Coded Caching: Global vs Local Content Popularity

Abhin Shah
MIT

Nikhil Karamchandani
IIT Bombay

Suhas Diggavi
UCLA

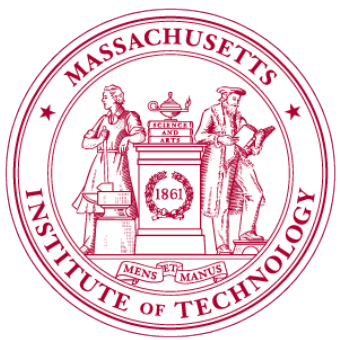

Work done at IIT Bombay

System model

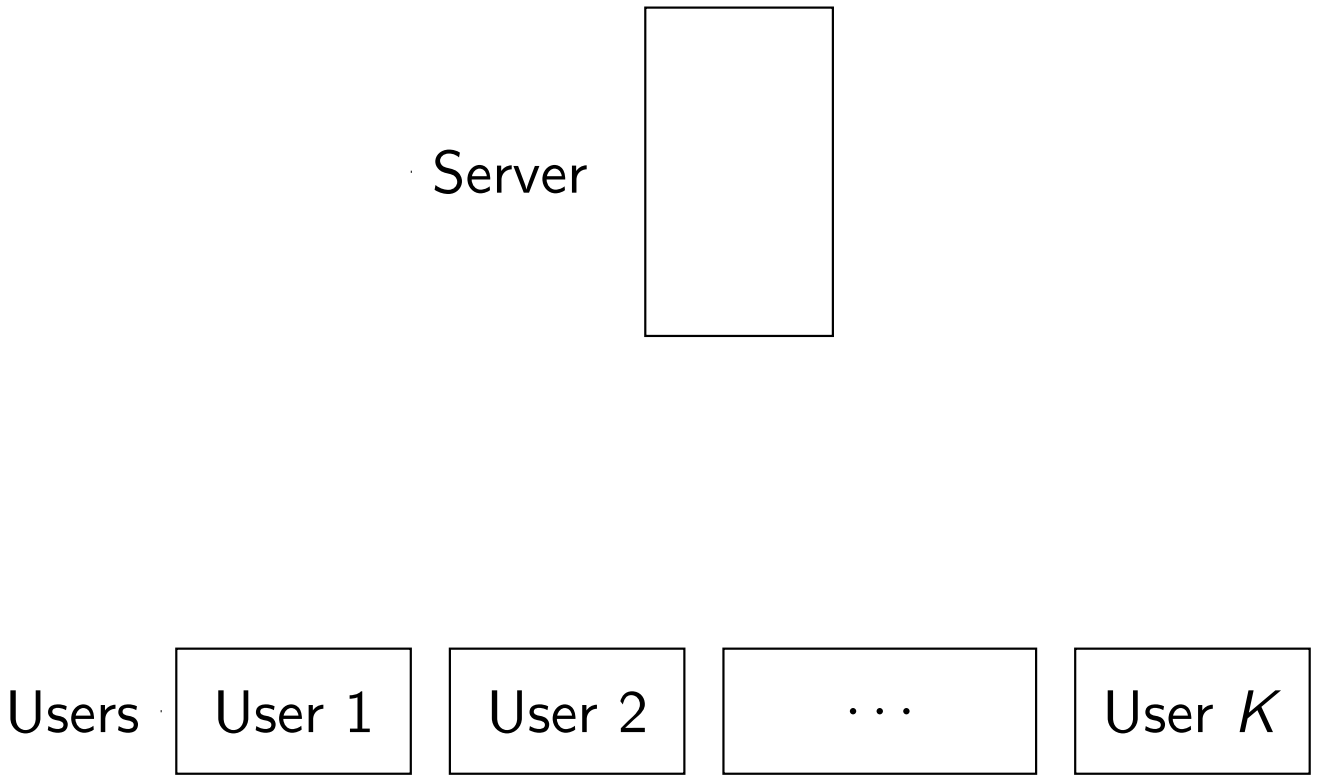

System model

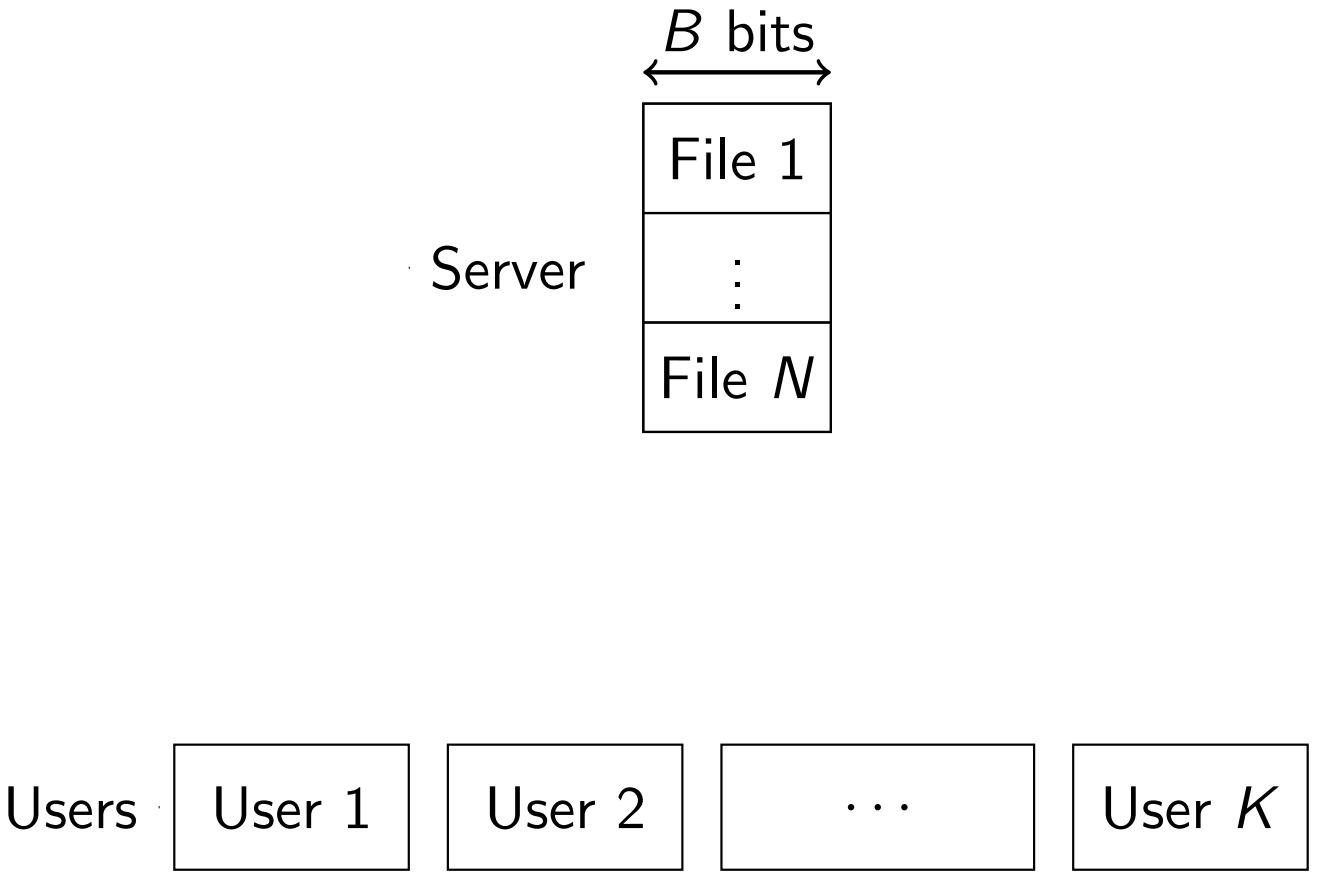

