{Y}_i$, and lift it to $N_i : F \to K_i^{\mathrm{pas}}$ vanishing on $A_i + \bar{K}_i$. Since $N_i(A_h) \subseteq A_h$ and $N_i N_h = 0$ for all i, h, the invertible map $I - \sum_i N_i$ preserves every side-information space and sends all three demand quotients to $K_i^{\mathrm{act}} \oplus \mathcal{Q}_i$.

Each K_i^{pas}, and each complement of $\sum_{i\neq j}\pi_{\mathcal{P}_j}L_i$ in $\mathcal{P}_j$, is now a side-only direct summand. Their deletion therefore preserves every $\mathcal{R}_L$. ∎

2) Independence of the coordinate choices: After removing the common side-information intersection, put $\mathcal{K}_{\mathrm{pair}} = \sum_{i<j}(A_i \cap A_j)$, $\bar{F} = F/\mathcal{K}_{\mathrm{pair}}$, $\bar{A}_i = (A_i + \mathcal{K}_{\mathrm{pair}})/\mathcal{K}_{\mathrm{pair}}$, and $\bar{U}_i = (U_i + \mathcal{K}_{\mathrm{pair}})/\mathcal{K}_{\mathrm{pair}}$. The subspace $\mathcal{V}_{\mathrm{tri}} = \sum_i[\bar{A}_i \cap (\bar{A}_j + \bar{A}_k)]$ does not depend on the coordinates. In each chosen normal form it is represented by the coefficient space $\mathcal{T} \otimes E$. Its lift to F need not be unique. The images of the private spaces form a direct sum in $\bar{F}/\mathcal{V}_{\mathrm{tri}}$. The quotients of $\bar{U}_i \cap \mathcal{V}_{\mathrm{tri}}$ by $\bar{A}_i \cap \mathcal{V}_{\mathrm{tri}}$ determine the subspaces $\mathcal{D}_i \leq \mathcal{T}$ under the fixed quotient identifications.

An invertible map preserving the three local subspaces $\mathcal{T}\otimes\ell_i$ changes the basis of $\mathcal{T}$ in the same way for all three users. Changing the private lifts adds terms in $\mathcal{T} \otimes \ell_j$. Thus the $\mathcal{D}_i$ undergo a common basis change, while G_0 changes only by source and target basis changes: the added term hB vanishes on $\ker B$. Hence $d_{\mathrm{com}}, d_{\mathrm{pure}}, d_{\mathrm{mix}}$ are invariant. The private projection ranks and (19) give the invariance of d_{priv} and each $d_{\mathrm{pair},i}$.

Lemma 18 (Removal of passive three-line blocks)**.** *Adding locally known block coordinates to private coordinates can separate a passive summand $\widetilde{N}_\varnothing \otimes E$, where $\widetilde{N}_\varnothing$ lifts a complement of $\operatorname{im} G_0$ in $\mathcal{T}_\varnothing$. Removing this summand preserves every $\mathcal{R}_L$ and leaves only the d_{mix}-dimensional lift $W_\varnothing$ outside $\mathcal{D}_1 + \mathcal{D}_2 + \mathcal{D}_3$.*

Proof: Choose $\mathcal{T}_\varnothing = \operatorname{im} G_0 \oplus N_\varnothing$ and lift the two summands to $W_\varnothing, \widetilde{N}_\varnothing \leq \mathcal{T}$, complementary to $\mathcal{D}_1+\mathcal{D}_2+\mathcal{D}_3$. Put $G_N = \pi_{N_\varnothing} G$. Since G_N vanishes on $\ker B$, we can write $G_N = hB$ for some $h : \bigoplus_j \mathcal{P}_j \to N_\varnothing$. Lift each $-h|_{\mathcal{P}_j}$ to $\widetilde{N}_\varnothing$. The resulting coordinate change preserves each side-information space and removes all demand terms in this lift. Thus $\widetilde{N}_\varnothing \otimes E$ is side-only and can be removed. ∎

B. Activity and Upload Additivity

The active pair spaces and $C \otimes E$ form a direct summand respecting every A_i and U_i: Proposition 2 separates the pair spaces, and the lifts of ϕ_i may be chosen with zero C-component. This gives (30)–(31).

Choose the complements and coordinate changes separately on each summand. Then $B = \bigoplus_a B_a$, $G = \bigoplus_a G_a$, and

$$G_0 = \bigoplus_a G_{0,a}, \qquad d_{\mathrm{mix}} = \sum_a \operatorname{rk} G_{0,a}. \tag{85}$$

The other rank parameters are additive as well. Their nonnegativity implies $d_{\mathrm{com}} = d_{\mathrm{pair},i} = 0$ in each residual component. Each component is also active: a side-only summand would split from the full residual task, contradicting (23) or the retained dimension $\dim \mathcal{T}_{\mathrm{res}} = d_{\mathrm{tri}}$. Coordinate invariance makes these conclusions independent of the chosen complements.

Proof of additivity in Proposition 9: Apply Theorem 3 and the preceding rank identities. On each active residual component, twice the ambient dimension minus the sum of the side-information dimensions equals the private dimension plus the number of retained three-line blocks, and therefore equals the minimum upload cost. Ambient and side-information dimensions add under direct sums, proving (56). ∎

Appendix B
Structural Details of the Minimum-Upload Proof

A. Common Functions from Signal-Space Intersections

Using $\mathcal{T}_{\mathrm{res}}$ fixed in Section III-A, set $B_0 = (\mathcal{D}_1 + \mathcal{D}_2 + \mathcal{D}_3) \cap \mathcal{T}_{\mathrm{res}}$ and $\overline{\mathcal{D}}_i = \mathcal{D}_i \cap B_0$. Since $C \leq \mathcal{D}_i$,

$$\begin{aligned} &\mathcal{D}_i = C \oplus \overline{\mathcal{D}}_i, \\ &B_0 = \overline{\mathcal{D}}_1 + \overline{\mathcal{D}}_2 + \overline{\mathcal{D}}_3, \qquad \dim B_0 = d_{\mathrm{pure}}. \end{aligned} \tag{86}$$

Use the retained lift $W_\varnothing$ of $\operatorname{im} G_0$ fixed in Section III-A, so that $\mathcal{T}_{\mathrm{res}} = B_0 \oplus W_\varnothing$. Thus $\dim W_\varnothing = d_{\mathrm{mix}}$ and $\dim \mathcal{T}_{\mathrm{res}} = d_{\mathrm{tri}}$.

For a subspace B of a coefficient space F, write $B^\perp = \{\ell \in F^* : \ell|_B = 0\}$, where F^* is the dual space of F.

Proof of Lemma 6: The residual task has $F_{\mathrm{res}} = \bigoplus_i \mathcal{P}_i \oplus (\mathcal{T}_{\mathrm{res}} \otimes E)$ and $A_{i,\mathrm{res}} = \mathcal{P}_i \oplus (\mathcal{T}_{\mathrm{res}} \otimes \ell_i)$. We prove (45) by contradiction. If the inclusion fails, there is an $f \in F_{\mathrm{res}}^*$ that vanishes on $U_{3,\mathrm{res}} + (U_{1,\mathrm{res}} \cap U_{2,\mathrm{res}})$ but not on $\mathcal{T}_{\mathrm{res}} \otimes E$. The identity $(U_{1,\mathrm{res}} \cap U_{2,\mathrm{res}})^\perp = (U_{1,\mathrm{res}})^\perp + (U_{2,\mathrm{res}})^\perp$ gives $f = f_1 + f_2$, with $f_i|_{U_{i,\mathrm{res}}} = 0$. For $d \in \mathcal{T}_{\mathrm{res}}$, the side-information constraints give $f_1(d\mathbf{v}) = f_2(d\mathbf{r}) = 0$ and $f(d\mathbf{v}) = f(d\mathbf{r}) =: \psi(d)$. The functional ψ on $\mathcal{T}_{\mathrm{res}}$ is nonzero. Applying f_1, f_2, f to the pure demand representatives $d\mathbf{r}, d\mathbf{v}, d\mathbf{v}$, respectively, shows that ψ vanishes on each $\overline{\mathcal{D}}_i$, and hence on B_0.

Set $\alpha_1 = f_2|_{\mathcal{P}_1}$, $\alpha_2 = f_1|_{\mathcal{P}_2}$, and $\alpha_3 = f_1|_{\mathcal{P}_3} = -f_2|_{\mathcal{P}_3}$. The last identity follows from $\mathcal{P}_3 \leq U_{3,\mathrm{res}}$. The maps ϕ_i here are taken after the passive-summand reduction, so G takes values in $\operatorname{im} G_0$. For $\xi_i \in L_i$, apply f_1, f_2, f to the three demand representatives specified by ξ_i and $\phi_i(\xi_i)$ in quotient coordinates $\mathbf{r}, \mathbf{v}, \mathbf{v}$. We obtain

$$\begin{aligned} \psi\phi_1(\xi_1) &= -\alpha_2(\xi_{12}) - \alpha_3(\xi_{13}), \\ \psi\phi_2(\xi_2) &= -\alpha_1(\xi_{21}) + \alpha_3(\xi_{23}), \\ \psi\phi_3(\xi_3) &= -\alpha_1(\xi_{31}) - \alpha_2(\xi_{32}). \end{aligned}$$

Because ψ vanishes on B_0, these equalities do not depend on the chosen representatives of $\phi_i(\xi_i)$. It also defines a functional on $\mathcal{T}_{\mathrm{res}}/B_0 \cong \operatorname{im} G_0$, again denoted by ψ. Adding the three equalities gives

$$\psi G\xi = -(\alpha_1, \alpha_2, \alpha_3)B\xi. \tag{87}$$

Restricting to $\ker B$ gives $\psi G_0 = 0$. Thus ψ vanishes on both B_0 and the lift $W_\varnothing$ of $\operatorname{im} G_0$, which span $\mathcal{T}_{\mathrm{res}}$. This contradicts $\psi \neq 0$. Relabeling the users proves the cyclic inclusions.

Now define

$$\Theta_{12} = X^{\mathsf{T}}(U_{1,\mathrm{res}} \cap U_{2,\mathrm{res}}). \tag{88}$$

Given $(M, \mathbf{T}_i^{\mathrm{res}}, \Omega)$, first recover $\mathbf{P}_i$ by (44). User i's decoder then recovers its residual demand and hence $X^{\mathsf{T}}U_{i,\mathrm{res}}$. Therefore

$$H(\Theta_{12} \mid M, \mathbf{T}_i^{\mathrm{res}}, \Omega) = 0, \qquad i = 1, 2. \tag{89}$$

Likewise, $(M, \mathbf{T}_3^{\mathrm{res}}, \Omega)$ determines $X^{\mathsf{T}}U_{3,\mathrm{res}}$. The inclusion (45) then gives

$$H(\mathbf{T}_1^{\mathrm{res}} \mid M, \Theta_{12}, \mathbf{T}_3^{\mathrm{res}}, \Omega) = 0. \tag{90}$$

Apply Lemma 4. Its two retained terms are J_1, J_2: $(M_i, \mathbf{T}_i^{\mathrm{res}}, \Omega)$ and $(\mathbf{P}_i, \mathbf{T}_i^{\mathrm{res}}, \Omega)$ determine each other. The cyclic intersections give the remaining bounds in (46), for nonlinear uploads and every blocklength. ∎

B. Residual Upload Construction

Lemma 19 (Completion of residual uploads)**.** *After the reduction in Lemma 18, a side-information-preserving coordinate change makes the residual demands feasible with the locally generated relay subspace*

$$\Big(\bigoplus_i \mathcal{P}_i\Big) \oplus P_{\mathrm{row}} \oplus W_{\varnothing} \otimes \langle \mathbf{v} \rangle, \qquad \dim P_{\mathrm{row}} = d_{\mathrm{pure}}, \quad (91)$$

Here $W_{\varnothing}$ is the retained lift, and $P_{\mathrm{row}} \le B_0 \otimes E$ has a locally available basis. The upload dimension is $U_{\mathrm{res}} = d_{\mathrm{priv}} + d_{\mathrm{pure}} + d_{\mathrm{mix}}$, attained in one instance. The coordinate change fixes the coefficient spaces of the three-line blocks and the pairwise-known blocks pointwise.

