Content Fragmentation: A New Inter-Data Center Content Redundancy Scheme to Saving Energy in Optical Cloud Networks

Group Meeting Presentation

Speaker: Yu Wu
1/12/2018

UC DAVIS
UNIVERSITY OF CALIFORNIA

Motivation

- The rapid growth of content-based cloud services (as video streaming, etc.) is raising concerns on the environmental sustainability of the supporting infrastructure.

- Traditional inter-DC content redundancy scheme -- Content Replication (CR) failed to save energy on storage, as it usually requires multiple replicas of a content to be distributed across DCs to guarantee resiliency.

- To explore energy optimization possibilities on storage, in this study, we propose a new scheme called **Content Fragmentation (CF) to achieve less storage redundancy, thus less energy consumption**.

- **However, the tradeoff is that CF consumes more energy than CR on core network and on content decoding.**

Content Fragmentation

- Content is fragmented into k equally-sized data blocks.
- Erasure code is used to perform content encoding, i.e., encoding k data blocks to r parity blocks of same size.
- A specific type of erasure code, called Reed-Solomon (RS), is used to provide inter-DC redundancy.
- RS allows any k blocks out of (k+r) blocks to reconstruct the original content.

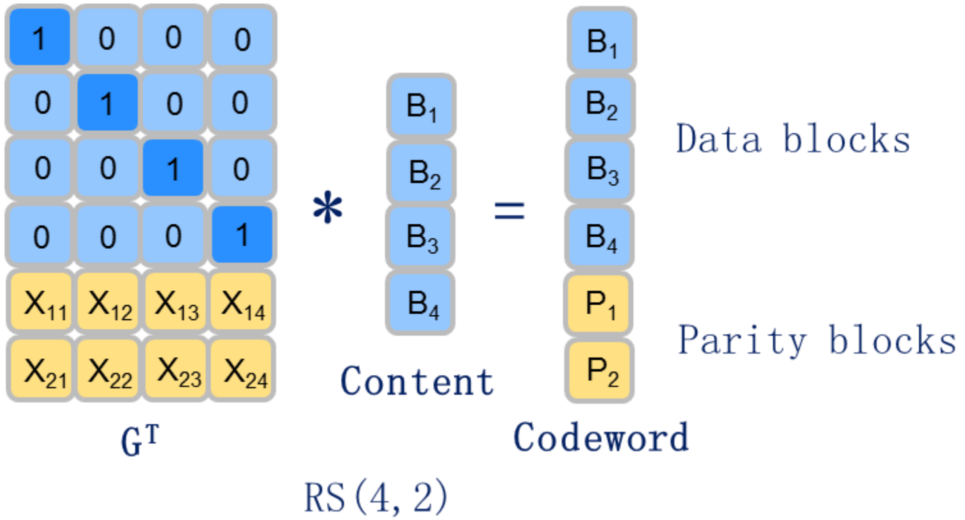

UC DAVIS
UNIVERSITY OF CALIFORNIA

CF Energy Consumption Tradeoff

- CF consumes less energy on **storage** than CR.
  - Industry standard CR redundancy level: 3 replicas per content.
    - ✓ 2x more storage usage.
  - Industry standard CF redundancy level: RS(10, 4).
    - ✓ 0.4x more storage usage.
- CF could potential consume more energy on **core network** and **content decoding** than CR.
  - Content retrieving follows reverse-multicast routing.
  - Content decoding consumes extra energy in CF if there is any data block missing.

UCDAVIS
UNIVERSITY OF CALIFORNIA

W. Wang, "Saving capacity with HDFS RAID," *Facebook Code*, 2014.

The Goals of This Study

- We propose and formulate the **content placement and routing assignment problem** with the goal to **minimize energy consumption** while **satisfying QoS constraints such as resiliency and latency** for both CF and CR schemes.

- This problem is solved optimally by MILP. Due to its complexity, we also propose an efficient meta-heuristic algorithm (M-CPRA).

- We provide guidelines on how to choose which scheme to use based on three factors: **(1) number of content requests**, **(2) resiliency**, and **(3) latency**.