System model

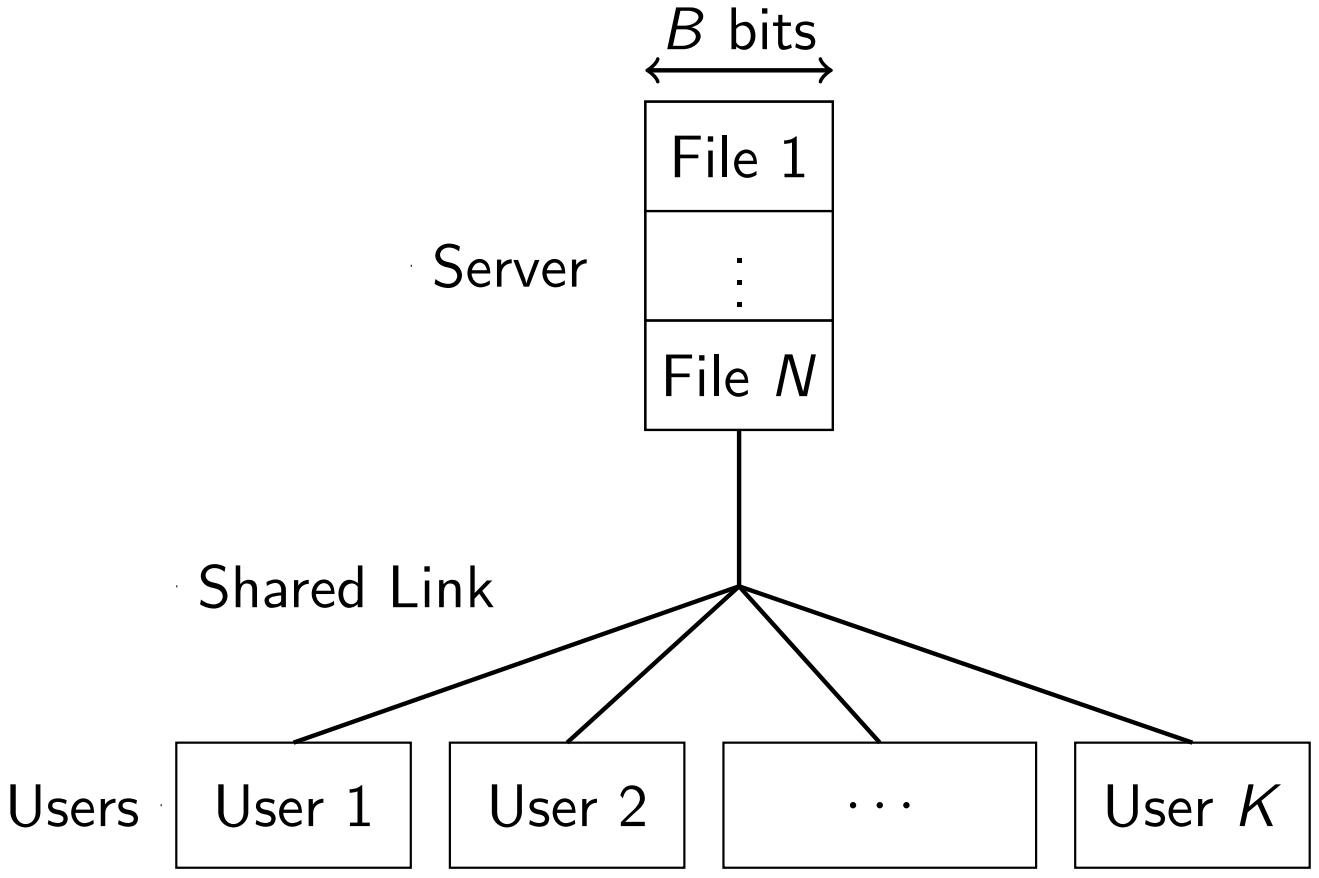

System model

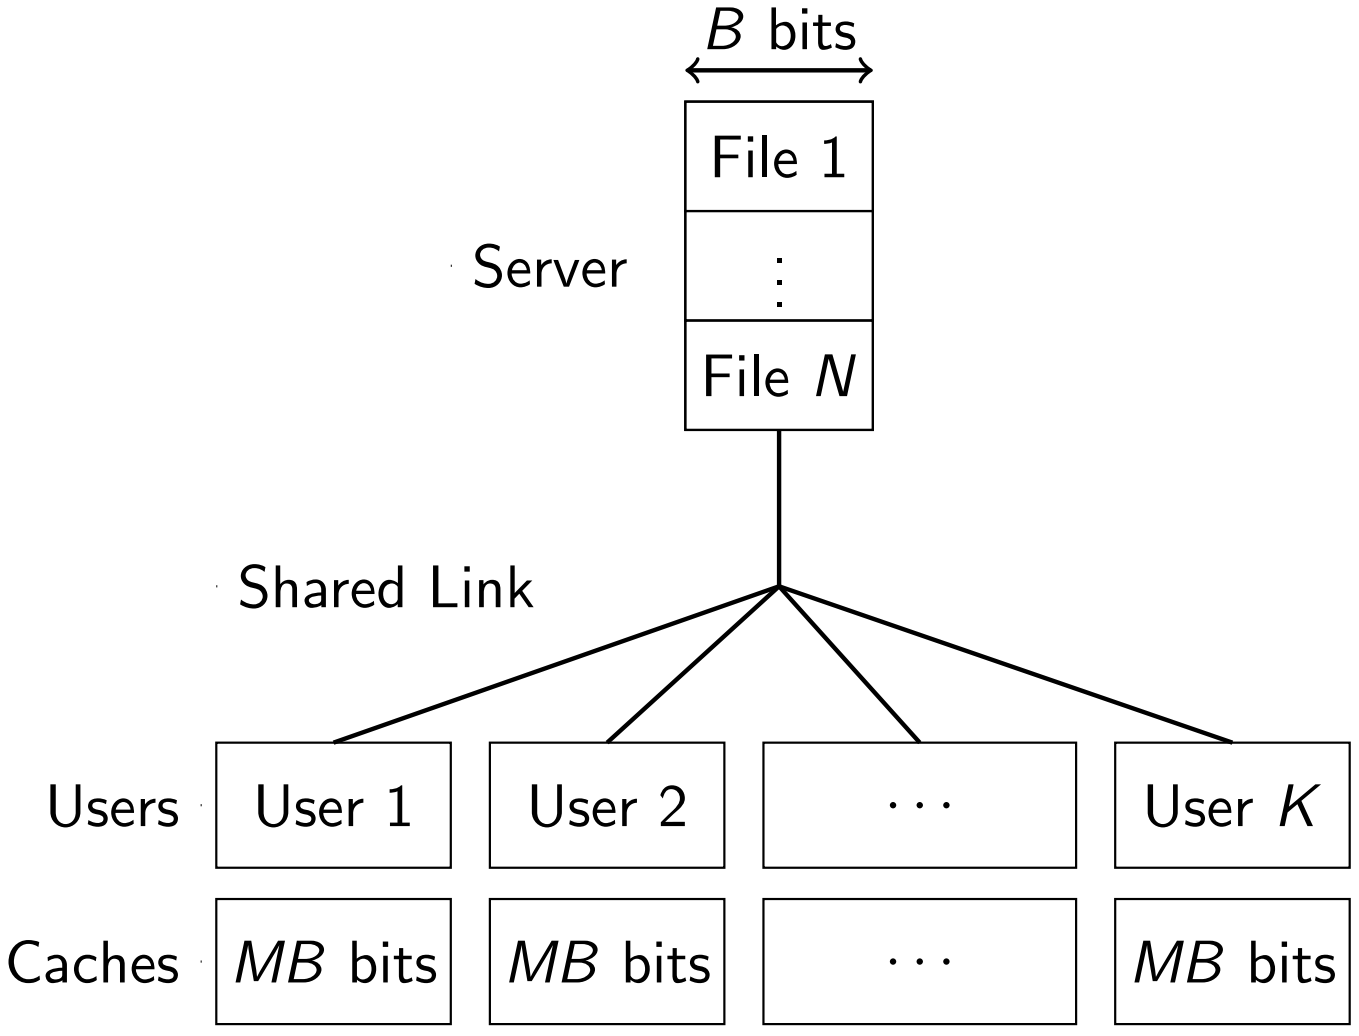

System model

