Domination Strategies for Free-Riding in Cross-Silo FL-Based Caching

Jiqing Gu, Chao Song*, Jianfeng Huang, Jie Wu, Rulin Hu, Li Lu

University of Electronic Science and Technology of China; Chengdu University of Information Technology; Temple University

ICASSP 2026

UESTC

Contents

Background: Edge Caching & Privacy Bottleneck

The Caching Imperative

- Explosive growth in video content (e.g., YouTube, TikTok).
- Proactive edge caching improves Quality of Experience (QoE) and offloads ISP traffic.
- Relies on predicting Zipf-like distributions of popular content.

The Privacy Bottleneck & FL Solution

- Problem: ISPs possess localized, heterogeneous data but refuse to share it due to privacy and market competition.
- Solution: Cross-Silo Federated Learning (FL).
- Mechanism: "Train locally, aggregate globally." Raw user data is never exposed.

The Core Problem: Free-Riders in Federated Caching

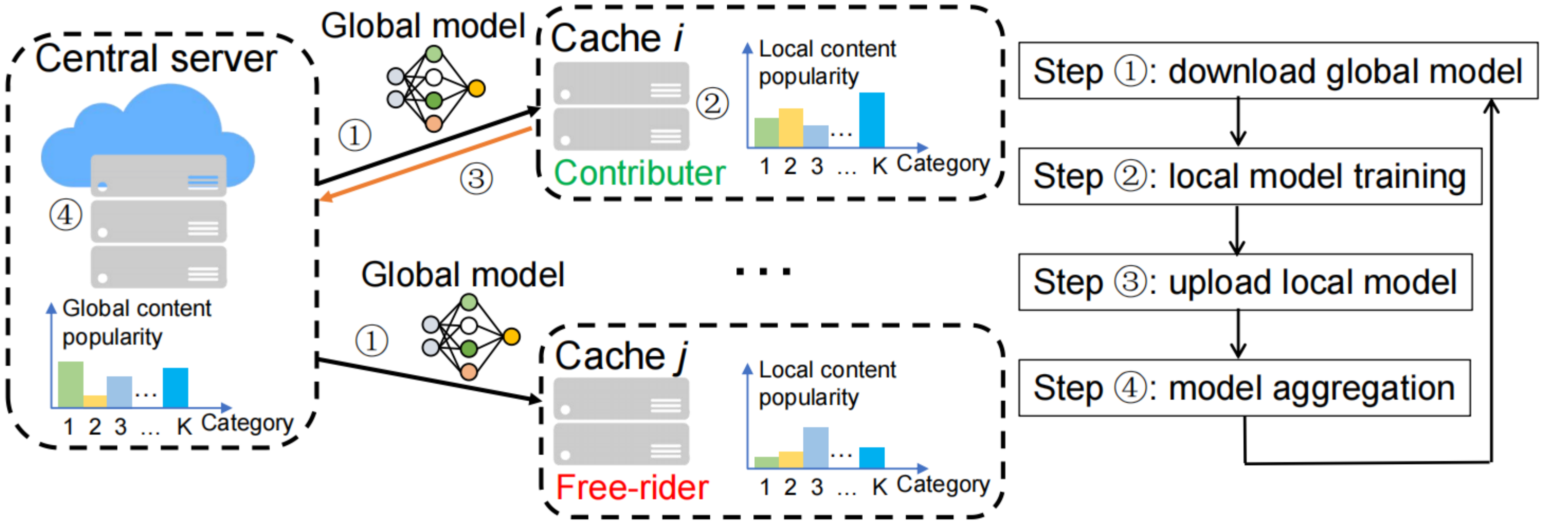

The Free-Rider Problem

- The global model acts like a **public good** accessible to all participants.
- Selfish local caches can simply **download the global model** to improve their own caching accuracy — **without contributing their own local updates** to the federation.

Critical Consequences:

❌ **Degraded Global Accuracy:** Model performance suffers from lack of diverse local data.

📉 **Reduced Collective Payoff:** System-wide efficiency drops for all nodes.

💥 **Long-term Instability:** Entire FL system risks collapse due to defection.