Proof: Write $(\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3) = (\mathbf{v}, \mathbf{r}, \mathbf{r}-\mathbf{v})$. Use a coordinate change that sends $p \in \mathcal{P}_j$ to $p + g_j(p)\mathbf{v}_j$, where $g_j : \mathcal{P}_j \to B_0 \oplus W_{\varnothing}$. Its coefficients in the three quotient coordinates are

$$(\sigma_{ij}) = \begin{pmatrix} 0 & 1 & 1 \\ 1 & 0 & -1 \\ 1 & 1 & 0 \end{pmatrix}.$$

Thus it changes the maps ϕ_i by

$$\phi_i(x) \longmapsto \phi_i(x) + \sum_{j \ne i} \sigma_{ij} g_j(\pi_{\mathcal{P}_j} x) \pmod{\mathcal{D}_i}. \quad (92)$$

Every off-diagonal entry is nonzero over every field. To cancel one coordinate of ϕ_i, extend its negative from L_i to $\mathcal{P}_j \oplus \mathcal{P}_k$ and write the extension as $h_j + h_k$. Then choose $g_j = \sigma_{ij}^{-1} h_j$ and $g_k = \sigma_{ik}^{-1} h_k$.

Apply the three-subspace decomposition [21, Lem. 2] to $(\overline{\mathcal{D}}_1, \overline{\mathcal{D}}_2, \overline{\mathcal{D}}_3)$ in B_0. An s-dimensional direct summand S in this decomposition corresponds to the coefficient space $S \otimes E$, which has dimension $2s$. A basis of S expresses this space as s three-line blocks. On each one-dimensional summand S, the users with a pure demand form a nonempty proper subset $\mathcal{I}$. Choose a user $h \notin \mathcal{I}$ to upload its local symbol and use the extension above to cancel the projection of ϕ_h onto S. The other users recover both source coordinates on $S \otimes E$, satisfying the corresponding parts of their demands.

On a two-dimensional summand S, choose a basis $\mathbf{a}, \mathbf{b}$ so that $\overline{\mathcal{D}}_i \cap S$ equals $\langle \mathbf{a} \rangle, \langle \mathbf{b} \rangle, \langle \mathbf{a} + \mathbf{b} \rangle$, for Users 1, 2, and 3, respectively. Upload the combinations with coefficient vectors $\mathbf{a} \otimes \mathbf{r}$ and $\mathbf{b} \otimes \mathbf{v}$. User 3 knows those with coefficient vectors $\mathbf{a} \otimes (\mathbf{r}-\mathbf{v})$ and $\mathbf{b} \otimes (\mathbf{r}-\mathbf{v})$, so it recovers all four source coordinates. Users 1 and 2 can lack only the $\mathbf{b}$- and $\mathbf{a}$-components of ϕ_1 and ϕ_2, respectively. Cancel them by extensions to $\mathcal{P}_2 \oplus \mathcal{P}_3$ and $\mathcal{P}_1 \oplus \mathcal{P}_3$. The two corrections use different target directions, so they are compatible even on their shared source $\mathcal{P}_3$. This uses two uploaded symbols.

On $W_{\varnothing}$, upload the d_{mix} linear combinations specified by a basis of $W_{\varnothing} \otimes \langle \mathbf{v} \rangle$. Users 2 and 3 recover the coordinates on $W_{\varnothing} \otimes E$. Cancel the $W_{\varnothing}$-component of ϕ_1 by an extension to $\mathcal{P}_2 \oplus \mathcal{P}_3$.

The corrections have independent target coordinates and combine into $I + N$. Since N maps private coefficient vectors into $(B_0 \oplus W_{\varnothing}) \otimes E$ and vanishes there, $N^2 = 0$. Thus $I + N$ is invertible and preserves every A_i. The uploads supported on $B_0 \otimes E$ generate a d_{pure}-dimensional space P_{row}. The three summands in (91) therefore have total dimension U_{res}. Taking the inverse image gives a locally generated, feasible relay subspace of the same dimension for the original task. ∎

C. Original-Field Omniscience Construction

Proof of Lemma 7: Recall that $d_{\mathrm{om},i} = d_{\mathrm{com}} + d_{\mathrm{pair},i}$ is the dimension missing at User i in the omniscience component, and let r_i be User i's allocated upload rate in symbols per instance. The pair constraints give $\sum_i r_i \ge \max\{\max_i d_{\mathrm{om},i}, \frac{1}{2}\sum_i d_{\mathrm{om},i}\} = U_{\mathrm{om}}$. If $2\max_i d_{\mathrm{om},i} \le \sum_i d_{\mathrm{om},i}$, choose $r_i = (d_{\mathrm{om},j} + d_{\mathrm{om},k} - d_{\mathrm{om},i})/2$, which attains the half-sum. Otherwise, for $d_{\mathrm{om},i} \ge d_{\mathrm{om},j} + d_{\mathrm{om},k}$, choose $r_i = 0$, $r_j = d_{\mathrm{om},k}$, and $r_k = d_{\mathrm{om},i} - d_{\mathrm{om},k}$. These rates are feasible and attain $d_{\mathrm{om},i}$. The chosen rates are integral unless the optimal half-sum is nonintegral; in that case, two instances give integer upload lengths.

For a feasible integer rate triple, write the one-instance blocks in (34) as row vectors: $\mathbf{T}_i^{\mathrm{om}} \in \mathbb{F}_q^{1 \times d_{\mathrm{com}}}$ and $\mathbf{K}_i \in \mathbb{F}_q^{1 \times d_{\mathrm{pair},i}}$, with $\mathbf{T}_3^{\mathrm{om}} = \mathbf{T}_2^{\mathrm{om}} - \mathbf{T}_1^{\mathrm{om}}$. First choose $B_i \in \mathbb{F}_q^{r_i \times d_{\mathrm{com}}}$ so that stacking any two gives column rank d_{com}. Put $t_i = \max\{d_{\mathrm{com}} - r_i, 0\}$, so $t_i + t_j \le d_{\mathrm{com}}$. Temporarily order the indices so that $t_1 \ge t_2 \ge t_3$, and define

$$\begin{aligned} \mathcal{N}_1 &= \langle \mathbf{e}_1, \dots, \mathbf{e}_{t_1} \rangle, \\ \mathcal{N}_2 &= \langle \mathbf{e}_{t_1+1}, \dots, \mathbf{e}_{t_1+t_2} \rangle, \\ \mathcal{N}_3 &= \langle \mathbf{e}_a + \mathbf{e}_{t_1+a} : 1 \le a \le t_3 \rangle. \end{aligned}$$

The $\mathcal{N}_i$ have dimensions t_i and pairwise zero intersections. Choose B_i with kernel $\mathcal{N}_i$, pad with zero rows to reach r_i rows, and restore the index order. This construction works over every field, including when some dimensions are zero.

User j transmits the row vector $M_j = \mathbf{T}_j^{\mathrm{om}} B_j^{\mathsf{T}} + \sum_{i \ne j} \mathbf{K}_i C_{ji}^{\mathsf{T}}$, where $C_{ji} \in \mathbb{F}_q^{r_j \times d_{\mathrm{pair},i}}$ are chosen below. After subtracting locally known terms and transposing the received vectors, the coefficient matrices for the d_{com} missing coordinates in the three-line blocks at Users 1, 2, and 3 are

$$\begin{pmatrix} B_2 \\ B_3 \end{pmatrix}, \qquad \begin{pmatrix} B_1 \\ -B_3 \end{pmatrix}, \qquad \begin{pmatrix} B_1 \\ B_2 \end{pmatrix}.$$

Each has column rank d_{com}. For each i, append $d_{\mathrm{pair},i}$ independent columns by selecting the stacked matrix $[C_{ji}^{\mathsf{T}}, C_{ki}^{\mathsf{T}}]^{\mathsf{T}}$. Such a completion exists because $r_j + r_k \ge d_{\mathrm{com}} + d_{\mathrm{pair},i}$. The choices are independent across i, since $\mathbf{K}_i$ is unknown only at User i. Each user therefore recovers its missing block coordinates and $\mathbf{K}_i$, hence the whole component.

Apply this construction at the chosen optimal rates. If the rates are half-integral, concatenate two copies of each coordinate block and use the integer parameters $2d_{\mathrm{com}}, 2d_{\mathrm{pair},i}, 2r_i$. The resulting two-instance construction uses only the original field. ∎

APPENDIX C
PROOF OF THE RESIDUAL RELAY CHOICES

We prove Proposition 10 by identifying the component dimensions, constructing upload-optimal relay subspaces, and verifying their profiles. We also establish the required choices over the defining finite field: existence over an algebraic closure alone does not ensure a relay subspace over the defining field.

A. Residual Dimensions and Descent

The inclusions $A_i \subseteq U_i \subseteq F$ form the three arms $A_i \to U_i \to F$ of the $\widetilde{E}_6$ quiver. For residual summands, all arrows are injective, $F = \sum_i A_i$, and the side-information spaces are pairwise disjoint. A representation is *absolutely indecomposable* if it remains indecomposable over an algebraic closure; its dimension vector is then a positive root by Kac's theorem [32].

Write the dimension vector as $\alpha = (f; \boldsymbol{u}; \boldsymbol{a})$, where $f = \dim F$, $u_i = \dim U_i$, $a_i = \dim A_i$, $\boldsymbol{u} = (u_i)_{i=1}^3$, and $\boldsymbol{a} = (a_i)_{i=1}^3$. Set $a_\Sigma = \sum_i a_i$ and $\chi = a_\Sigma - f$. The Tits form and null root are

$$q_{\mathrm{T}}(\alpha) = f^2 - f\sum_i u_i + \sum_i (a_i^2 - a_i u_i + u_i^2), \tag{93}$$

$$\delta = (3; 2, 2, 2; 1, 1, 1). \tag{94}$$

Completing squares gives

$$\begin{aligned} q_{\mathrm{T}}(\alpha) = &\sum_i \left(u_i - \frac{f + a_i}{2}\right)^2 \\ &+ \frac{3}{4}\sum_i \left(a_i - \frac{a_\Sigma}{3}\right)^2 + \frac{\chi^2}{4} \end{aligned} \tag{95}$$

This identity holds in every characteristic. Kac's theorem gives $q_{\mathrm{T}} \in \{0, 1\}$ for absolutely indecomposable representations of this affine quiver. The case $q_{\mathrm{T}} = 0$ forces $\alpha = r\delta$, $r \geq 1$, and $\chi = 0$.

Lemma 20 (Positive-excess components over the original field)**.** *Let $k = \mathbb{F}_q$, and let $\mathsf{M} = (A_i \subseteq U_i \subseteq F)_{i=1}^3$ be an indecomposable representation of the inward $\widetilde{E}_6$ quiver over k, with injective arrows and $F = \sum_i A_i$. If $\chi(\mathsf{M}) = \sum_i \dim A_i - \dim F > 0$, then M is absolutely indecomposable and its dimension vector is a real root.*

Proof: Choose a finite extension $\mathbb{K}/k$ such that the scalar extension $\mathsf{M}_{\mathbb{K}}$ splits into absolutely indecomposable summands. In each summand, the arrows remain injective and the side-information spaces span the ambient space, so its excess is nonnegative. Since excess is additive, some summand N has positive excess. Its dimension vector β is a real root by (95).

For every finite field, a real root has exactly one absolutely indecomposable isomorphism class [33, Thm. 1(a)]. Thus there is a k-representation N_0 with $\mathsf{N} \cong (\mathsf{N}_0)_{\mathbb{K}}$. Now regard the $\mathbb{K}$-vector spaces as k-vector spaces. Then $\mathsf{N}_0^{\oplus e}$ is a direct summand of $\mathsf{M}^{\oplus e}$, where $e = [\mathbb{K} : k]$. Since both M and N_0 are indecomposable over k, Krull–Schmidt gives $\mathsf{M} \cong \mathsf{N}_0$. ∎

By Lemma 20 and (95), every positive-excess indecomposable has $q_{\mathrm{T}} = 1$ and $\chi \in \{1, 2\}$. If $\chi = 2$, all other squares vanish: the a_i are equal and $u_i = (f + a_i)/2$, giving $R_2 + n\delta$. If $\chi = 1$, then $\sum_i (a_i - a_\Sigma/3)^2 \leq 1$, so the integers a_i differ by at most one. Classifying by $f \bmod 3$ and choosing the nearest integers to $(f + a_i)/2$ gives the other rows of Table II, up to a permutation of the arms. Here R_0, R_1, R_2 label base dimension vectors; R_n denotes the task defined in Appendix C-C.