CF Content Retrieving

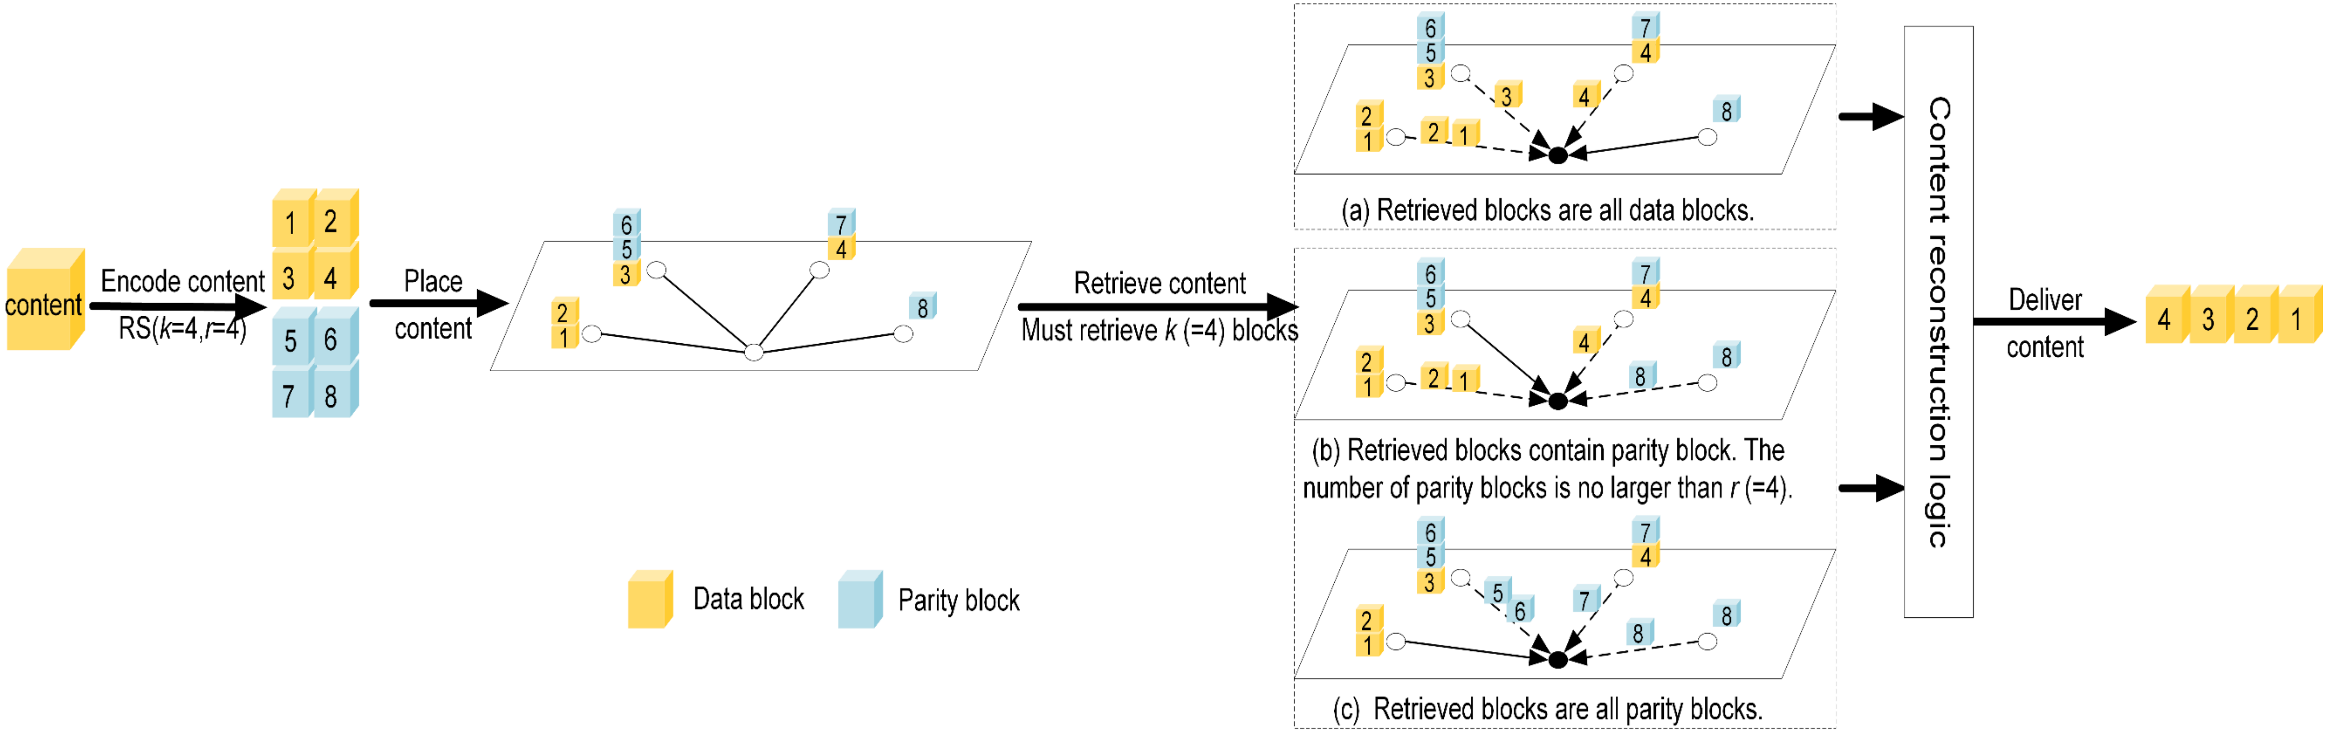

CF Content Decoding

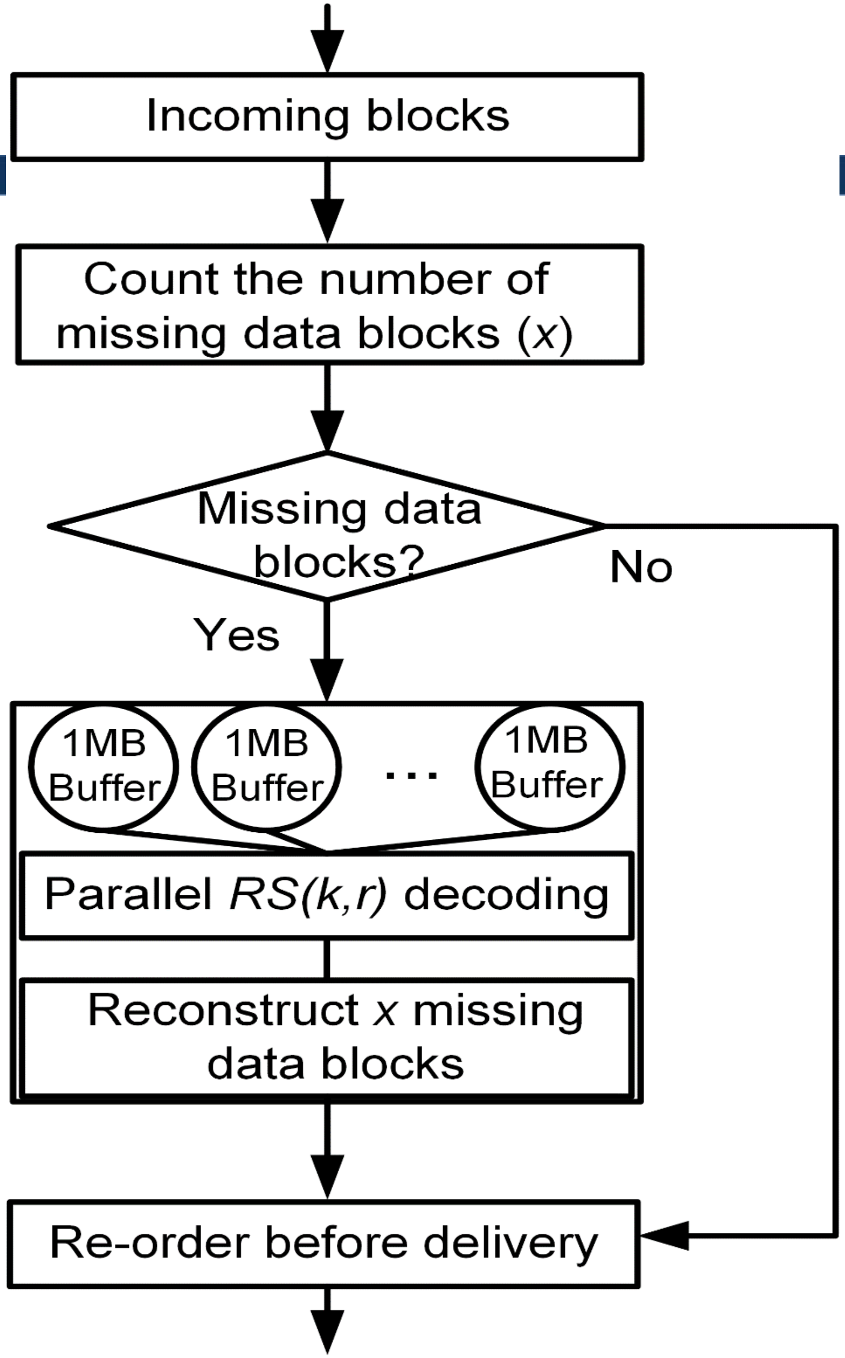

UCDAVIS
UNIVERSITY OF CALIFORNIA

Energy Consumption