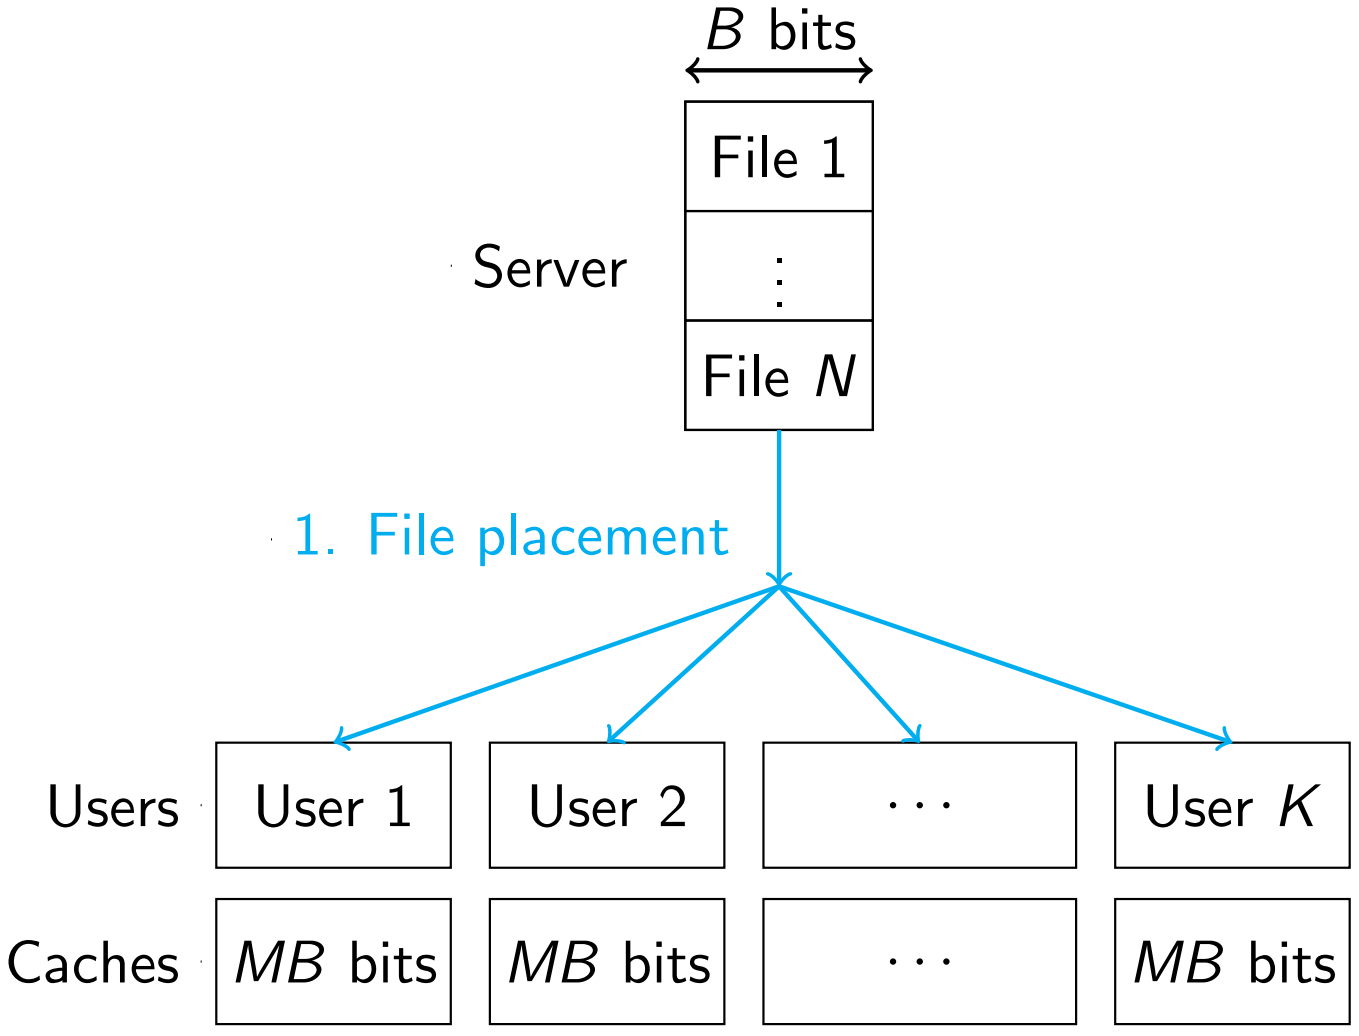

System model

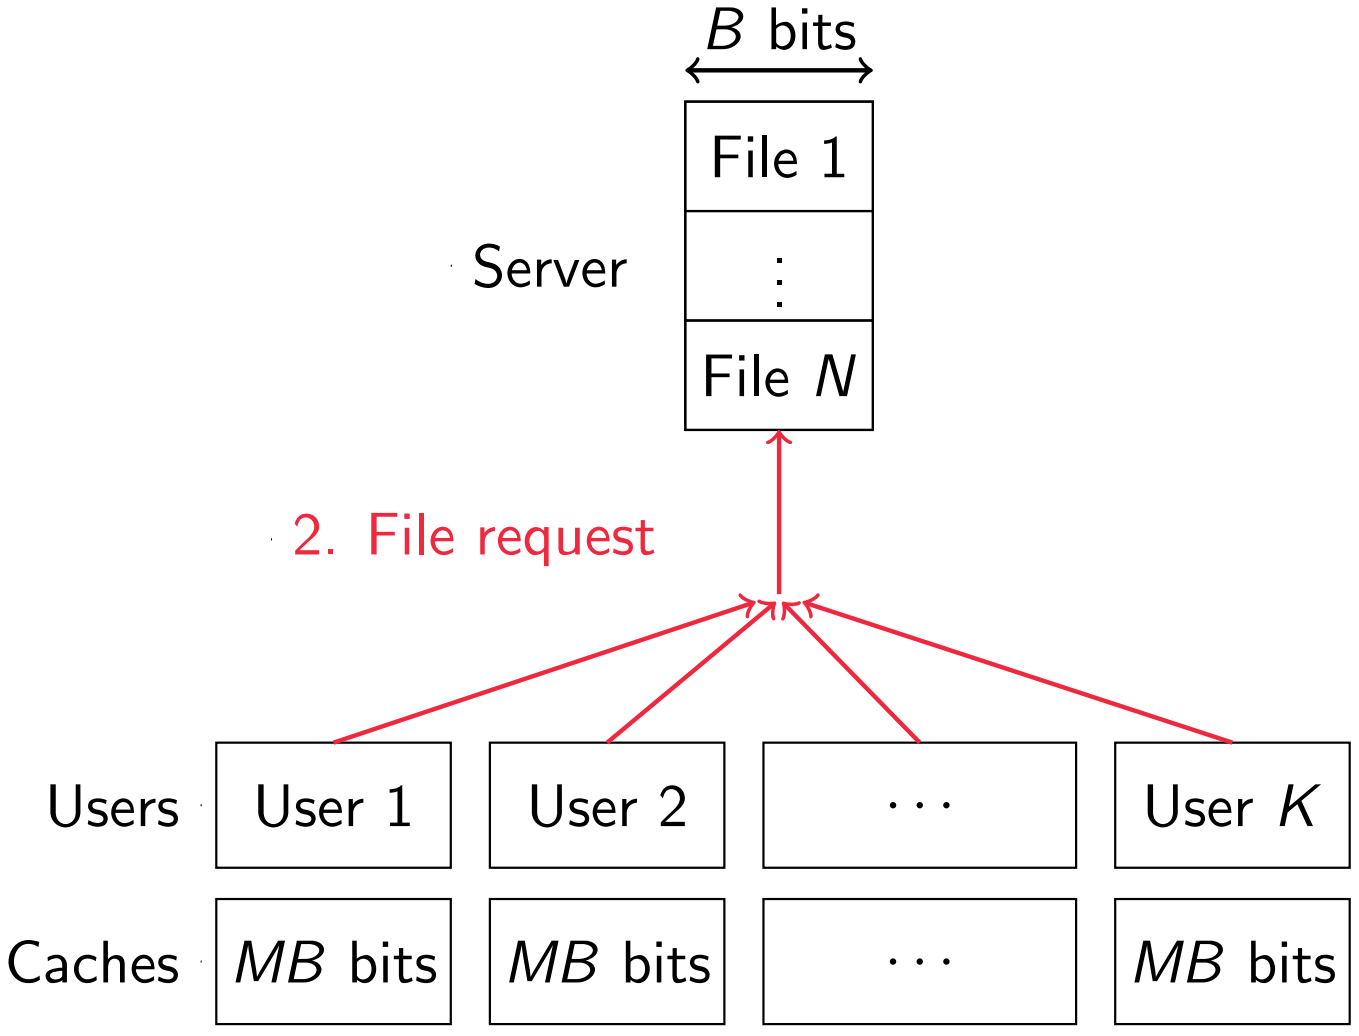

System model

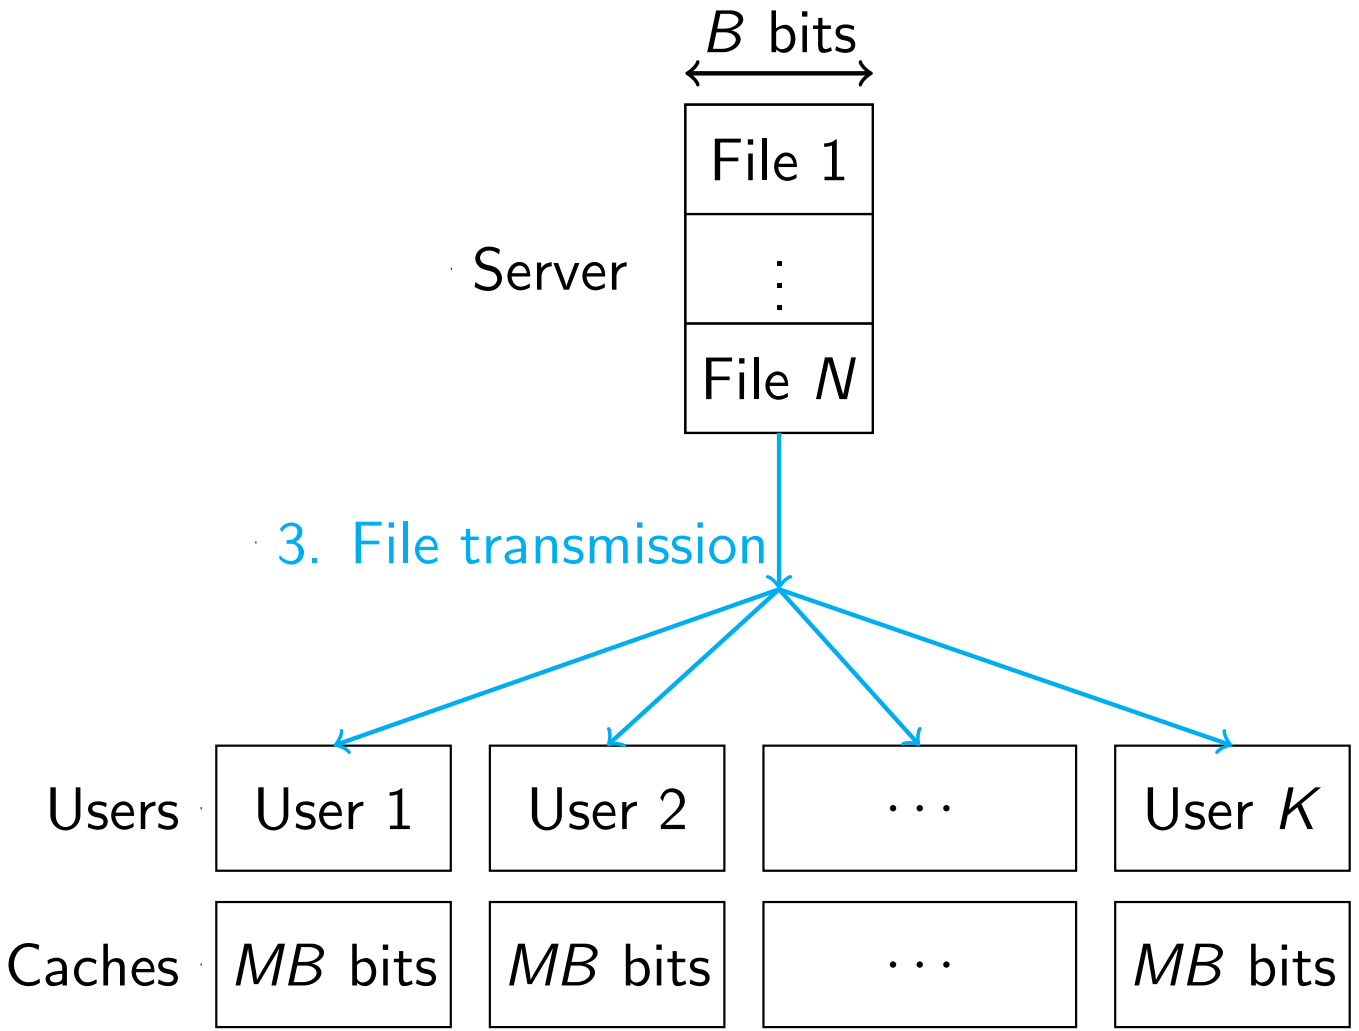

Goal