TABLE II
BASE DIMENSION VECTORS FOR POSITIVE-EXCESS REAL-ROOT FAMILIES. ARM PERMUTATIONS ARE UNDERSTOOD; $n_{\min}$ IS THE LEAST n ALLOWED BY THE INCLUSION INEQUALITIES

type	f	$\boldsymbol{u}$	$\boldsymbol{a}$	χ	$n_{\min}$
O	0	$(0,0,0)$	$(0,0,1)$	1	1
H	0	$(0,0,1)$	$(0,0,1)$	1	1
R_0	1	$(0,1,1)$	$(0,1,1)$	1	0
R_1	1	$(1,1,1)$	$(0,1,1)$	1	0
R_2	1	$(1,1,1)$	$(1,1,1)$	2	0
T_0	2	$(1,1,1)$	$(1,1,1)$	1	0
T_1	2	$(1,1,2)$	$(1,1,1)$	1	0
T_2	2	$(1,2,2)$	$(1,1,1)$	1	0
T_3	2	$(2,2,2)$	$(1,1,1)$	1	0

For $\gamma = (f; \boldsymbol{u}; \boldsymbol{a})$, the dimension vector $\gamma + n\delta = (f + 3n; \boldsymbol{u} + 2n\mathbf{1}_3; \boldsymbol{a} + n\mathbf{1}_3)$ satisfies the inclusion inequalities exactly when

$$n \geq n_{\min} := \max\{0, \max_i(a_i - u_i), \max_i(u_i - f)\}. \tag{96}$$

Thus O, H start at $n = 1$. The dimension vector $\gamma + n\delta$ is *admissible* if it admits subspaces $A_i \subseteq U_i \subseteq F$ with $F = \sum_i A_i$. A corresponding component is residual only if it survives the active and omniscience reductions. The case $\chi = 0$, including all imaginary roots, is handled by the argument for independent side information in Section IV-C.

Component families. For $n \geq 1$, the rows of Table I are defined as follows. The three $T_1^{(r)}$ families have $\dim F = 3n + 2$, $\dim A_i = n + 1$, and $\dim U_i = 2n + 1$, except that $\dim U_r = 2n + 2$. The $O^{(r)}$ families have $\dim F = 3n$, $\dim U_i = 2n$, and $\dim A_i = n$, except that $\dim A_r = n + 1$.

The indecomposable R_n has $\dim F = 3n + 2$, $\dim A_i = n + 1$, and $\dim U_i = 2n + 1$ for all i. Appendix C-C gives an explicit model. Write

$$\mathsf{R} \triangleq \{\mathsf{R}_n : n \geq 1\}. \tag{97}$$

The boundary tasks $\mathsf{Z}_{16}, \mathsf{Z}_{24}, \mathsf{Z}_{35}$ have a two-dimensional ambient space and three distinct one-dimensional side-information spaces. Only User 2, 3, or 1, respectively, has $U_i = F$; the other users have $U_i = A_i$.

B. Relay Constructions for the Nonzero-Defect Families

To establish the required rank conditions for each component, define the affine defect $\mathrm{def}(\alpha) = -3f + \sum_i (u_i + a_i)$. Both the defect and χ are invariant under addition of δ. For $T_1, T_2, T_3, R_0, R_1, R_2, O, H$, respectively, the defects are

$$1, 2, 3, 1, 2, 3, 1, 2. \tag{98}$$

For this orientation, the defect is the Euler-form value $\langle \delta, \alpha \rangle$. Positive-defect indecomposables are preinjective and have no self-extensions. Their orbits are therefore open and dense in the representation variety [34, Sec. 7, Lems. 2 and 4; Sec. 3, Lem. 1]. The family $T_0 + n\delta$ has zero defect, so we identify it separately through its endomorphisms in Appendix C-C1.

For the profile calculations, we order the arms as

$$\boldsymbol{u}(T_1 + n\delta) = (2n+2, 2n+1, 2n+1),$$
$$\boldsymbol{u}(T_2 + n\delta) = (2n+2, 2n+2, 2n+1).$$

The other families use the order in Table II; the T_1 and O configurations thus correspond to $T_1^{(1)}$ and $O^{(3)}$ in Table I. Activity and upload additivity follow independently from Appendix A-B.

Lemma 21 (Relays for the nonzero-defect families)**.** *Over any finite field K with $|K| \geq 3$, each of the eight nonzero-defect families, at every admissible dimension vector $\gamma + n\delta$ with $n \geq 1$, has upload-optimal, locally generated relay subspaces over K with the following first-family profile entries. Alternative choices of the indexed arm, together with arm permutations, realize the entries in Table I.*

α	$\boldsymbol{F}_{\text{full}}$	$\boldsymbol{F}_P - \boldsymbol{F}_{\text{full}}$
$T_1 + n\delta$	$n\mathbf{1}_6 + \boldsymbol{e}_1 + \boldsymbol{e}_2$	$\boldsymbol{e}_5$
$T_2 + n\delta$	$n\mathbf{1}_6 + \sum_{j=1}^{4} \boldsymbol{e}_j$	0
$T_3 + n\delta$	$(n+1)\mathbf{1}_6$	0
$R_0 + n\delta$	$n\mathbf{1}_6$	0
$R_1 + n\delta$	$n\mathbf{1}_6 + \boldsymbol{e}_1 + \boldsymbol{e}_2$	0
$R_2 + n\delta$	$n\mathbf{1}_6$	0
$O + n\delta$	$n\mathbf{1}_6 - \boldsymbol{e}_5 - \boldsymbol{e}_6$	$\boldsymbol{e}_5$
$H + n\delta$	$n\mathbf{1}_6$	0

The coordinates follow (50)*. For both full and induced tasks, all nine doubled second-family entries $2Q_{i\ell}$ equal $3n+1$ for T_1, T_2, T_3, $3n$ for R_0, R_1, and $3n-1$ for R_2, O, H.*

Proof: Step 1: construct relay kernels and list the required ranks. Work first over an algebraic closure k. The excess χ is one except for R_2, where it is two. Since the side-information spaces generate F and are pairwise disjoint, the relation space $\mathcal{K} = \ker(A_1 \oplus A_2 \oplus A_3 \xrightarrow{\sum} F)$ has dimension χ and injective projections $\pi_i : \mathcal{K} \to A_i$. Thus the side-information configuration has χ three-line blocks, occupying 2χ coefficient dimensions, and $a_i - \chi$ private coordinates on arm i.

We construct each relay subspace as the common kernel of one or two linear forms. In Table IV, the suffix $/s$ indicates that the form vanishes on U_s; we also construct $R_0/1$. For $\chi = 1$ and arm index s, choose $\ell \in U_s^{\perp}$ with $\ell\pi_j \neq 0$ on $\mathcal{K}$, for some $j \neq s$, and set $P = \ker \ell$. Since $\ell\pi_s = 0$ and $\sum_i \pi_i = 0$, the two maps $\ell\pi_i$, $i \neq s$, are nonzero and negatives of each other. Writing $\widehat{a}_i = \dim \widehat{A}_i$ for $\widehat{A}_i = P \cap A_i$, we have $\widehat{a}_s = a_s$, $\widehat{a}_i = a_i - 1$ for $i \neq s$, and $\mathcal{K} \cap \bigoplus_i \widehat{A}_i = 0$. Consequently $P = \bigoplus_i \widehat{A}_i$ has dimension $f - 1$. Feasibility follows from $U_s \leq P$ and $A_i + P = F$ for $i \neq s$.

For R_2, choose $\ell_1 \in U_1^{\perp}$ and $\ell_2 \in U_2^{\perp}$ such that

$$\operatorname{rk} \begin{bmatrix} \ell_1\pi_2 \\ \ell_2\pi_1 \end{bmatrix} = 2 \quad \text{on } \mathcal{K}, \tag{99}$$

and set $P = \ker \ell_1 \cap \ker \ell_2$. The relations $\ell_1\pi_3 = -\ell_1\pi_2$, $\ell_2\pi_3 = -\ell_2\pi_1$ give $\widehat{a}_i = (n, n, n-1)_i$ and $\mathcal{K} \cap \bigoplus_i \widehat{A}_i = 0$. Hence $P = \bigoplus_i \widehat{A}_i$ has dimension $f - 2$. Feasibility follows from $A_1 + P = \ker \ell_1$, $A_2 + P = \ker \ell_2$, and $A_3 + P = F$. In both constructions, $\dim P = f - \chi$ and $\dim(P \cap U_i) = \widehat{a}_i + u_i - a_i$.

We now state sufficient rank conditions for both the full task $\mathcal{V} = F$ and the induced task $\mathcal{V} = P$. Set $d = \dim \mathcal{V}$. For the induced task, put hats on all side-information and signal spaces, their intersections, and their dimensions, with $\widehat{U}_i = P \cap U_i$ and $\widehat{u}_i = \dim \widehat{U}_i$. Set $m_i = u_i - a_i = \widehat{u}_i - \widehat{a}_i$, the demand dimension remaining after conditioning on User i's side information, and require $U_i + U_j = \mathcal{V}$ for $i \neq j$ and $(U_1 \cap U_2) + U_3 = \mathcal{V}$. The dimension formula gives

$$p_{ij} := \dim(U_i \cap U_j) = u_i + u_j - d,$$
$$t := \dim U_{123} = u_1 + u_2 + u_3 - 2d,$$
$$U_{i(j,k)} = U_i, \qquad U_{ij} + U_{ik} = U_i.$$

We also require the following maximal conditional ranks:

$$\operatorname{rk}(U_{ij} \mid A_j) = \min\{m_j, p_{ij}\},$$
$$\operatorname{rk}(U_{123} \mid A_\ell) = \min\{m_\ell, t\}.$$

Substitution into (14) and (15) yields

$$F_{ijk} = \max\{u_i - a_i,\ d - a_i - a_j\}, \tag{100}$$
$$2Q_{i\ell} = \sum_{j=1}^{3} m_j - \min\{m_\ell, t\}. \tag{101}$$

The configurations below satisfy $0 \leq t \leq \min_i m_i$. Hence all nine doubled second-family entries equal $\sum_i m_i - t$.

Step 2: construct configurations over an algebraic closure. Tables III–IV give the $n = 1$ configurations except $R_0/1$. Row reduction verifies the required ranks in every characteristic. Each form defining a hyperplane relay vanishes on U_s and has relation values $(0, 1, -1)$, up to sign and arm order. For R_2, the determinant in (99) is one.

For $R_0/1$, choose $\theta \neq 0, 1$ and use the following four-dimensional configuration:

$$F = \langle x, y, z, w \rangle, \qquad P = \langle x, y, z \rangle,$$
$$(A_1, A_2, A_3) = (\langle x \rangle, \langle y, w \rangle, \langle z, x + w \rangle),$$
$$U_1 = \langle x, y + z \rangle, \qquad U_2 = \langle y, x + z, w \rangle,$$
$$U_3 = \langle z, x + \theta y, x + w \rangle.$$

The form w^* has relation values $(0, 1, -1)$. In P, the pair intersections are the lines $\langle x + y + z \rangle$, $\langle x + \theta y + \theta z \rangle$, and $\langle x + \theta y + z \rangle$. For the full task, the pair-intersection dimensions are $(1, 1, 2)$. In both tasks, the triple intersection is zero, every conditional rank for a pair intersection is one, and each pair of signal spaces spans the ambient space. Thus the required ranks hold, with $\boldsymbol{F} = \mathbf{1}_6$ and $2\boldsymbol{Q} = 3\mathbf{1}_9$.

For arbitrary $n \geq 1$, append $n - 1$ copies of the block D with $P = F = \langle \mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_3 \rangle$, $A_i = \langle \mathbf{e}_i \rangle$, and

$$U_1 = \langle \mathbf{e}_1, \mathbf{e}_2 + \mathbf{e}_3 \rangle, \quad U_2 = \langle \mathbf{e}_2, \mathbf{e}_1 + \mathbf{e}_3 \rangle,$$
$$U_3 = \langle \mathbf{e}_3, \mathbf{e}_1 + \theta \mathbf{e}_2 \rangle, \qquad \theta \neq 0, 1.$$

For this block, every pair of signal spaces spans F, the triple intersection is zero, and every conditional rank for a pair intersection is one. Appending one copy changes

$$(d, a_i, u_i, p_{ij}, t, m_i) \longmapsto (d{+}3, a_i{+}1, u_i{+}2, p_{ij}{+}1, t, m_i{+}1).$$

Both arguments in each minimum defining a pair-conditional rank increase by one, while the triple-conditional rank remains

TABLE III
BASE-CASE TASK CONFIGURATIONS AT $n = 1$. IN EACH ROW, $F = \langle \mathbf{e}_1, \dots, \mathbf{e}_f \rangle$, THE ENTRIES LIST BASES, AND $U_i = A_i + \langle B_i \rangle$. A DASH DENOTES AN EMPTY LIST.