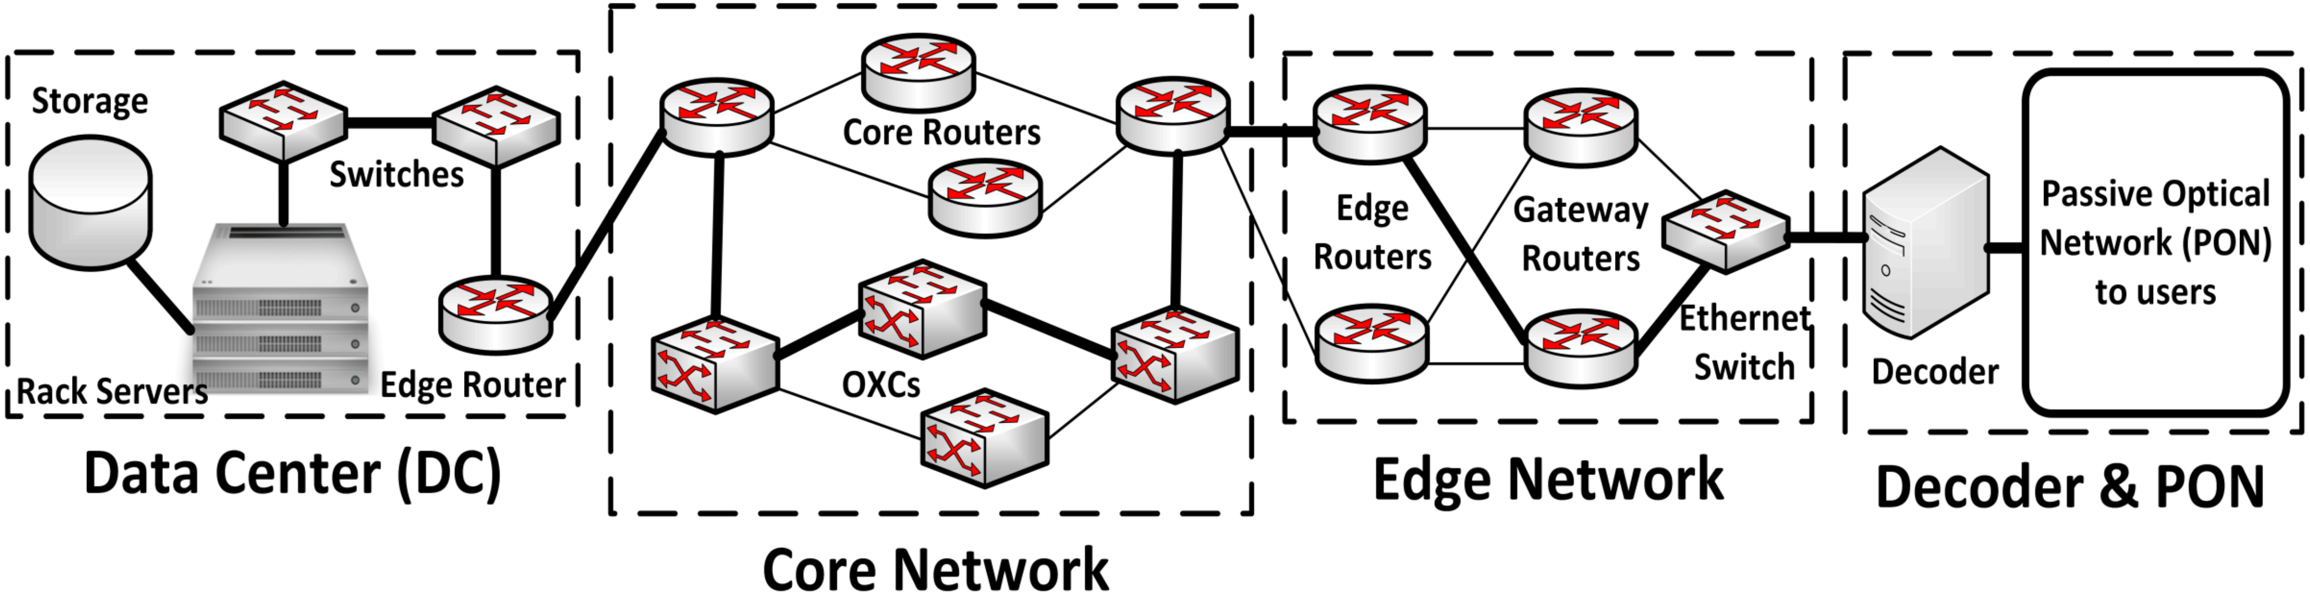

UCDAVIS
UNIVERSITY OF CALIFORNIA

QoS Constraint: Latency

- Latency Island
  - A set of nodes within the reach of the current node.