Problem Modeling: A Repeated Game

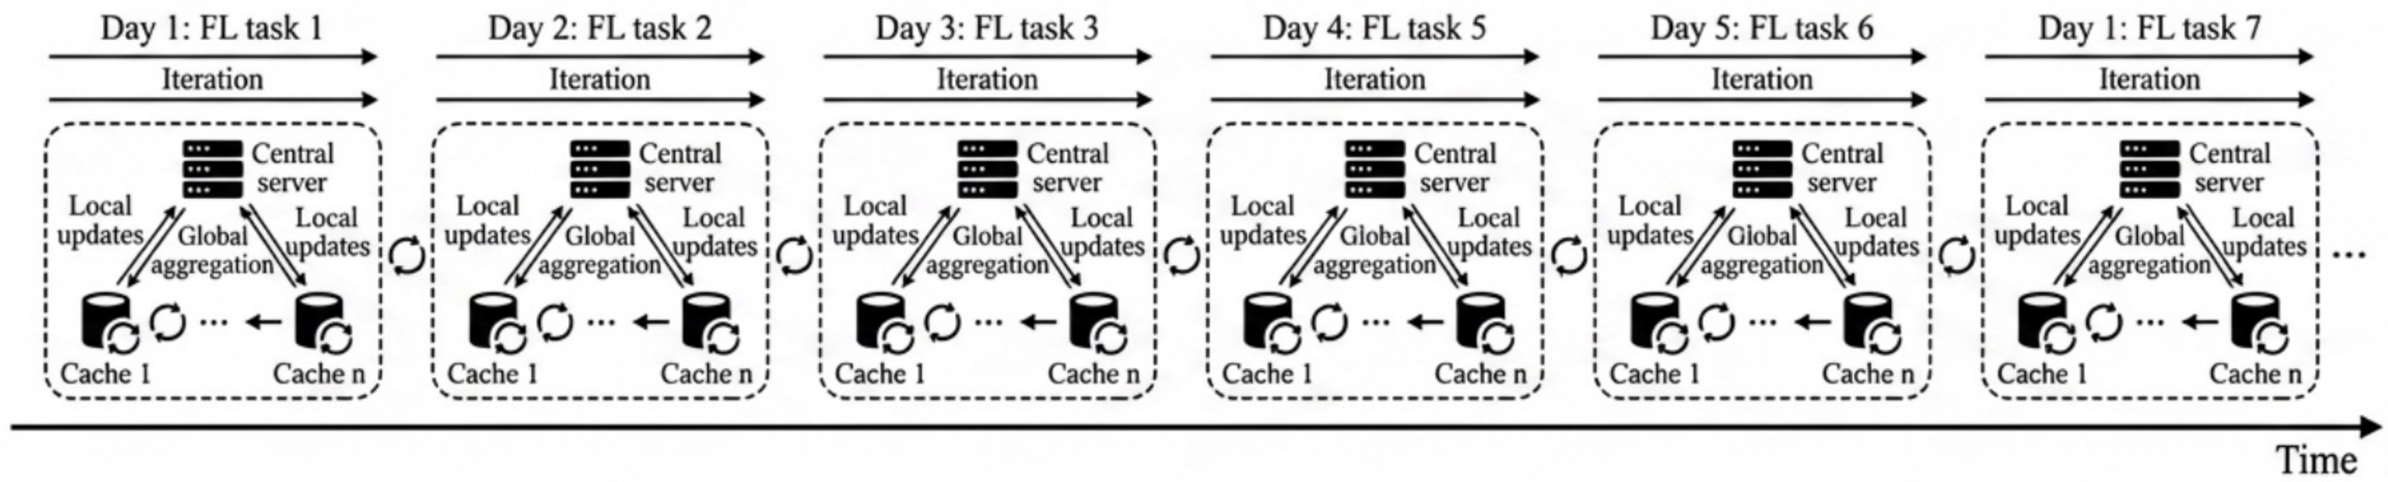

Game Theory Perspective

•**Interaction Model:** We model the caching node interaction as a **Repeated Game** (daily FL task iterations).

•**Utility Function:** $U_i(x) = f_i(x) - g_i(x)$

-$f_i(x)$:Shared Income (depends on **total** participation)

-$g_i(x)$:Individual Cost (depends on **personal** participation)

• **Key Insight: Income is shared, cost is individual → Natural incentive to free-ride.**

The Goal

Maximize & Stabilize Social Welfare (SW):

The sum of expected long-term payoffs for all caching nodes (players). The aim is to design mechanisms that align individual incentives with the collective good.

"Incentivize cooperation to avoid a suboptimal Nash equilibrium where social welfare collapses."