Type	f	A_1	A_2	A_3	B_1	B_2	B_3
T_1	5	$\mathbf{e}_1, \mathbf{e}_4$	$\mathbf{e}_2, \mathbf{e}_5$	$\mathbf{e}_3, \mathbf{e}_4 + \mathbf{e}_5$	$\mathbf{e}_2 + \mathbf{e}_5, \mathbf{e}_3 - \mathbf{e}_5$	$\mathbf{e}_1 + \mathbf{e}_3 + \mathbf{e}_4$	$\mathbf{e}_1 + \mathbf{e}_2$
T_2	5	$\mathbf{e}_1, \mathbf{e}_4$	$\mathbf{e}_2, \mathbf{e}_5$	$\mathbf{e}_3, \mathbf{e}_4 + \mathbf{e}_5$	$\mathbf{e}_2, \mathbf{e}_3$	$\mathbf{e}_1, \mathbf{e}_3 + \mathbf{e}_4$	$\mathbf{e}_1 + \mathbf{e}_2$
T_3	5	$\mathbf{e}_1, \mathbf{e}_4$	$\mathbf{e}_3, \mathbf{e}_4 + \mathbf{e}_5$	$\mathbf{e}_2, \mathbf{e}_5$	$\mathbf{e}_2 + \mathbf{e}_5, \mathbf{e}_3 + \mathbf{e}_5$	$\mathbf{e}_1, \mathbf{e}_2$	$\mathbf{e}_1 + \mathbf{e}_3, \mathbf{e}_4$
R_1	4	$\mathbf{e}_4$	$\mathbf{e}_1, \mathbf{e}_3$	$\mathbf{e}_2, \mathbf{e}_3 + \mathbf{e}_4$	$\mathbf{e}_1, \mathbf{e}_2$	$\mathbf{e}_2 + \mathbf{e}_4$	$\mathbf{e}_1 + \mathbf{e}_4$
R_2	4	$\mathbf{e}_1, \mathbf{e}_2$	$\mathbf{e}_3, \mathbf{e}_4$	$\mathbf{e}_1 + \mathbf{e}_3, \mathbf{e}_2 + \mathbf{e}_4$	$\mathbf{e}_3 + \mathbf{e}_4$	$\mathbf{e}_2$	$\mathbf{e}_1$
O	3	$\mathbf{e}_3$	$\mathbf{e}_1$	$\mathbf{e}_1 + \mathbf{e}_3, \mathbf{e}_2$	$\mathbf{e}_2$	$\mathbf{e}_2 + \mathbf{e}_3$	—
H	3	$\mathbf{e}_3$	$\mathbf{e}_1$	$\mathbf{e}_1 + \mathbf{e}_3, \mathbf{e}_2$	$\mathbf{e}_2$	$\mathbf{e}_2 + \mathbf{e}_3$	$\mathbf{e}_1$

TABLE IV
RELAY SUBSPACES FOR THE BASE-CASE CONFIGURATIONS IN TABLE III. HERE $\mathbf{e}_j^*(x) = x_j$. THE RELAY SUBSPACE IS THE COMMON KERNEL OF THE LISTED FUNCTIONALS.

Choice	Defining functionals
$T_1/2$	$\mathbf{e}_1^* - \mathbf{e}_4^*$
$T_1/3$	$\mathbf{e}_4^* - \mathbf{e}_5^*$
$T_2/3$	$\mathbf{e}_1^* - \mathbf{e}_2^* - \mathbf{e}_4^* + \mathbf{e}_5^*$
$T_3/2$	$\mathbf{e}_4^* - \mathbf{e}_5^*$
$R_1/1$	$\mathbf{e}_3^*$
R_2	$\mathbf{e}_3^* - \mathbf{e}_4^*,\ \mathbf{e}_1^*$
$O/1$	$\mathbf{e}_1^*$
$O/2$	$\mathbf{e}_2^* - \mathbf{e}_3^*$
$O/3$	$\mathbf{e}_1^* - \mathbf{e}_3^*$
$H/1$	$\mathbf{e}_1^*$

t. Thus the rank conditions still hold. Each copy adds $\mathbf{1}_6$ to $\boldsymbol{F}$ and 31_9 to $2\boldsymbol{Q}$, for both tasks.

Step 3: transfer the ranks to the prescribed component. Fix a dimension vector over an algebraic closure k. The flags $A_i \subseteq U_i \subseteq F$ with $F = \sum_i A_i$ and pairwise disjoint side-information spaces form a nonempty irreducible open subset $\mathcal{X}$ of a product of flag varieties. Adding $\ell \in U_s^{\perp}$, or a pair in $U_1^{\perp} \times U_2^{\perp}$ for R_2, gives a vector bundle over $\mathcal{X}$. The relation-rank conditions in Step 1 give an open locus with fixed full and induced side-information and signal dimensions. For both tasks, first impose the pair-sum conditions of Step 1. On this open locus, the pair intersections have constant dimension and form vector subbundles. The spanning conditions $(U_1 \cap U_2) + U_3 = \mathcal{V}$ for both tasks are then surjectivity conditions on bundle maps; on their open locus, the triple intersections also form vector subbundles. The remaining conditional-rank conditions require maximal rank of the bundle maps $U_{ij} \to U_j/A_j$ and $U_{123} \to U_\ell/A_\ell$, and their induced analogues, so they too are open. The configurations in Step 2 show that their common solution set is nonempty.

The prescribed indecomposable has a dense open orbit in $\mathcal{X}$, obtained from its open representation orbit by taking image flags. Its inverse image in the irreducible parameter bundle meets the open solution set. Transporting the forms by an isomorphism establishes the full-task ranks and nonzero restriction maps for the prescribed component over k. Step 4 selects forms over its defining finite field.

For $n \geq 1$, these indecomposables are neither passive nor omniscience components. Their minimum total upload cost therefore satisfies $U^* = f - \chi = \dim P$ by (56). The $T_1/2, T_1/3$ choices have increments $\boldsymbol{e}_5, \boldsymbol{e}_3$, and the three O choices have increments $\boldsymbol{e}_6, \boldsymbol{e}_5, \boldsymbol{e}_1 + \boldsymbol{e}_3$. Every other choice preserves the fully informed profile. Relabeling the arms gives Table I.

Step 4: choose the forms over the defining finite field. Let K be the defining field, with $Q = |K| \geq 3$. Scalar extension preserves ranks, so the full-task ranks and nonzero restriction maps from Step 3 hold over K. For $r \geq 1$, a union of at most three proper linear subspaces of K^r has size at most $1 + 3(Q^{r-1} - 1) < Q^r$. Thus the choices below exist over K.

The rank identities established in Step 3 give

$$A_i \cap U_{123} = 0, \qquad \dim(A_i \cap U_j) = \max\{0, p_{ij} - m_i\}. \quad (102)$$

Thus each nonzero remote side intersection $A_j \cap U_k$ meets U_s only at zero. Requiring a form to be nonzero there therefore excludes only a proper subspace of $U_s^{\perp}$.

For a hyperplane relay, choose $\ell \in U_s^{\perp}$ with $\ell\pi_j \neq 0$. Step 3 shows that the forms failing this condition form a proper subspace. Set $C_j = A_j \cap U_k$ and $C_k = A_k \cap U_j$, and require ℓ to be nonzero on each nonzero C_j, C_k. If both spaces are zero, require $\ell|_{U_{jk}} \neq 0$ instead. This also excludes a proper subspace, since $\dim U_{jk} - \dim(U_{jk} \cap U_s) = f - u_s > 0$ for every listed arm. In either case, the remote pair loses one dimension. At most three proper subspaces are excluded, so the count above guarantees such an ℓ over K.

The kernel construction gives a locally generated feasible relay subspace with the stated dimension drops. The pairs containing s and the triple intersection U_{123} are unchanged. For the remote pair,

$$\begin{aligned} \mathrm{rk}\left(\widehat{U}_{jk} \mid \widehat{A}_j\right) &= p_{jk} - 1 - \dim C_j + \mathrm{rk}(\ell|_{C_j}) \\ &= \min\{m_j, p_{jk} - 1\}. \end{aligned}$$

Here $\dim C_j = \max\{0, p_{jk} - m_j\}$ by (102); the same calculation applies to k. The other pair-conditional ranks and the triple-conditional ranks are unchanged because their intersections lie in $U_s \subseteq P$. The induced pair dimensions are $\widehat{u}_i + \widehat{u}_j - (f-1)$, and the triple dimension is $\sum_i \widehat{u}_i - 2(f-1)$. Thus the induced task also satisfies the required spanning conditions.

For a relay subspace of codimension two, the family is $R_2 + n\delta$, with $f = 3n+1$, $a_i = n+1$, $u_i = 2n+1$, $p_{ij} = n+1$, $t = 1$, and $m_i = n$. On $U_1^{\perp} \times U_2^{\perp}$, write the relation determinant as the bilinear form

$$B(\ell_1, \ell_2) = \det \begin{bmatrix} \ell_1\pi_2 \\ \ell_2\pi_1 \end{bmatrix}.$$

Step 3 shows that B is nonzero after scalar extension, so it is already nonzero over K. Choose ℓ_1 nonzero on $A_2 \cap U_3$ and $A_3 \cap U_2$, with $B(\ell_1, \cdot) \neq 0$. These conditions exclude three proper subspaces by (102) and $B \neq 0$. Next choose ℓ_2 nonzero on $A_1 \cap U_3$ and $A_3 \cap U_1$, with $B(\ell_1, \ell_2) \neq 0$. Again at most three proper subspaces are excluded, so such an ℓ_2 exists over K.

The nonzero determinant gives $\dim P = 3n-1$, $\widehat{\boldsymbol{a}} = (n, n, n-1)$, and $\widehat{\boldsymbol{u}} = (2n, 2n, 2n-1)$. Both forms vanish on U_{12} and U_{123}, so these spaces keep dimensions $n+1$ and one. Each of U_{13}, U_{23} loses one dimension: one form vanishes there, while the other is nonzero on a specified line inside it. The intersections with both endpoint side-information spaces lose their single dimension. Thus all pair-conditional ranks are n, and all triple-conditional ranks are one. The dimension formulas also verify the required spanning conditions. Both constructions therefore realize the stated fifteen-entry profiles over K, for every n.

■

The case $n = 0$. The canonical T_1 boundary has

$$(A_1, A_2, A_3) = (\langle x\rangle, \langle y\rangle, \langle x+y\rangle),$$
$$(U_1, U_2, U_3) = (A_1, A_2, \langle x, y\rangle).$$

The one-dimensional relay subspaces $\langle x\rangle$ and $\langle y\rangle$ are upload-optimal, preserve every second-family entry, and have increments $\boldsymbol{e}_4$ and $\boldsymbol{e}_2$. Arm relabeling gives the three boundary rows of Table I. For T_2, the same side-information lines with signal spaces (A_1, F, F) admit the profile-preserving relay subspace $P = A_1$. For $n = 0$, T_0, R_0, R_2 are passive, R_1, T_3 are omniscience, and O, H violate the inclusion inequalities. These constructions use the base field; a relay subspace of codimension two is needed only for $n \geq 1$.

C. The R_n Component

The indecomposable with dimension vector $T_0 + n\delta$ is denoted by $A(3n+2, 0)$ in [35, Prop. 4.4 and Sec. 5.3]. We describe it below in terms of the users' side information and demands. The proof is valid over every finite field, including characteristic two.

For $n \geq 1$ over any finite field K, set $A = I_n + N$, where N is a single nilpotent Jordan block. Choose a column $h \in K^{n\times 1} \setminus \operatorname{im} N$ and a row $\ell_0 \in K^{1\times n}$ spanning the left kernel of N, normalized by

$$\ell_0(A - I) = 0, \qquad \ell_0 h = 1. \tag{103}$$

For independent source coordinates $\mathbf{X}_1, \mathbf{X}_2, \mathbf{X}_3 \in K^n$ and $V, R \in K$, put $T = R - V$. Each $\mathbf{X}_i$ is private to User i; V, R, T are the local observations of one three-line block. The side information and demands of R_n are

User	Side information	Demand
1	$(\mathbf{X}_1, V)$	$\mathbf{X}_2 + \mathbf{X}_3$
2	$(\mathbf{X}_2, R)$	$\mathbf{X}_1 + \mathbf{X}_3$
3	$(\mathbf{X}_3, T)$	$\mathbf{X}_1 + A\mathbf{X}_2 + hV$

The five-bit task (1) is the case $n = 1$, $K = \mathbb{F}_2$.

1) Identification over the defining field: Let $E_i : K^n \to F$ embed private block i, and let $\mathbf{v}, \mathbf{r}$ be the coefficient directions of the scalar coordinates V, R. Put $J = N^{\mathsf{T}}$; then J is a single nilpotent Jordan block with $h^{\mathsf{T}}|_{\ker J} \neq 0$. The side-information spaces force each endomorphism Φ to act by a common scalar α on $\langle \mathbf{v}, \mathbf{r}\rangle$, since it preserves the three lines $\langle \mathbf{v}\rangle$, $\langle \mathbf{r}\rangle$, and $\langle \mathbf{r} - \mathbf{v}\rangle$. Preserving U_1, U_2 then requires

$$\begin{aligned} \Phi E_1 x &= E_1 B x + \mathbf{v} g x, \\ \Phi E_2 x &= E_2 B x - \mathbf{r} g x, \\ \Phi E_3 x &= E_3 B x + (\mathbf{r} - \mathbf{v}) g x, \end{aligned} \tag{104}$$

where $B \in K^{n\times n}$ and $g \in K^{1\times n}$. Preservation of U_3 gives

$$BJ = JB, \qquad h^{\mathsf{T}}(B - \alpha I) = -gJ. \tag{105}$$

Since J is a single Jordan block, $BJ = JB$ gives $B = p(J)$. Restricting the second identity to $\ker J$ gives $p(0) = \alpha$. Hence $\Phi - \alpha I$ is block triangular with nilpotent diagonal blocks, and is therefore nilpotent. Every endomorphism is thus a scalar plus a nilpotent map. The same calculation holds after any scalar extension, so the endomorphism criterion in [33, Lem. 2] proves absolute indecomposability in every characteristic.