- Latency
  - All of the serving blocks (data/parity) need to be placed within latency island.

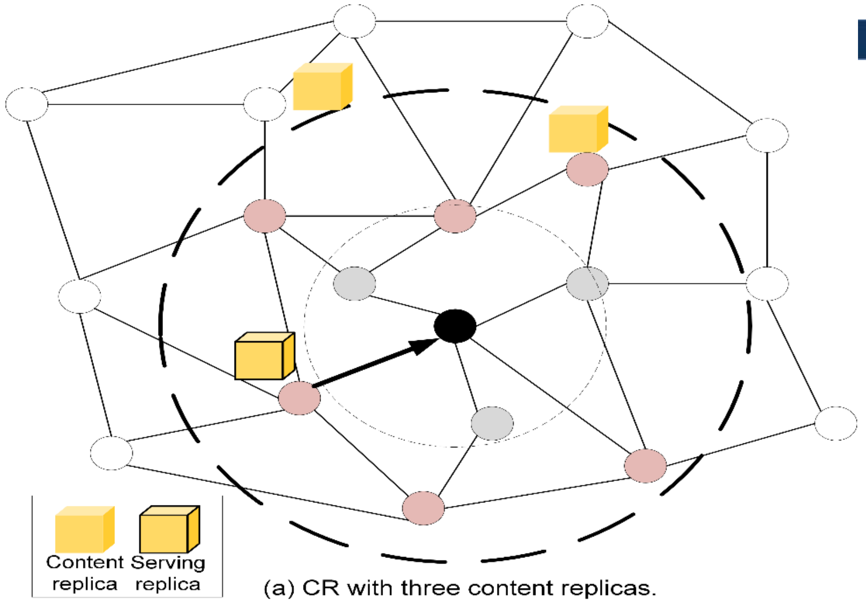

(a) CR with three content replicas.

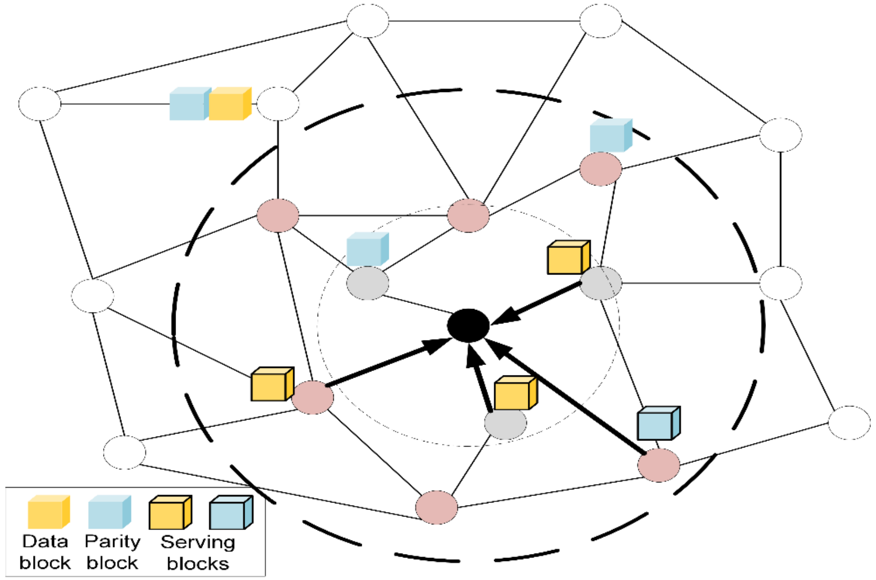

(b) CF with RS(k=4, r=4) content encoding.

UC DAVIS
UNIVERSITY OF CALIFORNIA

QoS Constraint: Resiliency

- We consider resiliency for single-node failure.

- If a content-hosting node is down, we need a backup plan for serving blocks at failure node so that the service is not disrupted.