Our Solution: The MMZD Strategy

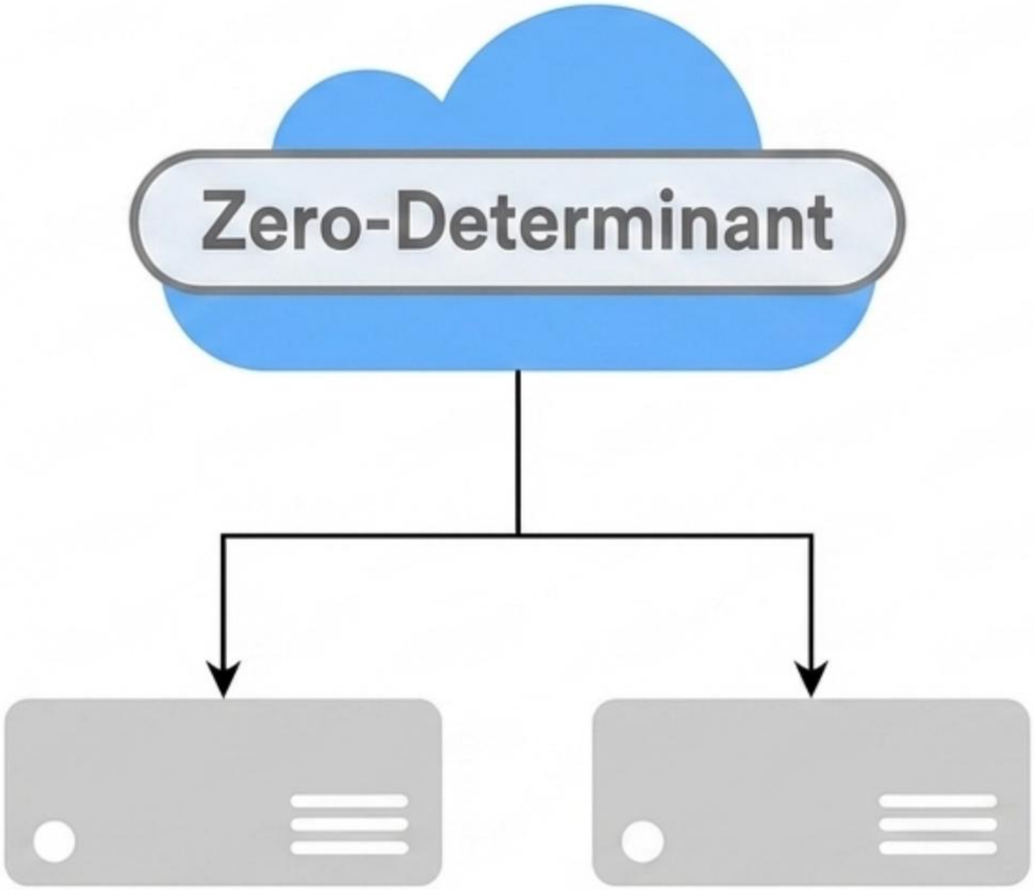

The Strategy

Multi-player Multi-action Zero-Determinant (MMZD).
A class of probabilistic, conditional strategies designed to tackle free-riding in federated environments.

Mechanism of Action

• Enables a single player (e.g., a dominant ISP) to **unilaterally set the sum of expected payoffs** or enforce a specific payoff ratio.
• Operates **independently of opponents' strategies** — no need to predict or react to others' moves.

The MMZD adopter dynamically modulates its aggregation participation to counteract free-riders, effectively **locking social welfare at a stable maximum.**

Theoretical Guarantee: Bounding the Social Welfare

$$max\ \gamma = \sum_{i=1}^{n} \alpha_i E_i\,,$$

$$\text{s.t.} \begin{cases} 0 \leq p_{j,s}^{1} \leq 1, \mathrm{j} \in \{0,1,2,\ldots,\mathrm{r}\}, s \in \{1,2,\ldots,(r+1)^n\} \\ \hat{\mathbf{p}}^1 = \lambda\left(\sum_{i=1}^{n} \alpha_i \mathbf{u}^i - \gamma \mathbf{1}\right) \\ \lambda \neq 0 \end{cases}$$

Theorem Guarantee

Social welfare can be mathematically forced to converge to a specific target value γ, ensuring predictability in the system's equilibrium.

Control Parameter λ

This key hyperparameter dictates the convergence rate of the system towards the stable state. Tuning λ optimizes both speed and stability.

Key Result

MMZD guarantees social welfare is protected from collapse, remaining locked within a strictly defined, controllable range regardless of initial conditions.