Its dimension vector is

$$(3n+2; 2n+1, 2n+1, 2n+1; n+1, n+1, n+1) = T_0 + n\delta,$$

with $q_{\mathrm{T}} = 1$. Real-root uniqueness identifies this model with the required family over an algebraic closure [32] and, by Lemma 20 and [33, Thm. 1(a)], over the original finite field.

2) Three upload-optimal relay subspaces: Indecomposability and $\dim A_i = n+1 < \dim U_i = 2n+1 < \dim F = 3n+2$ exclude passive and omniscience summands. Thus R_n is active residual, and (56) gives $U^*(\mathsf{R}_n) = 3n+1$. Consider

$$\begin{aligned} P_V &= \langle \mathbf{X}_1, V, \mathbf{X}_2, \mathbf{X}_3\rangle, \\ P_R &= \langle \mathbf{X}_1, \mathbf{X}_2, R, \mathbf{X}_3\rangle, \\ P_T &= \langle \mathbf{X}_1 + hV, \mathbf{X}_2, \mathbf{X}_3, T\rangle. \end{aligned} \tag{106}$$

Each relay subspace has $3n+1$ independent local generators. All demands have coefficient vectors in P_V. Given the relay input specified by P_R, User 3 obtains $V = R - T$. With P_T, User 2 uses the same identity to recover $\mathbf{X}_1$, and User 3 forms $(\mathbf{X}_1 + hV) + A\mathbf{X}_2$. The remaining demands are recovered directly. Thus all three relay subspaces are feasible and upload-optimal.

3) Broadcast profiles: The fully informed and induced profiles are

$$\begin{aligned} \boldsymbol{F}_{\mathrm{full}} &= n\mathbf{1}_6, & 2\boldsymbol{Q}_{\mathrm{full}} &= 3n\mathbf{1}_9, \\ \boldsymbol{e}(P_V) &= \boldsymbol{e}_4 + \boldsymbol{e}_6, & \boldsymbol{e}(P_R) &= \boldsymbol{e}_2 + \boldsymbol{e}_5, \\ \boldsymbol{e}(P_T) &= \boldsymbol{e}_1 + \boldsymbol{e}_3. \end{aligned} \tag{107}$$

All three relay subspaces preserve $\boldsymbol{Q}$: $2\boldsymbol{Q}_P = 3n\mathbf{1}_9$.

For the rank calculation, suppose $m_i = m$, $U_{123} = 0$, and the image of each pair sum $U_{ij} + U_{ik}$ spans the demand quotient U_i/A_i. Since $U_{ij} + U_{ik} \subseteq U_{i(j,k)} \subseteq U_i$, both spaces have conditional rank m. The LCBC profile entries become

$$F_{ijk} = 2m - \mathrm{rk}(U_{ij} \mid A_j), \qquad 2Q_{i\ell} = 3m. \tag{108}$$

To check the hypotheses, write coefficient rows as $(x_1, x_2, x_3; v, r)$. The full pair spaces are $U_{ij} = \{B_{ij}(w) : w \in K^{1\times n}\}$, where

$$\begin{aligned} B_{12}(w) &= (w, w, w; 0, 0), \\ B_{13}(w) &= (w, wA, wA; wh, 0), \\ B_{23}(w) &= (w, wA, w; 0, wh). \end{aligned} \tag{109}$$

Each pair space has dimension n and intersects both incident side-information spaces only at zero, since A is invertible. In the triple intersection, $w(A-I) = 0$ and $wh = 0$ force $w = 0$ by (103), so $U_{123} = 0$.

The relay hyperplanes are $P_V : \ r = 0$, $P_R : \ v = 0$, and $P_T : \ v + r = x_1 h$. They impose $wh = 0$ only on U_{23}, U_{13}, U_{12}, respectively, reducing that pair dimension to $n-1$. All intersections with incident side-information spaces remain trivial, and each user retains an unchanged n-dimensional pair space whose image spans its demand quotient. Applying (108) with $m = n$ before and after restriction gives (107), including preservation of all nine second-family entries.

For the positive T_1-, O-, and R families, the increment choices are independent of n, while each pair margin increases by $n-1$ from its $n = 1$ value. This follows from Lemma 21 and (107), and proves the lower bounds used in (59).

D. Completion of the Relay-Choice Proof

Completion of the proof of Proposition 10: Zero-excess components are profile-preserving by Section IV-C. For positive excess, Table II is exhaustive. Lemma 21, the boundary constructions, and the R_n relay subspaces cover every row. With arm permutations, they give all stated increments, preserve all nine second-family entries, and preserve the fully informed profile of every unlisted residual component. ∎

Appendix D
Allocation and Directed-Edge Extraction

A. Proof of the Aggregate Selection Criterion

Proof of Proposition 11: Every relay choice for R uses one unit of capacity in each orientation, so (65) is necessary. To prove sufficiency, we first allocate the relay choices for R and then the other types. Set $\boldsymbol{a} = (s_1, s_4, s_5)$ and $\boldsymbol{b} = (s_6, s_2, s_3)$. The three relay choices for R $\boldsymbol{e}_1 + \boldsymbol{e}_3$, $\boldsymbol{e}_2 + \boldsymbol{e}_5$, $\boldsymbol{e}_4 + \boldsymbol{e}_6$ have respective capacities

$$c_{\mathsf{R},1} = \min\{a_1, b_3\}, \quad c_{\mathsf{R},2} = \min\{a_3, b_2\}, \quad c_{\mathsf{R},3} = \min\{a_2, b_1\}.$$

By (64), $a_i + b_i \geq \mu_{\mathsf{R}}$. Expanding the minima gives $\sum_i c_{\mathsf{R},i} \geq \mu_{\mathsf{R}}$: the sums using only entries of $\boldsymbol{a}$ or only entries of $\boldsymbol{b}$ satisfy (65), while every mixed sum contains a pair $a_i + b_i \geq \mu_{\mathsf{R}}$. Assign weights to the choices in this order, up to their capacities, to obtain $0 \leq \lambda_{\mathsf{R},i} \leq c_{\mathsf{R},i}$ with $\sum_i \lambda_{\mathsf{R},i} = \mu_{\mathsf{R}}$.

These choices use at most μ_{R} units in each opposite pair. By (64), the remaining pair capacity covers $\sum_{t\neq\mathsf{R}} \mu_t \underline{\kappa}_{ab}(t)$. For each pair, reserve $\mu_t \underline{\kappa}_{ab}(t)$ for each remaining type, allocating capacity from one endpoint and then the other. The reserved capacities suffice for boundary types; a T_1-type needs only μ_t units in its pair, which receives $2\mu_t$. For $t = O^{(1)}$, let c_j be its reserved capacities, so $c_1 + c_6, c_2 + c_4 \geq \mu_t$. Give $\boldsymbol{e}_4 + \boldsymbol{e}_6$ weight $\lambda_{\mathrm{dbl}} = (\mu_t - c_1 - c_2)_+$. Then $\lambda_{\mathrm{dbl}} \leq \min\{c_4, c_6\}$ and $\mu_t - \lambda_{\mathrm{dbl}} \leq c_1 + c_2$, so the remaining weight can be assigned to the singleton choices within their capacities. Relabeling treats the other O-types.

This construction uses only addition, subtraction, minima, and maxima, so it preserves integrality. Allocating weights to summands in order likewise distributes each type's weights among its individual summands. ∎

B. Simultaneous Extraction of Private Demands

We prove the extraction assertion in Lemma 12, including its exact multiplicity formula. The deficit implication then follows from the entrywise profile calculation in (73).

Proof: Put $\mathcal{B} = A_j + U_{ij}$ and $\mathcal{N} = A_k \cap U_j$. The number of extractable edge copies is the dimension of the image of $\mathcal{N}$ in $F/\mathcal{B}$:

$$\mu_{k\to j} = \dim \mathcal{N} - \dim(\mathcal{N} \cap \mathcal{B}). \tag{110}$$

Indeed, the modular law gives $(U_i + A_j) \cap U_j = \mathcal{B}$. Choose a complement W of $\mathcal{N} \cap \mathcal{B}$ in $\mathcal{N}$, then a decomposition $F = W \oplus K_0$ with $U_i + A_j \subseteq K_0$. Projection onto W kills U_i, A_i, A_j and has image in $A_k \cap U_j$, so it preserves every task subspace. Its image is $\mathsf{E}_{k\to j}^{\oplus \dim W}$ and its kernel is a complementary task. The dimension in (110) is nonnegative and additive, equals one on $\mathsf{E}_{k\to j}$, and vanishes on the other five edge types. This proves maximality and permits successive extraction of all six types without changing the remaining multiplicities.

Substitution in (14) and (15) now gives

$$p_{k\to j} = \mu_{k\to j} - \dim \frac{U_{jk} \cap \mathcal{B}}{U_{123} + (A_k \cap \mathcal{B})} \leq \mu_{k\to j}. \tag{111}$$

The quotient is well-defined: $A_k \cap \mathcal{B} = \mathcal{N} \cap \mathcal{B}$, and both denominator spaces lie in $U_{jk} \cap \mathcal{B}$. ∎

Appendix E
Profile Verification for Demand-Cycle Repair

We verify the pair intersections and quotient-spanning properties used in Lemma 13. Append the demand-cycle coordinates to the coefficient rows of R_n given in Appendix C-C. The full pair spaces consist of $B_{ij}(w) + \beta \mathbf{c}_{ij}^{\sigma}$, where $w \in \mathbb{F}_q^{1\times n}$, $\beta \in \mathbb{F}_q$, and $B_{ij}(w)$ is given in (109). The cycle coordinate rows are

σ	$\mathbf{c}_{12}^{\sigma}$	$\mathbf{c}_{13}^{\sigma}$	$\mathbf{c}_{23}^{\sigma}$	restricted pair
$+$	y	z	x	13
$-$	z	x	y	23.

The triple intersection is zero before restriction: the triple intersection in R_n is zero, and the cycle pair spaces use distinct coordinates. It remains zero after restriction.

For P^+, the functional $\ell^+(\xi) = \xi_V + \xi_z$ vanishes on U_{12} and U_{23}; on U_{13}, it imposes $\beta = -wh$. For P^-, the functional $\ell^-(\xi) = \xi_R + \xi_y$ vanishes on U_{12} and U_{13}, and imposes the same constraint on U_{23}. The intersections with side-information spaces and the losses in conditional rank are those computed in the proof of Lemma 13.

To verify quotient spanning, Users 1 and 2 use unchanged pair spaces: U_{12} and U_{23}, respectively, in the positive orientation, and U_{13} and U_{12} in the reverse orientation. Since A is invertible, these pairs span the vector demands, and their cycle coordinate is unrestricted.

In User 3's vector-demand and cycle-demand coordinates, the images of its two incident pairs are $\{(w, -wh) : w \in \mathbb{F}_q^{1\times n}\}$ and $\{(w, 0) : w \in \mathbb{F}_q^{1\times n}\}$, in opposite orders for the two orientations. Since $h \neq 0$, $w \mapsto wh$ is onto $\mathbb{F}_q$, so their sum is $\mathbb{F}_q^{1\times n} \times \mathbb{F}_q$. The pair sums therefore span each demand quotient after restriction, and hence before restriction. Together with (76), this proves (77).

Finally, in the full task, $\mathrm{rk}(U_{ij} \mid A_j) = n$ if $\mathbf{c}_{ij}^{\sigma} \in A_j$, and $\mathrm{rk}(U_{ij} \mid A_j) = n + 1$ otherwise: the pair space in R_n has dimension n and intersects each incident side-information space only at zero. Substitution into (77) yields the two profiles in (79) and the common second-family value $Q_{i\ell} = 3(n+1)/2$.

APPENDIX F
ORIGINAL-FIELD RELAY AND BROADCAST CONSTRUCTIONS

We prove the properties of the induced side-information spaces and the broadcast constructions over the original field used in Section VI.

A. Direct-Sum Structure at Minimum Upload

The following direct-sum property supports the binary broadcast construction.