UCDAVIS
UNIVERSITY OF CALIFORNIA

QoS Constraint: Resiliency

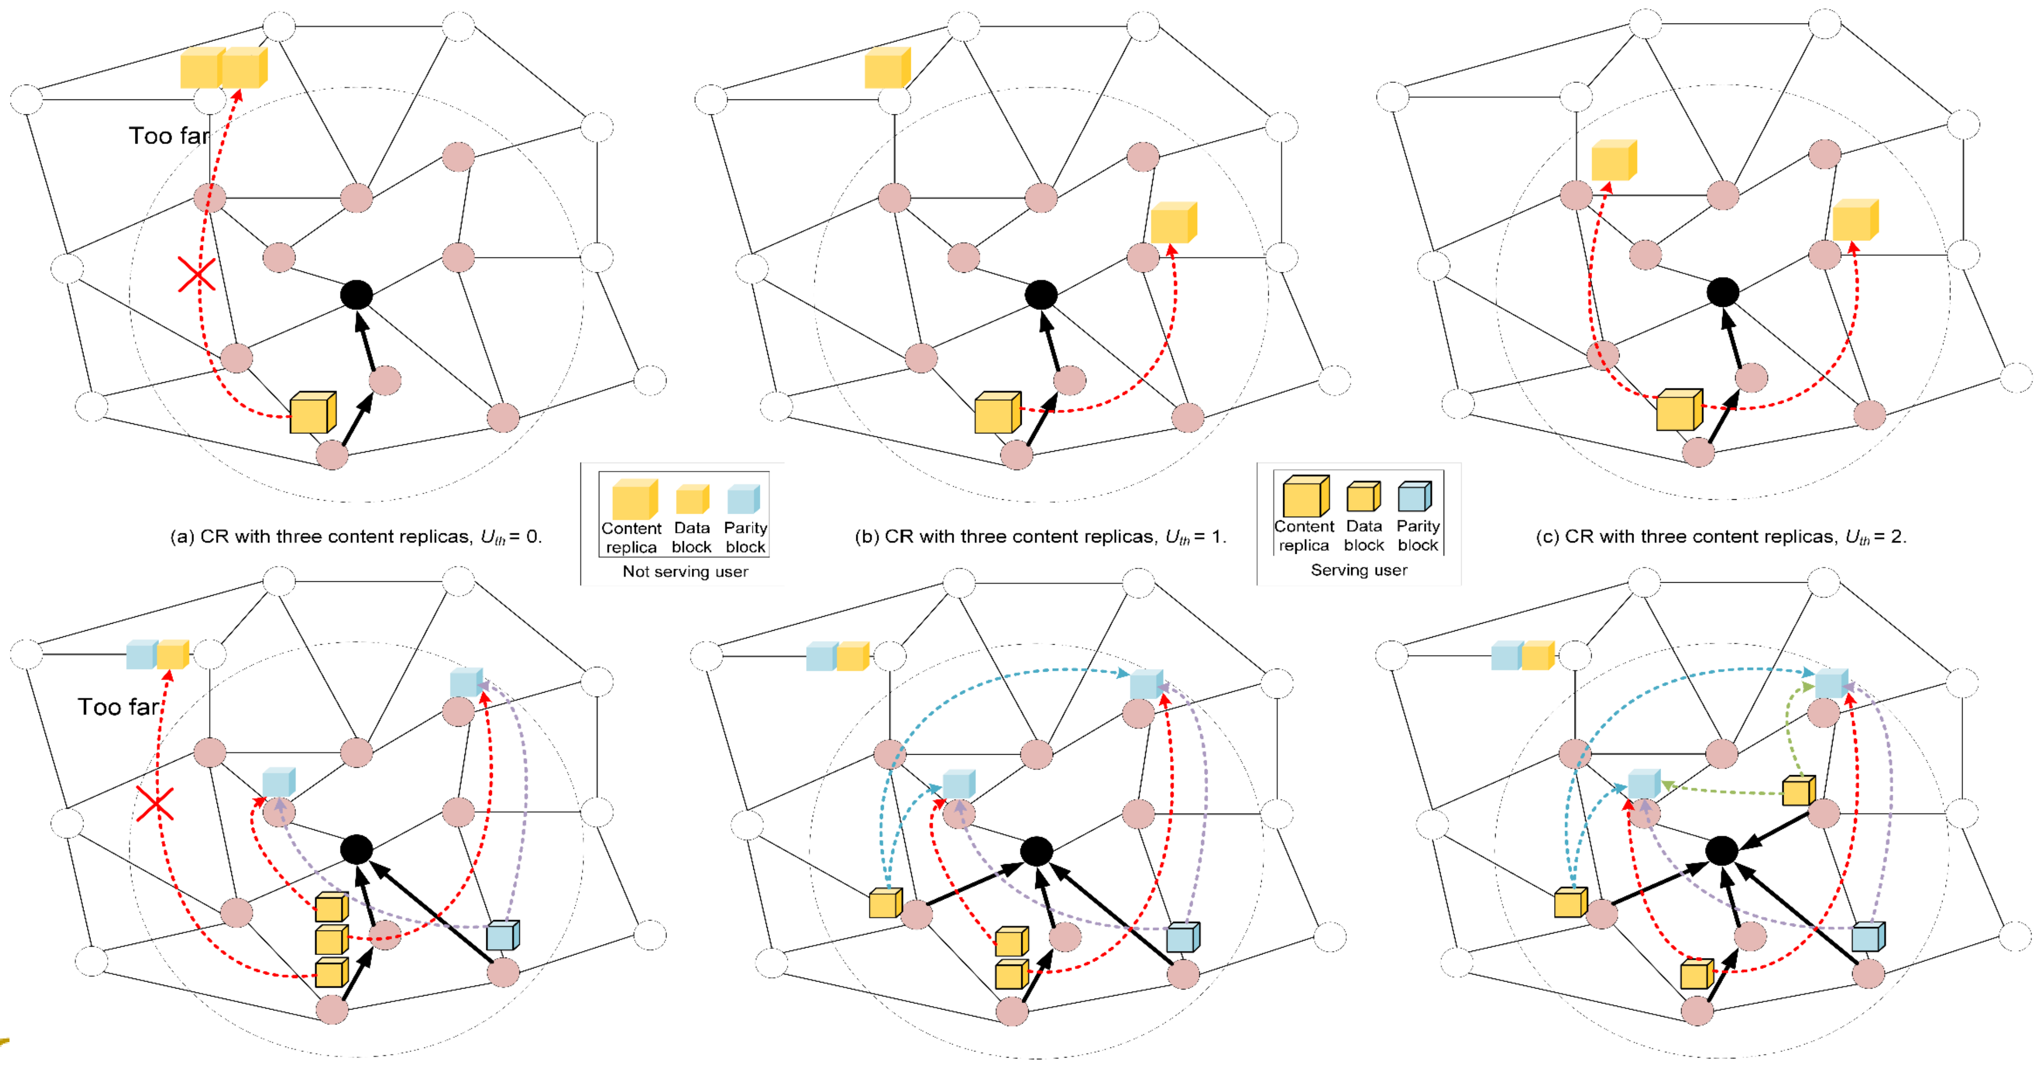

(a) CR with three content replicas, $U_{th} = 0$.

(b) CR with three content replicas, $U_{th} = 1$.

(c) CR with three content replicas, $U_{th} = 2$.

(d) CF with RS(k=4, r=4) content encoding, $U_{th} = 0$.

(e) CF with RS(k=4, r=4) content encoding, $U_{th} = 1$.

(f) CF with RS(k=4, r=4) content encoding, $U_{th} = 2$.

M-CPRA Algorithm

Probabilistically Rearrange Placement (PRP)

Input: Local optimal placement from last neighborhood (replica/block-hosting DC set (HS) + set of replicas/blocks placed at DC d in HS (ϕ_d));

Output: resulting content placement (HS^{*}+ ϕ^{*}_{d} *(equal* HS *and* ϕ_d *initially)*).