Experimental Setup

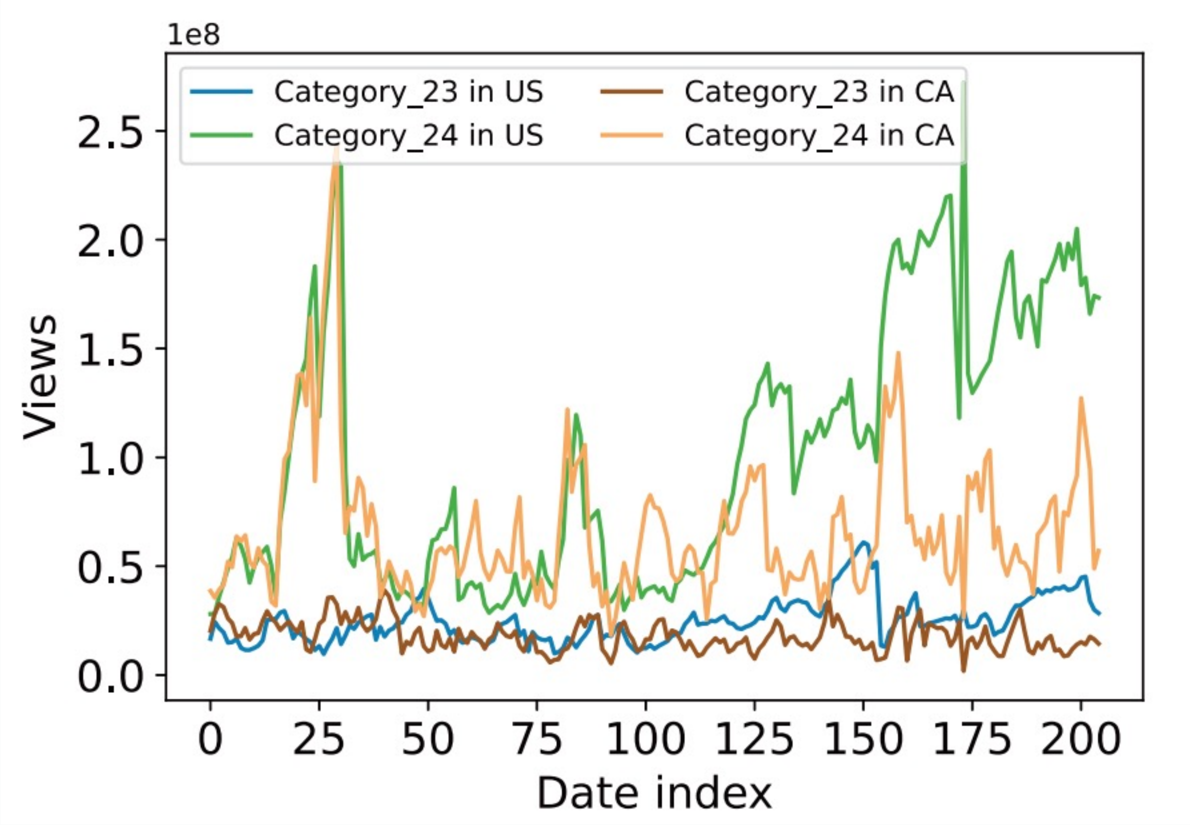

Dataset Overview

Source: Real-world YouTube Dataset detailing daily popular videos across 10 countries.

Characteristic: Highly volatile, time-varying local requests reflecting real-world dynamics.

Framework & Opponent Strategies

• **Infrastructure:** FedML open research library | **Model:** LSTM (Sequence Prediction) | **Aggregation:** FedAvg

• **Challengers:** Tested against four distinct opponent strategies: ALLC (Cooperate), ALLD (Defect), Random (No Pattern), and Evolved (Adaptive).