Lemma 22 (Direct-sum structure at minimum upload)**.** *Let $P \leq F_{\mathrm{res}}^{\oplus L}$ be a feasible, locally generated relay subspace for an active residual task over $\mathbb{F}_q$. If $\dim P = LU^*(\mathsf{M}_{\mathrm{res}})$, then*

$$P = \bigoplus_{i=1}^{3} (P \cap A_{i,\mathrm{res}}^{\oplus L}).$$

Proof: Put $\widehat{A}_{i,\mathrm{res}} = P \cap A_{i,\mathrm{res}}^{\oplus L}$, $d_P = \dim P$, and $\chi_P = \sum_i \dim \widehat{A}_{i,\mathrm{res}} - d_P$. The residual side-information spaces are pairwise disjoint: common and passive pair spaces have been removed, and active pair spaces belong to the omniscience component. Thus the $\widehat{A}_{i,\mathrm{res}}$ are pairwise disjoint and span P. Their three-subspace decomposition (Appendix A) has χ_P three-line blocks and $d_P - 2\chi_P$ private coordinates.

If $\chi_P > 0$, consider two copies of this relay input. The users upload all private coordinates and use Lemma 7 on the three-line blocks. This lets every user recover both copies of the relay input, with total upload length

$$2(d_P - 2\chi_P) + 3\chi_P = 2d_P - \chi_P < 2LU^*(\mathsf{M}_{\mathrm{res}}).$$

Every new upload is a function of the induced side information represented by $\widehat{A}_{i,\mathrm{res}}$, hence of the original local side information. Recovering the relay input satisfies the original demands by (48), contradicting Theorem 3 at blocklength $2L$. Thus $\chi_P = 0$. This auxiliary doubling is only a converse argument; it is not part of the constructed code. ∎

B. Binary Implementations of the Original Relay Choices

Lemma 23 (Binary relay choices)**.** *Every choice in Table I can be implemented over $\mathbb{F}_2$, with the stated first-family increment and unchanged second-family entries.*

Proof: The boundary and R_n relay subspaces work over every field. By user relabeling, it suffices to treat $T_1^{(1)}$ and $O^{(3)}$. Use the hyperplane construction in Lemma 21, with $\{s, j, k\} = [3]$. The linear form $\ell \in U_s^{\perp}$ must satisfy the nonvanishing relation condition in that construction and be nonzero on each nonzero intersection $C_j = A_j \cap U_k$, $C_k = A_k \cap U_j$. By (102), the possible pairs $(\dim C_j, \dim C_k)$, up to interchange, are

family	s	$(\dim C_j, \dim C_k)$
$T_1^{(1)}$	$2, 3$	$(0, 1)$
$O^{(3)}$	$1, 2$	$(0, 1)$
$O^{(3)}$	3	$(0, 0)$.

The last row instead requires $\ell|_{U_{12}} \neq 0$. Thus each choice excludes at most two proper subspaces. Two proper subspaces cannot cover a vector space over any field, so a valid ℓ exists even over $\mathbb{F}_2$. The rank calculations in Lemma 21 then give the stated profiles. ∎

C. Additional Binary Choices and Integer Allocation

For the families that are profile-preserving in Lemma 21, we allow limited first-family increments over $\mathbb{F}_2$ and use the component pair margins to accommodate them. Integer capacities suffice for all the additional choices below.

Lemma 24 (Additional binary relay choices)**.** *For each family in Table V and each listed increment bound $\boldsymbol{d}$, there is an upload-optimal scalar relay subspace over $\mathbb{F}_2$ with unchanged second-family entries and first-family increment at most $\boldsymbol{d}$, coordinatewise. The listed sufficient pair-capacity requirement $\boldsymbol{c} = (c_{16}, c_{24}, c_{35})$ satisfies $c_{ab} \leq \kappa_{ab}$ for $ab = 16, 24, 35$. Every capacity vector $\boldsymbol{b} \in \mathbb{Z}_{\geq 0}^6$ with*

$$(b_1 + b_6, b_2 + b_4, b_3 + b_5) \geq \boldsymbol{c}$$

admits a listed increment bound $\boldsymbol{d} \leq \boldsymbol{b}$. The three remaining cases $T_3 + \delta$, $R_2 + \delta$, and $H + \delta$ admit upload-optimal scalar relays that do not increase any profile entry.

Proof: For the excess-one families, the side-information decomposition contains one three-line block. Let $t_i \in A_i$ be nonzero generators of its three local lines, chosen so that $t_1 + t_2 + t_3 = 0$. Choose $\ell \in U_s^{\perp}$ with $\ell(t_j) \neq 0$, where $\{s, j, k\} = [3]$, and put $P = \ker \ell$. The kernel calculation in Lemma 21 gives

$$\begin{gathered} P = \bigoplus_i (P \cap A_i), \qquad \dim P = f - 1, \\ \widehat{a}_i = a_i - \mathbf{1}\{i \neq s\}. \end{gathered} \tag{112}$$

This relay is feasible and upload-optimal, with unchanged conditional demand dimensions m_i.

For every arm used below, the map $\ell \mapsto \ell(t_j)$ on $U_s^{\perp}$ is nonzero. To check the additional arms, use the primitive configurations from the proof of Lemma 21. The plane $\langle t_1, t_2, t_3 \rangle$ of this block is not contained in U_s for $T_2/1$,

TABLE V
ADDITIONAL BINARY CHOICES. FAMILIES ARE DEFINED IN TABLE II, WITH ARM ORDER AS IN APPENDIX C-B. THE LISTED VECTORS BOUND FIRST-FAMILY INCREMENTS COORDINATEWISE. USER PERMUTATIONS ACT JOINTLY ON THE COMPONENT, PROFILE COORDINATES, AND PAIR-CAPACITY COORDINATES. THE PAIR-CAPACITY REQUIREMENTS ARE SUFFICIENT FOR ALLOCATION.

Family	Range	Increment bounds	Pair-capacity requirement
$R_0 + n\delta$	$n \geq 1$	$\{\boldsymbol{e}_4, \boldsymbol{e}_6, \boldsymbol{e}_1 + \boldsymbol{e}_3, \boldsymbol{e}_2 + \boldsymbol{e}_5\}$	$(1,1,1)$
$T_2 + n\delta$	$n \geq 1$	$\{\boldsymbol{e}_1, \boldsymbol{e}_6\}$	$(1,0,0)$
$R_1 + n\delta$	$n \geq 1$	$\{\boldsymbol{e}_3, \boldsymbol{e}_5\}$	$(0,0,1)$
$T_3 + n\delta$, $H + n\delta$	$n \geq 2$	$\{\boldsymbol{e}_2, \boldsymbol{e}_4\}$	$(0,1,0)$
$R_2 + n\delta$	$n \geq 2$	$\{\boldsymbol{e}_4 + \boldsymbol{e}_5\}$, with user permutations	$(1,1,1)$

$T_3/1,2$, $R_1/2,3$, and $H/1,2$, or for any of the three R_0 arms. Appending the blocks D used there preserves this property. This nonempty open rank condition therefore holds on the dense indecomposable orbit over the algebraic closure. Since scalar extension preserves ranks, the restriction map is also nonzero over $\mathbb{F}_2$.

Set $C_j = A_j \cap U_k$ and $C_k = A_k \cap U_j$. By (102), each nonzero C_j, C_k is a line disjoint from U_s. Choose one of these nonzero lines and require ℓ to be nonzero on it. If both spaces are zero, require $\ell|_{U_{jk}} \neq 0$ instead; this is possible since $\dim U_{jk} - \dim U_{123} = f - u_s > 0$. Together with $\ell(t_j) \neq 0$, either choice excludes only two proper subspaces of $U_s^{\perp}$. Their union cannot cover $U_s^{\perp}$, so such an ℓ exists over $\mathbb{F}_2$.

In either case U_{jk} loses one dimension, while the two pair intersections containing s and the triple intersection are unchanged. The dimension identities in Lemma 21 give, for distinct i, j, k,

$$\widehat{U}_i + \widehat{U}_j = P, \qquad \widehat{U}_{ij} + \widehat{U}_{ik} = \widehat{U}_i. \tag{113}$$

The triple conditional ranks are unchanged as well, so (15) gives $\boldsymbol{Q}_P = \boldsymbol{Q}_{\text{full}}$. For the remote pair,

$$\operatorname{rk}\Big(\widehat{U}_{jk} \mid \widehat{A}_j\Big) = p_{jk} - 1 - \dim C_j + \operatorname{rk}(\ell|_{C_j}).$$

If ℓ vanishes on a nonzero C_j, the resulting rank loss adds one only to F_{kjs}; if it vanishes on a nonzero C_k, it adds one only to F_{jks}. All other first-family entries follow from (100) and (112).

For $R_0 + n\delta$, arm $s = 1$ gives the upper bounds $\boldsymbol{e}_4$ and $\boldsymbol{e}_6$, while arms $s = 2, 3$, for which $(\dim C_j, \dim C_k) = (0,0)$, give $\boldsymbol{e}_2 + \boldsymbol{e}_5$ and $\boldsymbol{e}_1 + \boldsymbol{e}_3$. For $T_2 + n\delta$, arm $s = 3$ gives the upper bound $\boldsymbol{e}_1$; arm $s = 1$ has $(\dim C_j, \dim C_k) = (0,1)$, up to interchange, and gives the upper bound $\boldsymbol{e}_6$. For $R_1 + n\delta$, arms $s = 2, 3$ have the same dimension pair, up to interchange, and give the upper bounds $\boldsymbol{e}_5, \boldsymbol{e}_3$. For both $T_3 + n\delta$ and $H + n\delta$, arms $s = 1, 2$ give, respectively, the upper bounds $\boldsymbol{e}_4, \boldsymbol{e}_2$.

The side-information decomposition of $R_2 + n\delta$ contains two three-line blocks. Use the determinant $B(\ell_1, \ell_2)$ from the proof of Lemma 21. Choose $\ell_1 \in U_1^{\perp}$ with $B(\ell_1, \cdot) \neq 0$ and $\ell_1|_{A_2 \cap U_3} \neq 0$. With ℓ_1 fixed, choose $\ell_2 \in U_2^{\perp}$ with $B(\ell_1, \ell_2) \neq 0$ and $\ell_2|_{A_3 \cap U_1} \neq 0$. At each step, the two excluded sets are proper subspaces by $B \neq 0$, $B(\ell_1, \cdot) \neq 0$, and (102). Both choices therefore exist over $\mathbb{F}_2$. For $P = \ker \ell_1 \cap \ker \ell_2$, the relation determinant gives

$$\dim P = 3n - 1, \quad \widehat{\boldsymbol{a}} = (n, n, n-1), \quad \widehat{\boldsymbol{u}} = (2n, 2n, 2n-1).$$

The pair U_{12} and the triple retain dimensions $n + 1$ and one. The imposed nonvanishing conditions reduce the dimensions of U_{23} and U_{13} to n. These dimensions give (113); the triple remains disjoint from every induced side-information space. Hence all nine values $2Q_{i\ell} = 3n - 1$ are unchanged.

Two restrictions remain uncontrolled: if ℓ_1 vanishes on $A_3 \cap U_2$, the increment increases by $\boldsymbol{e}_4$; if ℓ_2 vanishes on $A_1 \cap U_3$, it increases by $\boldsymbol{e}_5$. Thus $\boldsymbol{F}_P - \boldsymbol{F}_{\text{full}} \leq \boldsymbol{e}_4 + \boldsymbol{e}_5$ coordinatewise. By symmetry of the dimension vector and real-root uniqueness over $\mathbb{F}_2$ (Lemma 20), user permutations give all the listed bounds on the same component.

The remaining cases have zero pair margins and constant full profiles: every entry equals two for $T_3 + \delta$, and one for $H + \delta$ and $R_2 + \delta$. It suffices to attain these broadcast costs with upload-optimal relays. For $T_3 + \delta$, let $T = U_{123}$. By (102), $\dim T = 2$, $T \cap A_i = 0$, and $U_i = A_i \oplus T$. Thus broadcasting a basis of T solves all demands. Some U_s does not contain the plane of the three-line block: otherwise that plane lies in T, contradicting $T \cap A_i = 0$. For this arm, $P = U_s$ is a locally generated hyperplane by (112). Four uploads make T available and attain the component's upload optimum.

For $H + \delta$ and $R_2 + \delta$, the same rank identities give a line $T = U_{123}$, disjoint from every A_i, with $U_i = A_i + T$. In the first case $A_1 \oplus A_3 = F$; in the second $A_1 \oplus A_2 = F$. Decompose a generator of T into these two local spaces, upload its two components, and broadcast their sum. Both attain upload two and broadcast one.