Randomly generate number of replicas/blocks to be rearranged ($\omega^{(\cdot)} \leq |I^{(\cdot)}|$)

repeat

Count the number of LIs each DC d belongs to in HS' (δ_d).

Calculate the probability of DC d in HS' being selected to remove a content replica/block. k_1 and k_2 are coefficients used to scale the sum to be 1.

$$\Pr^{R}(d) = k_1 \cdot \frac{1/\delta_d}{\sum_{j \in HS'} 1/\delta_j} + k_2 \cdot \frac{\theta_d}{\sum_{j \in HS'} \theta_j} \quad \forall d \in HS' \tag{19}$$

Probabilistically select a DC d from HS' according to the above distribution, remove the replica/block with least index number from it, and update HS' and ϕ'_d if necessary.

until $\omega^{(\cdot)}$ replicas/blocks get removed

for each removed replica/block

Use same procedures to calculate $\Pr^{R}(d)$ for each DC d in DC set DS.

Calculate the probability of DC d in DS being selected to add a content replica/block.

$$\Pr^{A}(d) = \frac{1 - \Pr^{R}(d)}{|DS| - 1} \quad \forall d \in DS \tag{20}$$

Probabilistically select a DC d from DS according to the above distribution, add the removed replica/block to it, and update HS' and ϕ'_d if necessary.