Result 1: The Harm of Free-Riding

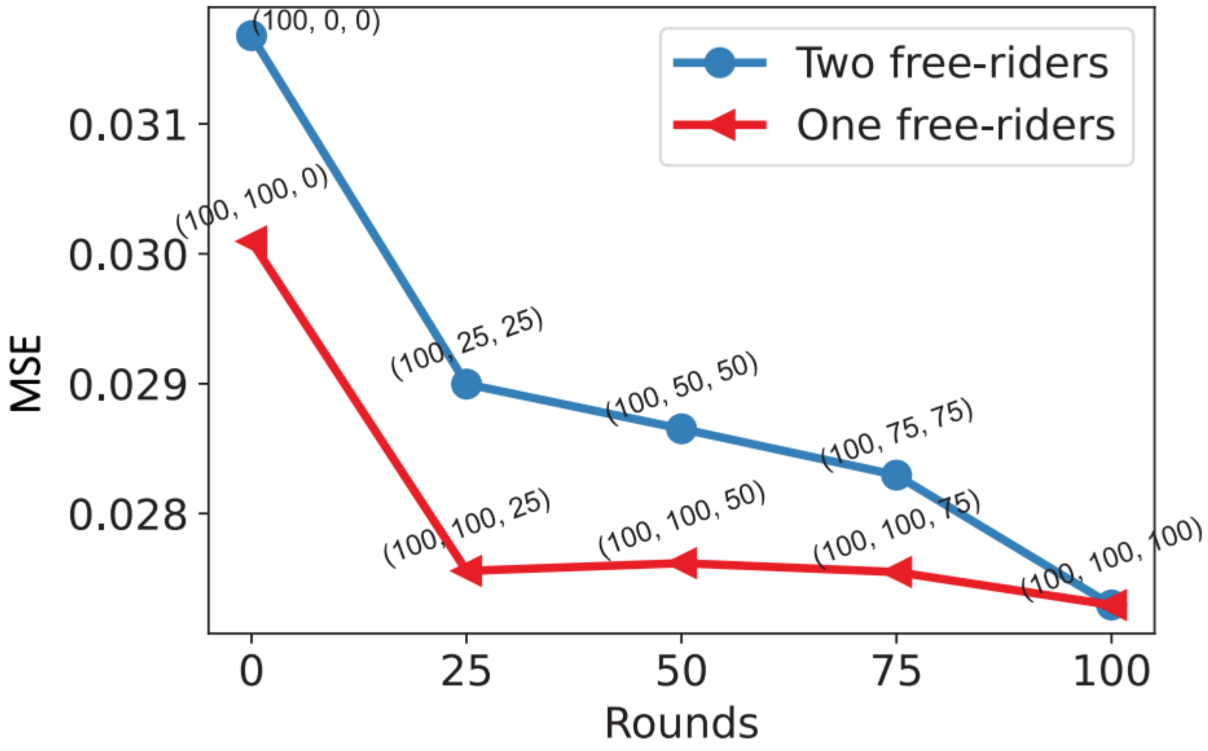

MSE Soars with Free-Riders

MSE increases severely when the number of free-riders increases from one to two, showing a sharp rise in prediction error.

Perfect Participation is Key

Prediction error plummets only when the participation probability is total (prob = 1), indicating the necessity of full cooperation.

Key Takeaways

- Low cooperation yields volatile, high error rates.
- Free-riders actively destroy the predictive utility of the global model.

Result 2: MMZD Stabilizes Social Welfare

Welfare over Time

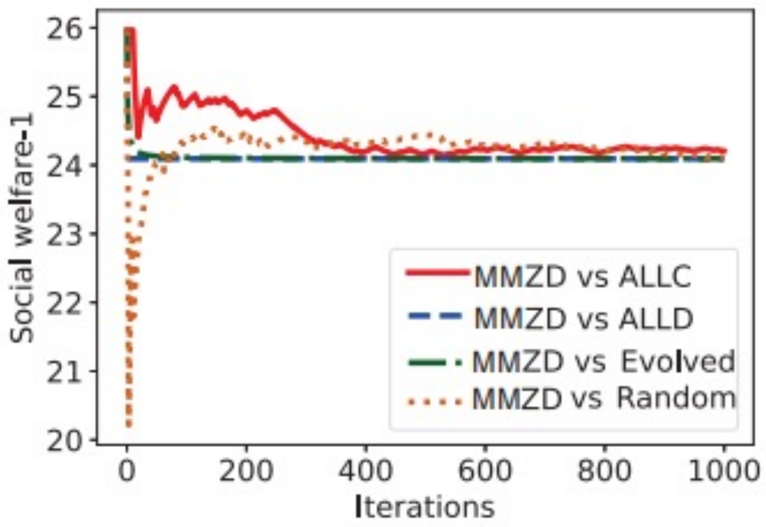

(a) Social welfare control under MMZD strategy.