We verify the pair-capacity requirements. The pair margins for $R_0 + n\delta, T_2 + n\delta, R_1 + n\delta$ are, respectively,

$$(n, n, n), \qquad (n, n-1, n), \qquad (n-1, n-1, n),$$

and those for $T_3 + n\delta, H + n\delta, R_2 + n\delta$ are $(n-1, n-1, n-1)$. They dominate the listed requirements. For T_2, R_1, T_3, H, the listed increment bounds are unit vectors, so one unit of capacity in the indicated pair suffices. For $R_0 + n\delta$, use $\boldsymbol{e}_4$ or $\boldsymbol{e}_6$ whenever its coordinate has positive capacity. Otherwise the first two pair constraints force positive capacities at coordinates 1 and 2. The third supplies coordinate 3 or 5, allowing $\boldsymbol{e}_1 + \boldsymbol{e}_3$ or $\boldsymbol{e}_2 + \boldsymbol{e}_5$.

For $R_2 + n\delta$, identify F_{ijk} with the directed edge $j \to i$. The pairs $(1,6), (2,4), (3,5)$ contain the two edges leaving vertices $2, 3, 1$, respectively. Choosing one positive-capacity coordinate per pair gives a graph with outdegree one at each vertex. It contains a directed path through all three vertices: either the graph is a three-cycle, or the third vertex enters

a two-cycle. These paths are exactly the user permutations of $\boldsymbol{e}_4 + \boldsymbol{e}_5$, proving that the stated pair-capacity requirement suffices. ∎

Lemma 25 (Scalar binary relay at integral upload)**.** *Let* M *be a binary task in active reduced form with* $U^* \in \mathbb{Z}$*. For every integer* $D \geq D^*$*, there is a feasible locally generated relay subspace* $P \leq F$ *such that*

$$\dim P = U^*, \qquad B_\nu(\mathsf{M}_P) \leq D \quad (\nu \in [15]).$$

Proof: Use the original binary relay choices of Lemma 23 and the additional choices of Lemma 24. Zero-excess components use $P = F$, and the $n = 0$ boundary of $T_2 + n\delta$ uses $P = A_1$ as in Appendix C-B. The three zero-margin cases use the last clause of Lemma 24. Extract all private-demand edges before allocating the relays; by Lemma 12, this leaves the positive-excess components and their pair-capacity requirements unchanged. Let $\mathcal{G}$ index the additional components in Table V, and let $c_{ab}(\gamma)$ be component γ's listed requirement in pair (a, b). Profile additivity gives the strengthened pair constraints

$$s_a + s_b \geq \sum_t \mu_t \underline{\kappa}_{ab}(t) + \sum_{\gamma \in \mathcal{G}} c_{ab}(\gamma), \tag{114}$$

for $(a, b) \in \{(1, 6), (2, 4), (3, 5)\}$, because each requirement is bounded by the corresponding component margin.

Since D is integral, all slacks $s_j = D - F_j(\mathsf{M})$ are integers. If an orientation constraint fails, its deficit ε_σ is an integer. Choose that many components in R and that many extracted cycles, and apply Lemma 13. The calculation in (82)–(84) leaves integer nonnegative slacks, reduces each original pair requirement by ε_σ, and decreases one pair capacity by the same amount. The additional pair-capacity requirements are unchanged. Both repaired orientation constraints therefore hold along with all strengthened pair constraints.

Allocate the remaining R components first, as in Appendix D-A. That construction uses only addition, subtraction, minima, and maxima, so integer inputs yield integer weights. Its choices consume at most one unit per opposite pair per component. The remaining pair capacities cover the requirements for the other original component families and all additional pair-capacity requirements. Reserve these capacities using a greedy allocation with integer amounts, allocating one component at a time. The same proof selects original relay choices within their reserved capacities, and Lemma 24 provides an additional choice within each additional component's reservation. Hence the scalar residual relay uses exactly U_{res} uploads, preserves or reduces every second-family entry, and keeps each first-family increment within its corresponding slack.

Since U_{res} is integral and $U^* = U_{\mathrm{res}} + U_{\mathrm{om}}$, the omniscience optimum U_{om} is also integral. Lemma 7 gives an optimal scalar upload for that component. Its induced profile is unchanged by (60). Combining the two relays gives the stated dimension and profile bounds. ∎

D. Separation of the Induced Binary Side-Information Spaces

We complete the calculation of the induced side-information spaces in Proposition 14. Use the relay subspace assembled in its proof and write its residual and omniscience summands as P_{res} and P_{om}.

Proof: Lemma 22 gives

$$P_{\mathrm{res}} = \bigoplus_i (P_{\mathrm{res}} \cap A_{i,\mathrm{res}}). \tag{115}$$

For the omniscience component, take the optimal rates r_i of Lemma 7, whose sum is $U_{\mathrm{om}} \in \mathbb{Z}$. The local upload spaces $P_{\mathrm{om},i} \leq A_{i,\mathrm{om}}$, of dimensions r_i, form a direct sum: any dependence would allow one uploaded symbol to be omitted without changing the information available to the relay, contradicting upload optimality. Set $\widehat{A}_{i,\mathrm{om}} = P_{\mathrm{om}} \cap A_{i,\mathrm{om}}$. Omniscience feasibility gives $\dim \widehat{A}_{i,\mathrm{om}} = U_{\mathrm{om}} - d_{\mathrm{om},i}$. If $2 \max_i d_{\mathrm{om},i} \leq \sum_i d_{\mathrm{om},i}$, then $r_i = U_{\mathrm{om}} - d_{\mathrm{om},i}$. Since $P_{\mathrm{om},i} \subseteq \widehat{A}_{i,\mathrm{om}}$ and their dimensions agree, $\widehat{A}_{i,\mathrm{om}} = P_{\mathrm{om},i}$, so the $\widehat{A}_{i,\mathrm{om}}$ form a direct sum. Otherwise, relabel users so that $d_{\mathrm{om},1} > d_{\mathrm{om},2} + d_{\mathrm{om},3}$. Then $U_{\mathrm{om}} = d_{\mathrm{om},1}$, $r_1 = 0$, $\widehat{A}_{1,\mathrm{om}} = 0$, and $\widehat{A}_{2,\mathrm{om}} + \widehat{A}_{3,\mathrm{om}} = P_{\mathrm{om}}$. In either case, (115) gives (81). ∎

E. Integer LCBC Allocations

For the broadcast arguments in this appendix, write $\mathsf{M} = (A_i \subseteq U_i \subseteq F)_{i=1}^3$ for the LCBC task and $D^* = D^*(\mathsf{M})$ for its minimum broadcast cost. Here A_i is User i's side-information space, U_i is its signal space, and $m_i = \dim U_i - \dim A_i$ is its conditional demand dimension. A general LCBC task need not satisfy $F = \sum_i A_i$. All dimensions and ranks in this part of the appendix refer to this task. Use the conditional ranks of [21, Sec. IX-A]: $r_{123} = \min_\ell \mathrm{rk}(U_{123} \mid A_\ell)$, $r_{ij} = \min_{\ell \in \{i,j\}} \mathrm{rk}(U_{ij} \mid A_\ell)$, $r_{ij,ik} = \mathrm{rk}(U_{ij} + U_{ik} \mid A_i)$, and $r_{i(j,k)} = \mathrm{rk}\big(U_{i(j,k)} \mid A_i\big)$, where $\{i, j, k\} = [3]$, and put $m_\Sigma = \sum_i m_i$. These satisfy $r_{ij,ik} \leq r_{i(j,k)} \leq m_i$.

The allocation variables $\lambda_{123}, \lambda_{12}, \lambda_{13}, \lambda_{23}, \lambda_{\mathrm{al}}$ are nonnegative and satisfy [21, Thm. 2]

$$\begin{aligned} \lambda_{123} \leq r_{123}, \qquad \lambda_{123} + \lambda_{ij} \leq r_{ij}, \\ \lambda_{123} + \lambda_{ij} + \lambda_{ik} \leq r_{ij,ik}, \\ \lambda_{123} + \lambda_{ij} + \lambda_{ik} + \lambda_{\mathrm{al}} \leq r_{i(j,k)}. \end{aligned} \tag{116}$$

For an integer allocation, λ_{123} common dimensions serve all three users, λ_{ij} pair dimensions serve Users i, j, and $2\lambda_{\mathrm{al}}$ aligned transmissions supply λ_{al} demand dimensions to each user. The remaining individual allocation is $\lambda_i := m_i - \lambda_{123} - \lambda_{ij} - \lambda_{ik} - \lambda_{\mathrm{al}}$. The LCBC linear program therefore minimizes

$$\begin{aligned} \lambda_{123} + \lambda_{12} + \lambda_{13} + \lambda_{23} + 2\lambda_{\mathrm{al}} + \sum_i \lambda_i \\ = m_\Sigma - 2\lambda_{123} - \lambda_{12} - \lambda_{13} - \lambda_{23} - \lambda_{\mathrm{al}}, \end{aligned} \tag{117}$$

over (116); its optimum is D^*.

Lemma 26 (Integer LCBC allocation)**.** *Every three-user LCBC task admits an integer allocation satisfying* (116)*, with* $\lambda_{123} = r_{123}$ *and total broadcast length* $\lceil D^* \rceil$.

Proof: Set $b_i = r_{i(j,k)}$. The waterfilling formulation [21, Sec. IX-B] gives

$$\begin{aligned} D^* &= m_\Sigma - r_{123} - h^*, \\ h^* &= \max_{\boldsymbol{w}\in\mathcal{W}} \min_i (b_i + w_i), \\ \mathcal{W} &= \{\boldsymbol{w} \in \mathbb{R}_+^3 :\ w_i \le r_{jk} - r_{123}, \\ &\qquad\qquad w_j + w_k \le r_{ij,ik} - r_{123}\}. \end{aligned} \tag{118}$$

Set $h = \lfloor h^* \rfloor$ and, for an optimizer $\boldsymbol{w}^*$, let $\bar{w}_i = (h - b_i)_+$. Then $0 \le \bar{w}_i \le w_i^*$, so $\bar{\boldsymbol{w}} \in \mathcal{W}$ by downward closure. Moreover, $\min_i(b_i + \bar{w}_i) \ge h$. This minimum is an integer at most h^*, hence it equals h. Set $\lambda_{123} = r_{123}$, $\lambda_{jk} = \bar{w}_i$, and

$$\lambda_{\mathrm{al}} = \min_i (b_i - \bar{w}_j - \bar{w}_k) - r_{123}.$$

The pair-sum constraints and $r_{ij,ik} \le b_i$ give $\lambda_{\mathrm{al}} \ge 0$, while $\lambda_i \ge m_i - b_i \ge 0$. Thus all allocation variables are nonnegative integers. Finally, $\lambda_{\mathrm{al}} = h - \sum_i \bar{w}_i - r_{123}$, and (117) gives length

$$m_\Sigma - r_{123} - \lfloor h^* \rfloor = \lceil D^* \rceil.$$

■

F. Original-Field Coefficient Selection

Lemma 27 (Homogeneous column coordination)**.** *Let $C_1(\theta), \dots, C_t(\theta)$ be square matrices over $\mathbb{F}_Q$, with parameters partitioned into nonempty vector groups $\theta_1, \dots, \theta_G$. In each matrix a group is either absent or supplies exactly one column $L_{ig}\theta_g$, with no constant term; distinct groups supply distinct columns. Suppose an assignment making C_i invertible is known for each i. If $Q \ge t$, a common assignment making all matrices invertible can be found deterministically.*

Proof: Process the parameter groups one at a time. Keep one invertible assignment for each C_i, with all assignments agreeing on the processed groups. For the next group $x \in \mathbb{F}_Q^r$, fix the other groups at assignment i. The determinant is then a nonzero constant or a nonzero homogeneous linear form in x. The kernels of at most $t \le Q$ such forms cannot cover $\mathbb{F}_Q^r$, since their union has size at most $1 + t(Q^{r-1} - 1) < Q^r$ for $r \ge 1$. Choose x outside these kernels and use this value in every assignment. All matrices remain invertible. After the last group, all assignments agree. ■

For aligned columns, use the full relation space

$$\mathcal{L} = \{(h_1, h_2, h_3) \in U_1 \times U_2 \times U_3 : h_1 + h_2 = h_3\}. \tag{119}$$

Its ith coordinate projection is $U_{i(j,k)}$. Let $\mathbf{L}_i$ be the coordinate blocks of a column-basis matrix of $\mathcal{L}$, so $\mathbf{L}_1 + \mathbf{L}_2 = \mathbf{L}_3$.

G. Scalar Broadcasts over Fields of Size at Least Three

Every feasible integer allocation in (116) can be implemented by a scalar linear code over the task's defining field $\mathbb{F}_q$ when $q \ge 3$.

Proof of Proposition 15: Fix a feasible integer allocation and a column-basis matrix V_i' of each A_i. Let $\mathbf{U}_{123}, \mathbf{U}_{ij}$ be column-basis matrices of the corresponding intersections. With variable coefficient matrices $\mathbf{N}_{123}, \mathbf{N}_{ij}, \mathbf{N}_{\mathrm{al}}$, form

$$C_i = [V_i', \mathbf{U}_{123}\mathbf{N}_{123}, \mathbf{U}_{ij}\mathbf{N}_{ij}, \mathbf{U}_{ik}\mathbf{N}_{ik}, \mathbf{L}_i\mathbf{N}_{\mathrm{al}}, \mathbf{H}_i, \mathbf{J}_i].$$

The shared blocks have the allocated column counts. We first choose the fixed completion blocks $\mathbf{H}_i, \mathbf{J}_i$ and an invertible assignment for each C_i separately.