end

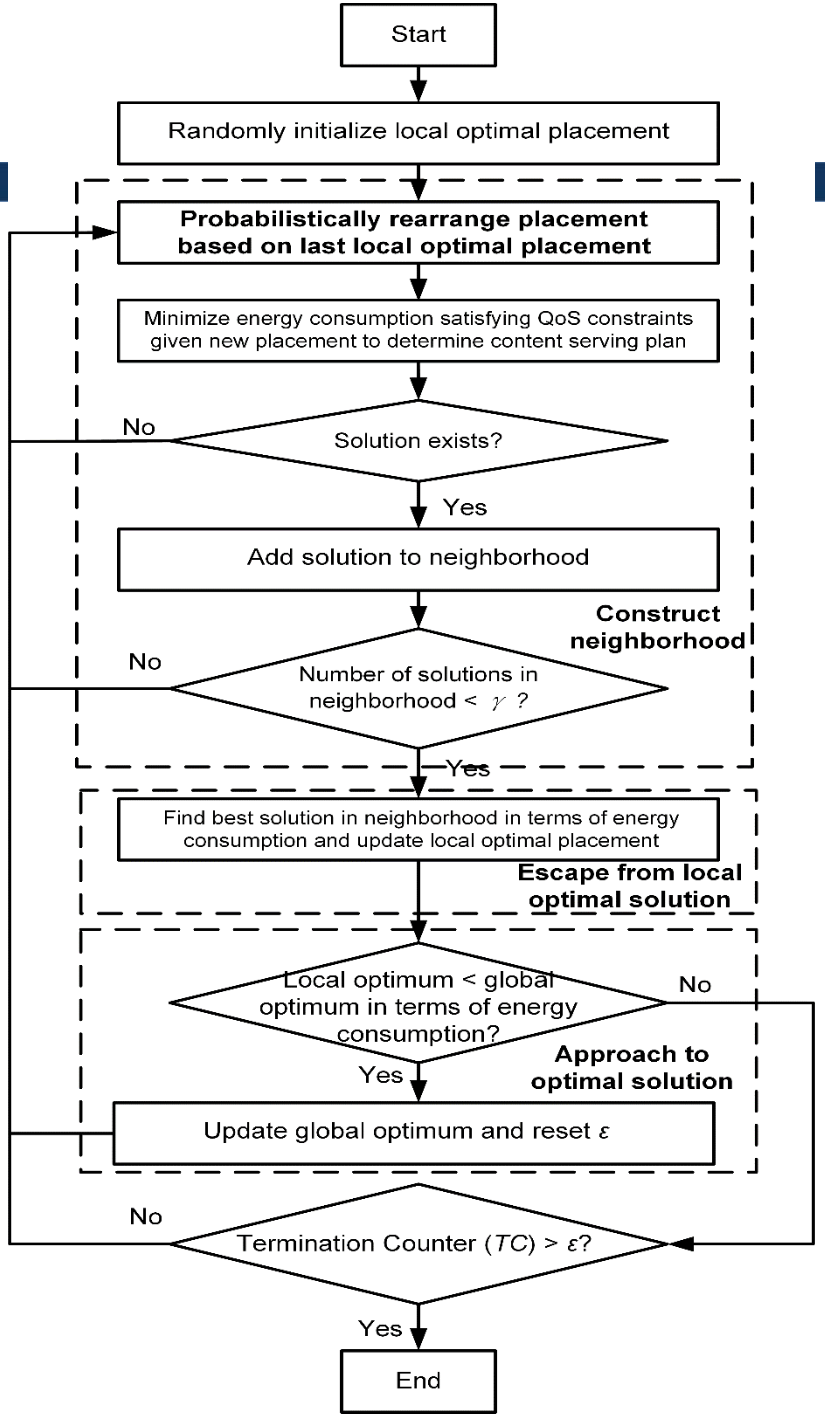

UC DAVIS
UNIVERSITY OF CALIFORNIA

CF or CR

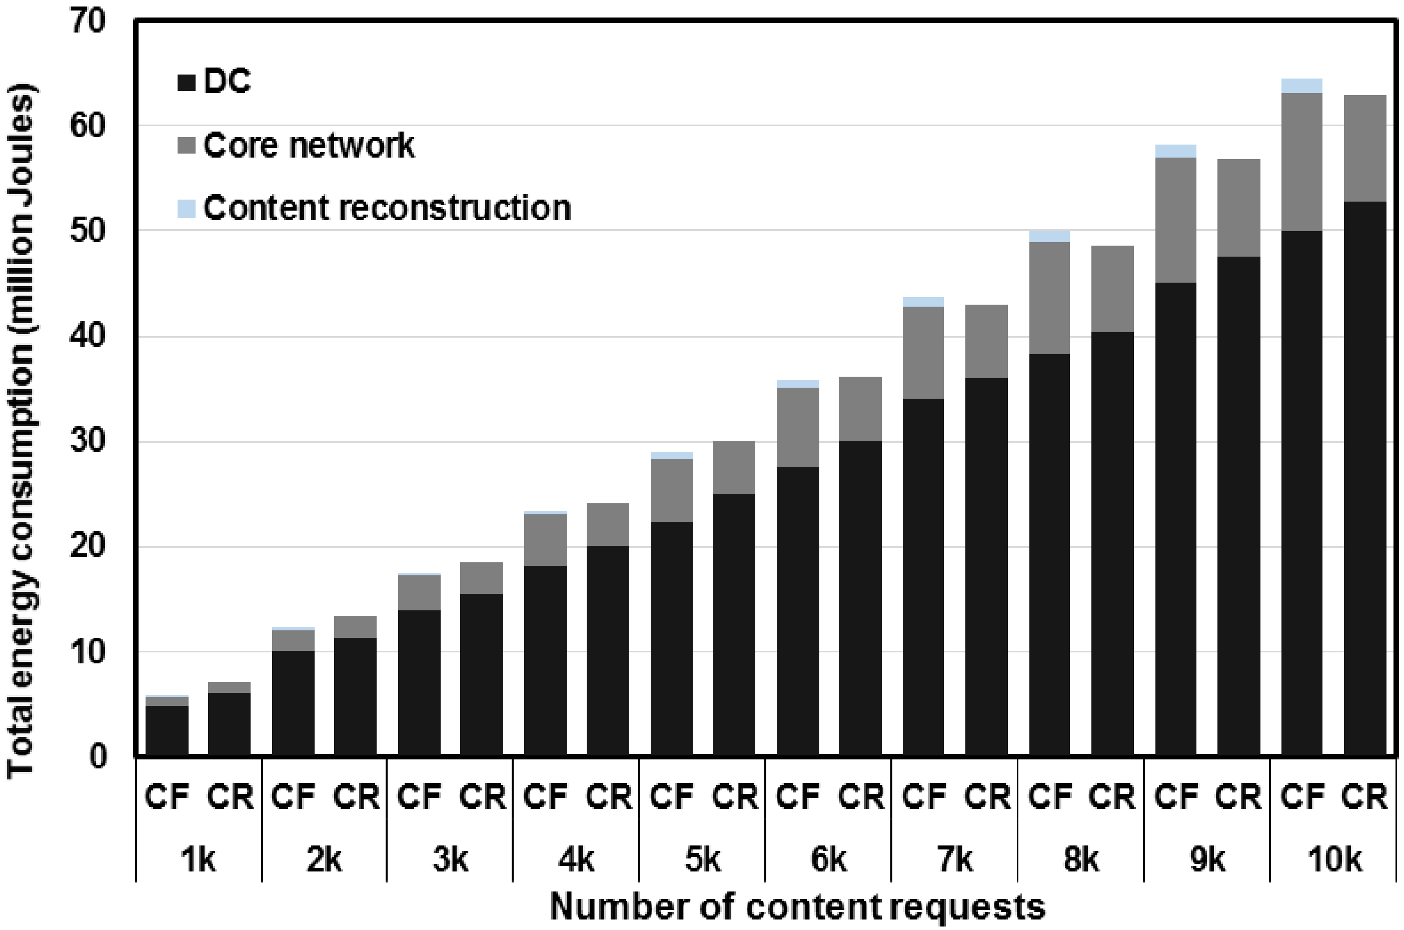

UCDAVIS
UNIVERSITY OF CALIFORNIA

CF or CR