Payoff Distribution

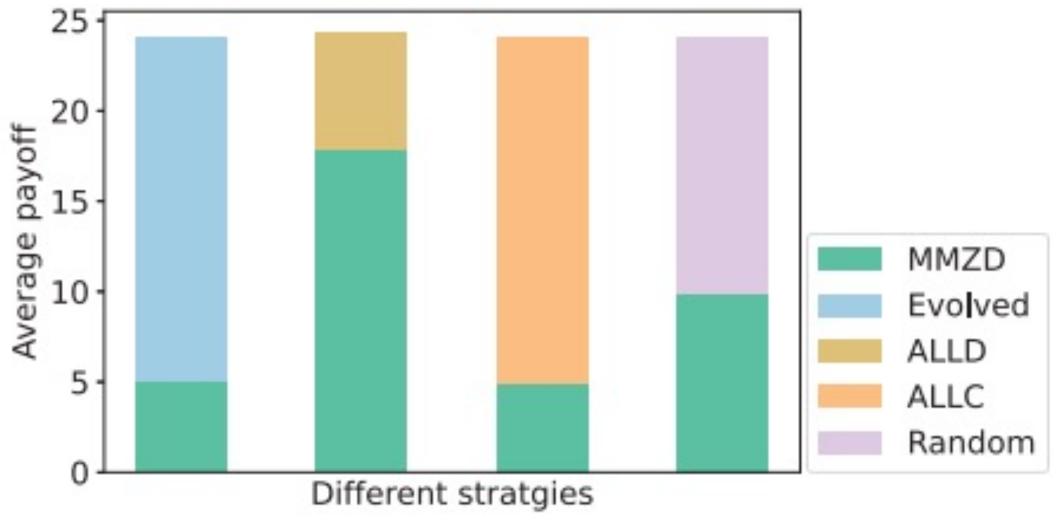

(b) Cooperation probability of the evolved cache under pinning strategy.

● **Welfare Stability:** MMZD successfully maintains the social welfare at its maximal possible value (~24), entirely preventing system collapse against all tested opponent strategies.

● **Payoff & Convergence:** When opponents defect (ALLD), the MMZD cache shoulders the burden. Evolved learning players eventually recognize the ineffectiveness of defection and converge toward full cooperation (ALLC).

Conclusion & Core Insights

1. The Vulnerability

Free-riding poses a fatal threat to privacy-preserving Cross-Silo FL in edge caching systems.

2. The Framework

By modeling the interaction as a repeated game, we can deploy the **Multi-player Multi-action Zero-Determinant (MMZD)** strategy.

3. The Guarantee

We proved mathematically and empirically that a **single** market-power cache can **unilaterally** dominate free-riders.

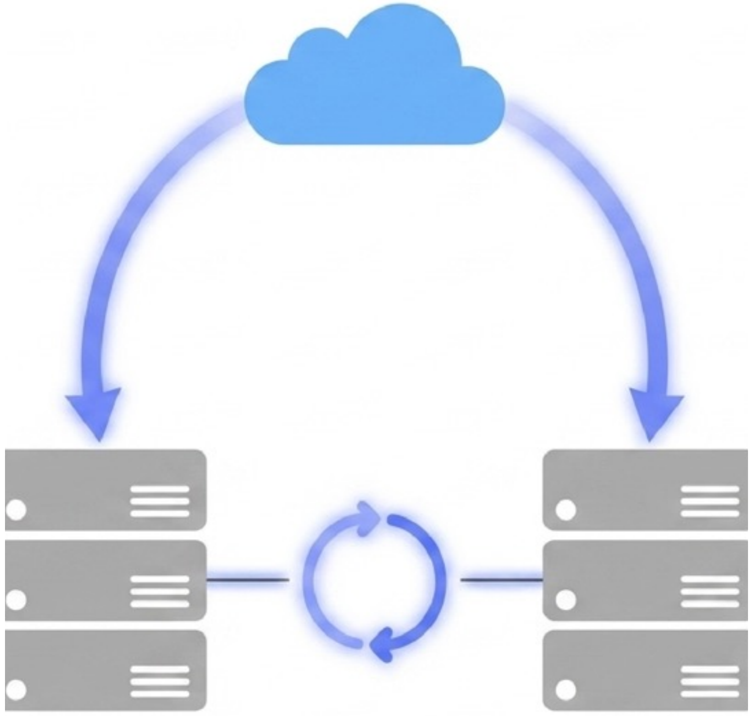

The Outcome

The system maintains **stable, maximal social welfare** and **robust content prediction accuracy** even under highly adversarial conditions.

Thank You!

Q&A: chaosong@uestc.edu.cn