For User i, choose λ_{123} common columns independent modulo A_i. They lie in both pair spaces, so adjoining them to A_i reduces each pair rank and their sum rank by λ_{123}. The allocation constraints and [21, Lem. 4] then give both pair blocks, jointly independent modulo A_i and the common block. All these columns lie in $U_{i(j,k)}$, whose remaining conditional rank is at least λ_{al}. Choose λ_{al} more independent columns there and lift them through the ith projection of $\mathcal{L}$. Choose $\mathbf{H}_i$, with λ_i columns, to complete a basis of U_i, then $\mathbf{J}_i$ to complete a basis of F. Groups absent from C_i may be zero. This construction works for any common block independent modulo A_i; $\mathbf{H}_i, \mathbf{J}_i$ remain fixed in the next step.

Each column of $\mathbf{N}_{123}, \mathbf{N}_{ij}$, or $\mathbf{N}_{\mathrm{al}}$ is one parameter group and supplies at most one homogeneous column to each C_i. For $q \ge 3$, Lemma 27 makes the three assignments agree. Transmit the linear combinations specified by the columns of the common and pair blocks, the individual blocks $\mathbf{H}_i$, and $\mathbf{L}_1\mathbf{N}_{\mathrm{al}}$ and $\mathbf{L}_2\mathbf{N}_{\mathrm{al}}$. User 3 adds the aligned transmissions. Removing the untransmitted $\mathbf{J}_i$ leaves a basis of U_i, so every user can decode. The stated length follows from Lemma 26. ■

H. A Common Block for Separated Side-Information Spaces

Lemma 28 (A maximum disjoint subspace)**.** *Let $T_1, T_2, T_3 \le T$ over a field $\mathbb{K}$, with $T_1 \cap (T_2 + T_3) = 0$. There is $C \le T$ such that*

$$\dim C = \dim T - \max_i \dim T_i, \qquad C \cap T_i = 0.$$

Proof: Write $T_{23} = T_2 \cap T_3$, choose $T_i = T_{23} \oplus E_i$ for $i = 2, 3$, and extend to $T = T_1 \oplus T_{23} \oplus E_2 \oplus E_3 \oplus R$. Set $m = \max_i \dim T_i$, $c = \dim T_{23}$. Map T_{23} onto the first c coordinates of $\mathbb{K}^m$, each E_i injectively to the last $m - c$, T_1 injectively to $\mathbb{K}^m$, and R to zero. This map is injective on every T_i and hence surjective on a largest one. Its kernel is C, including when $m = 0$. ■

I. Scalar Broadcasts for the Selected Binary Relay

Proof of Proposition 16: By Lemma 26, choose an integer allocation of length $\lceil D^* \rceil$ with $\lambda_{123} = r_{123}$, using the notation of Appendix F-E. Set $T = U_{123}$ and $T_i = T \cap A_i$. The assumed separation $A_1 \cap (A_2 + A_3) = 0$ implies $T_1 \cap (T_2 + T_3) = 0$. By Lemma 28, there is $C \le T$ with $\dim C = r_{123}$ and $C \cap A_i = 0$. Fix a basis of C as the common block in every decoder. The individual basis completion in the proof of Proposition 15 gives a separate invertible assignment for each decoder, with this block fixed.

Coordinate the remaining groups as in Lemma 27. Pair groups affect at most two decoders, so $Q = 2$ suffices. For an aligned group, we show that all three determinant forms can simultaneously equal one.

For $h = (h_1, h_2, h_3) \in \mathcal{L}$, write $\ell_i(h) = f_i(h_i)$, where the cofactor functional $f_i : U_i \to \mathbb{F}_2$ vanishes on A_i. Each form is nonzero because its current decoder matrix is invertible.

The system $\ell_i(h) = 1$, $i \in [3]$, is inconsistent only if an odd number of its coefficient rows sum to zero. No row is zero, so the only possible obstruction is $\ell_1 + \ell_2 + \ell_3 = 0$.

Suppose this sum were zero on $\mathcal{L}$. Then the functional $(u_1, u_2, u_3) \mapsto \sum_i f_i(u_i)$ would vanish on $\mathcal{L}$, the kernel of the surjective sum map $(u_1, u_2, u_3) \mapsto u_1 + u_2 + u_3$ from $U_1 \oplus U_2 \oplus U_3$ to F. It would therefore induce a functional $f : F \to \mathbb{F}_2$ with $f|_{U_i} = f_i$. But f vanishes on every A_i, and $\sum_i A_i = F$, so $f = 0$. This contradicts $\ell_i \neq 0$. Thus solving a binary linear system at each step makes all assignments agree and gives a code of length $\lceil D^* \rceil$. ∎

REFERENCES

[1] S. Li, M. A. Maddah-Ali, Q. Yu, and A. S. Avestimehr, "A fundamental tradeoff between computation and communication in distributed computing," *IEEE Trans. Inf. Theory*, vol. 64, no. 1, pp. 109–128, Jan. 2018.

[2] P. Kairouz, H. B. McMahan, B. Avent, A. Bellet, M. Bennis, A. N. Bhagoji *et al.*, "Advances and open problems in federated learning," *Found. Trends Mach. Learn.*, vol. 14, no. 1–2, pp. 1–210, Jun. 2021.

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

[4] Q. Yu, M. A. Maddah-Ali, and A. S. Avestimehr, "The exact rate-memory tradeoff for caching with uncoded prefetching," *IEEE Trans. Inf. Theory*, vol. 64, no. 2, pp. 1281–1296, Feb. 2018.

[5] B. Chor, O. Goldreich, E. Kushilevitz, and M. Sudan, "Private information retrieval," *J. ACM*, vol. 45, no. 6, pp. 965–982, Nov. 1998.

[6] S. Ulukus, S. Avestimehr, M. Gastpar, S. A. Jafar, R. Tandon, and C. Tian, "Private retrieval, computing, and learning: Recent progress and future challenges," *IEEE J. Sel. Areas Commun.*, vol. 40, no. 3, pp. 729–748, Mar. 2022.

[7] A. Elmahdy and S. Mohajer, "On the fundamental limits of coded data shuffling for distributed machine learning," *IEEE Trans. Inf. Theory*, vol. 66, no. 5, pp. 3098–3131, May 2020.

[8] A. G. Dimakis, K. Ramchandran, Y. Wu, and C. Suh, "A survey on network codes for distributed storage," *Proc. IEEE*, vol. 99, no. 3, pp. 476–489, Mar. 2011.

[9] S. Kadhe, N. Rajaraman, O. O. Koyluoglu, and K. Ramchandran, "FastSecAgg: Scalable secure aggregation for privacy-preserving federated learning," arXiv:2009.11248, Sep. 2020. [Online]. Available: https://arxiv.org/abs/2009.11248

[10] J. Körner and K. Marton, "How to encode the modulo-two sum of binary sources (corresp.)," *IEEE Trans. Inf. Theory*, vol. 25, no. 2, pp. 219–221, Mar. 1979.

[11] A. Orlitsky and J. R. Roche, "Coding for computing," *IEEE Trans. Inf. Theory*, vol. 47, no. 3, pp. 903–917, Mar. 2001.

[12] V. Doshi, D. Shah, M. Médard, and M. Effros, "Functional compression through graph coloring," *IEEE Trans. Inf. Theory*, vol. 56, no. 8, pp. 3901–3917, Aug. 2010.

[13] A. Giridhar and P. R. Kumar, "Computing and communicating functions over sensor networks," *IEEE J. Sel. Areas Commun.*, vol. 23, no. 4, pp. 755–764, Apr. 2005.

[14] R. Appuswamy, M. Franceschetti, N. Karamchandani, and K. Zeger, "Network coding for computing: Cut-set bounds," *IEEE Trans. Inf. Theory*, vol. 57, no. 2, pp. 1015–1030, Feb. 2011.

[15] R. Appuswamy and M. Franceschetti, "Computing linear functions by linear coding over networks," *IEEE Trans. Inf. Theory*, vol. 60, no. 1, pp. 422–431, Jan. 2014.

[16] Z. Bar-Yossef, Y. Birk, T. S. Jayram, and T. Kol, "Index coding with side information," *IEEE Trans. Inf. Theory*, vol. 57, no. 3, pp. 1479–1494, Mar. 2011.

[17] F. Arbabjolfaei and Y.-H. Kim, "Fundamentals of index coding," *Found. Trends Commun. Inf. Theory*, vol. 14, no. 3–4, pp. 163–346, 2018.

[18] N. Lee, A. G. Dimakis, and R. W. Heath, Jr., "Index coding with coded side-information," *IEEE Commun. Lett.*, vol. 19, no. 3, pp. 319–322, Mar. 2015.

[19] E. Byrne and M. Calderini, "Error correction for index coding with coded side information," *IEEE Trans. Inf. Theory*, vol. 63, no. 6, pp. 3712–3728, Jun. 2017.

[20] H. Sun and S. A. Jafar, "On the capacity of computation broadcast," *IEEE Trans. Inf. Theory*, vol. 66, no. 6, pp. 3417–3434, Jun. 2020.

[21] Y. Yao and S. A. Jafar, "The capacity of 3 user linear computation broadcast," *IEEE Trans. Inf. Theory*, vol. 70, no. 6, pp. 4414–4438, Jun. 2024.

[22] ——, "On the generic capacity of K-user symmetric linear computation broadcast," *IEEE Trans. Inf. Theory*, vol. 70, no. 5, pp. 3693–3717, May 2024.

[23] Y. Ma and D. Tuninetti, "An achievable scheme for the K-user linear computation broadcast channel," in *Proc. IEEE Int. Symp. Inf. Theory (ISIT)*, Ann Arbor, MI, USA, Jun. 2025, pp. 1–6.

[24] S. Tan and S. A. Jafar, "Wireless linear computation broadcast," *IEEE Transactions on Communications*, vol. 74, pp. 15 832–15 845, 2026.

[25] Y. Yao and S. A. Jafar, "Capacity of 2-user linear computation broadcast with partially informed server," *IEEE Commun. Lett.*, vol. 26, no. 9, pp. 2014–2017, Sep. 2022.

[26] J. Ravi and B. K. Dey, "Function computation through a bidirectional relay," *IEEE Trans. Inf. Theory*, vol. 65, no. 2, pp. 902–916, Feb. 2019.

[27] S. Li, Q. Yu, M. A. Maddah-Ali, and A. S. Avestimehr, "A scalable framework for wireless distributed computing," *IEEE/ACM Trans. Netw.*, vol. 25, no. 5, pp. 2643–2654, Oct. 2017.

[28] Q. Yan, X. Tang, M. Tao, and Q. Huang, "A fundamental tradeoff among storage, computation, and communication for distributed computing over star network," *IEEE Trans. Commun.*, vol. 71, no. 10, pp. 5686–5697, Oct. 2023.

[29] Y. Wang and Y. Wu, "Coded distributed computing with pre-set data placement and output functions assignment," *IEEE Trans. Inf. Theory*, vol. 71, no. 3, pp. 2195–2217, Mar. 2025.

[30] I. Csiszár and P. Narayan, "Secrecy capacities for multiple terminals," *IEEE Trans. Inf. Theory*, vol. 50, no. 12, pp. 3047–3061, Dec. 2004.

[31] N. Milosavljevic, S. Pawar, S. El Rouayheb, M. Gastpar, and K. Ramchandran, "Efficient algorithms for the data exchange problem," *IEEE Trans. Inf. Theory*, vol. 62, no. 4, pp. 1878–1896, Apr. 2016.

[32] V. G. Kac, "Root systems, representations of quivers and invariant theory," in *Invariant Theory*, ser. Lecture Notes in Mathematics, F. Gherardelli, Ed. Berlin, Germany: Springer, 1983, vol. 996, pp. 74–108.

[33] ——, "On complexity of representations of quivers," *C. R. Math.*, vol. 357, no. 11–12, pp. 841–845, 2019.

[34] W. Crawley-Boevey, "Lectures on representations of quivers," Lecture notes, 1992. [Online]. Available: https://www.math.uni-bielefeld.de/~wcrawley/quivlecs.pdf

[35] C. Herrmann, J. Lorand, and A. Weinstein, "(Co)isotropic triples and poset representations," *Rend. Sem. Mat. Univ. Padova*, vol. 146, pp. 57–162, 2021, author preprint: arXiv:1906.04882. [Online]. Available: https://arxiv.org/abs/1906.04882