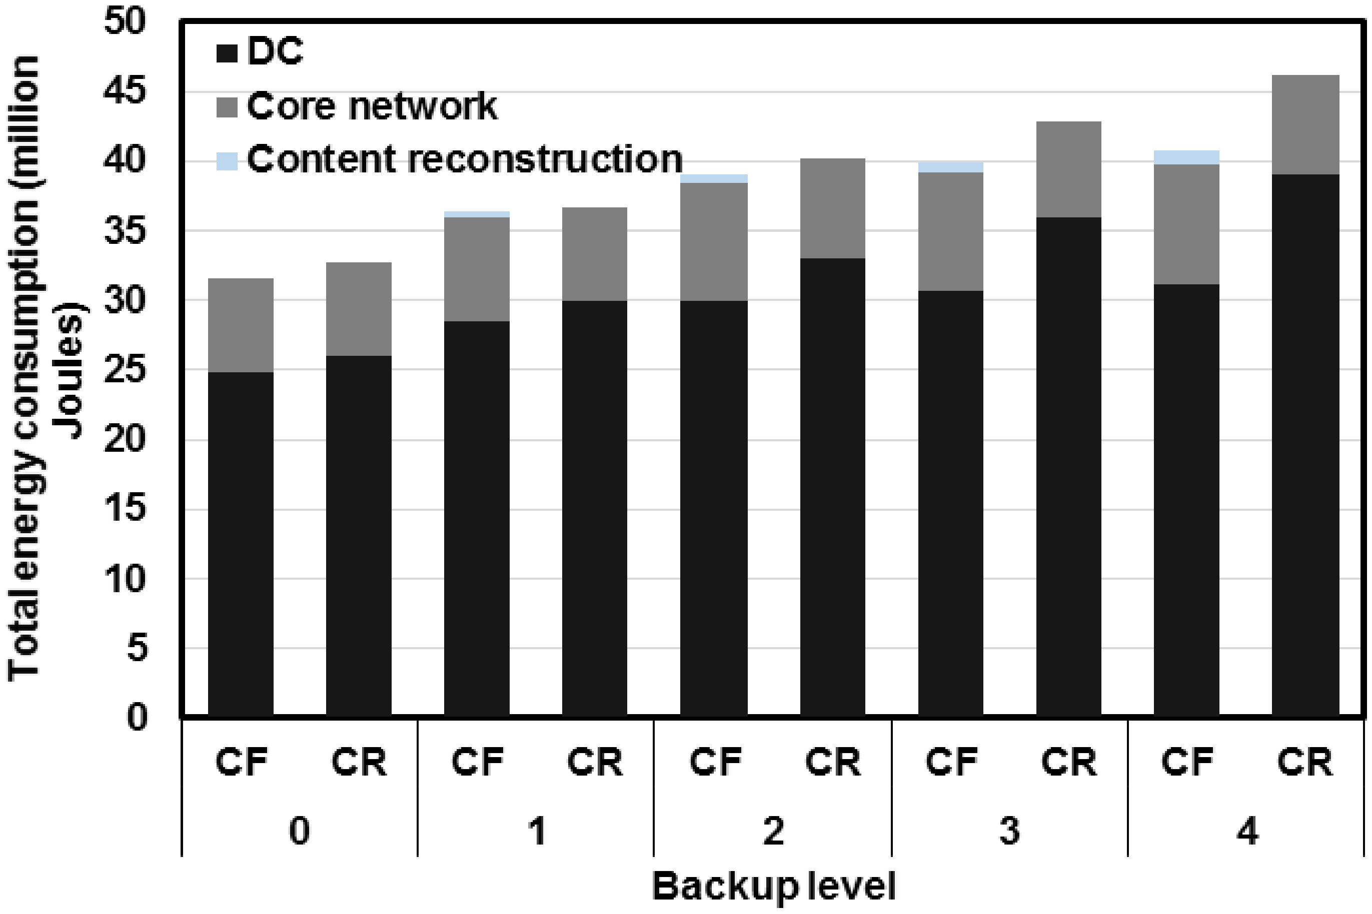

UC DAVIS
UNIVERSITY OF CALIFORNIA

CF or CR

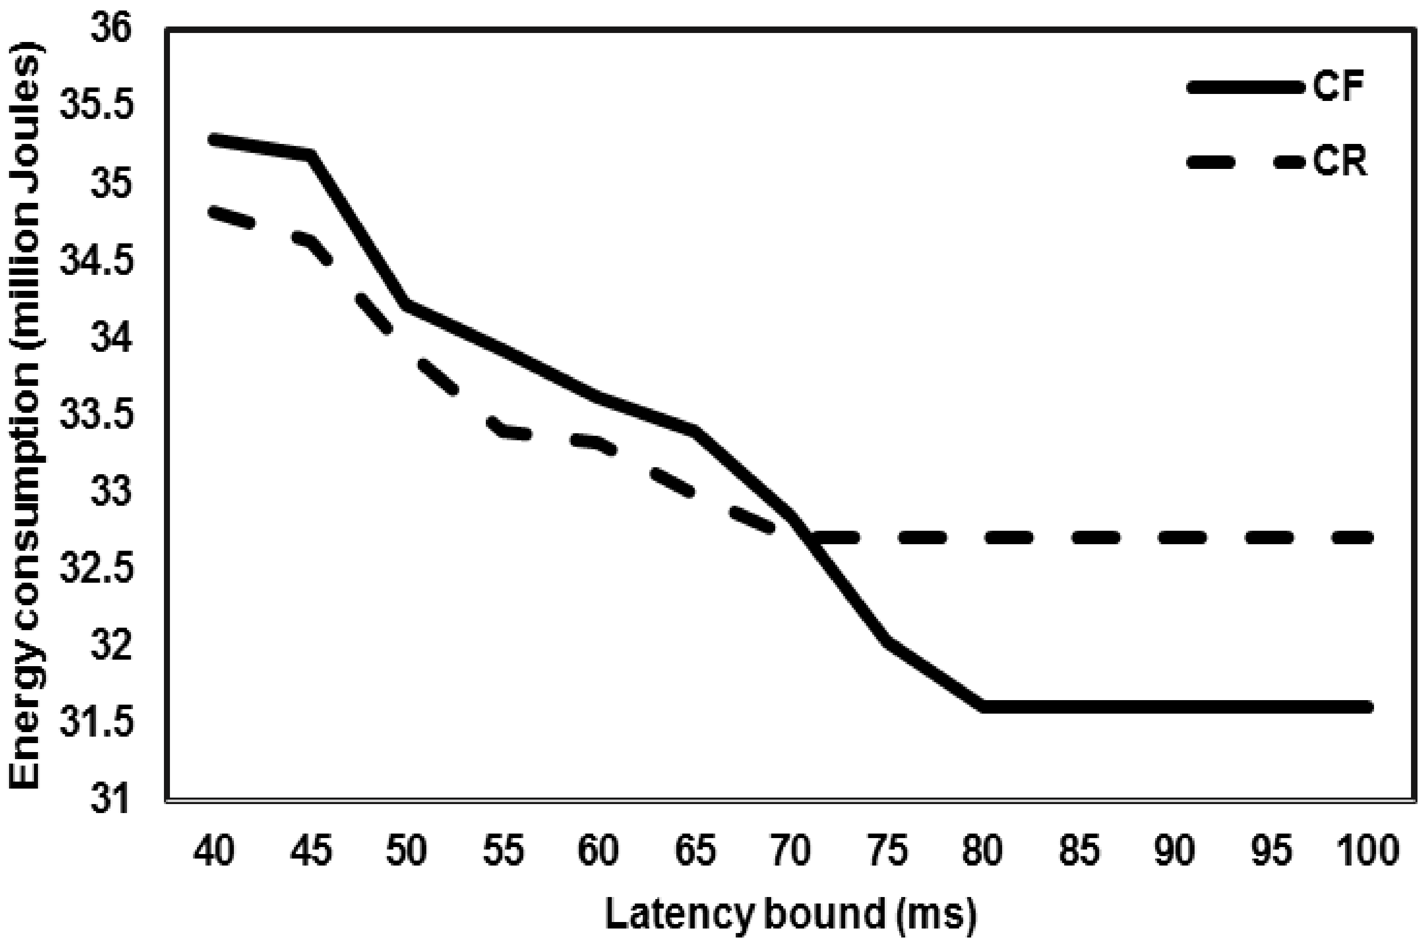

UCDAVIS
UNIVERSITY OF CALIFORNIA

Thank you!