- The worst-case rate

Goal

- The worst-case rate
  $\rightarrow$ Maximum size of server transmission required over all possible user demands.

Goal

- The worst-case rate
  $\rightarrow$ Maximum size of server transmission required over all possible user demands.
- The goal

Goal

- The worst-case rate
  $\rightarrow$ Maximum size of server transmission required over all possible user demands.
- The goal
  $\rightarrow$ Minimize the worst-case transmission rate over the shared link.

Traditional Uncoded Caching

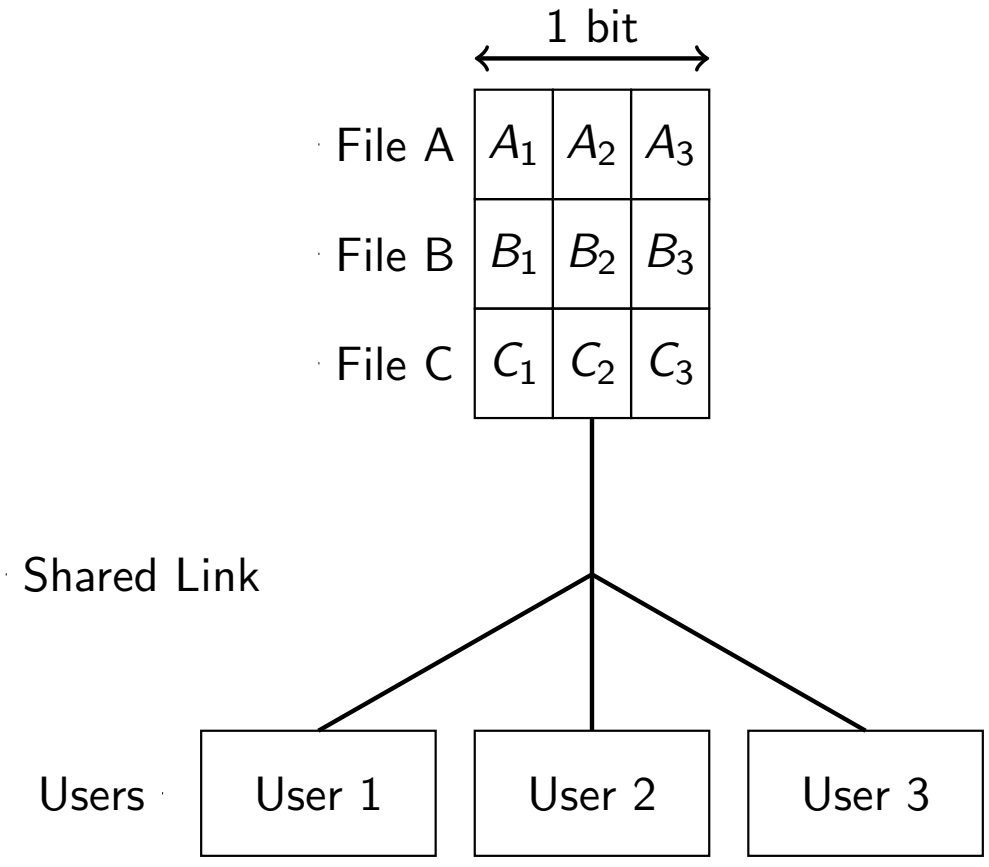

Traditional Uncoded Caching

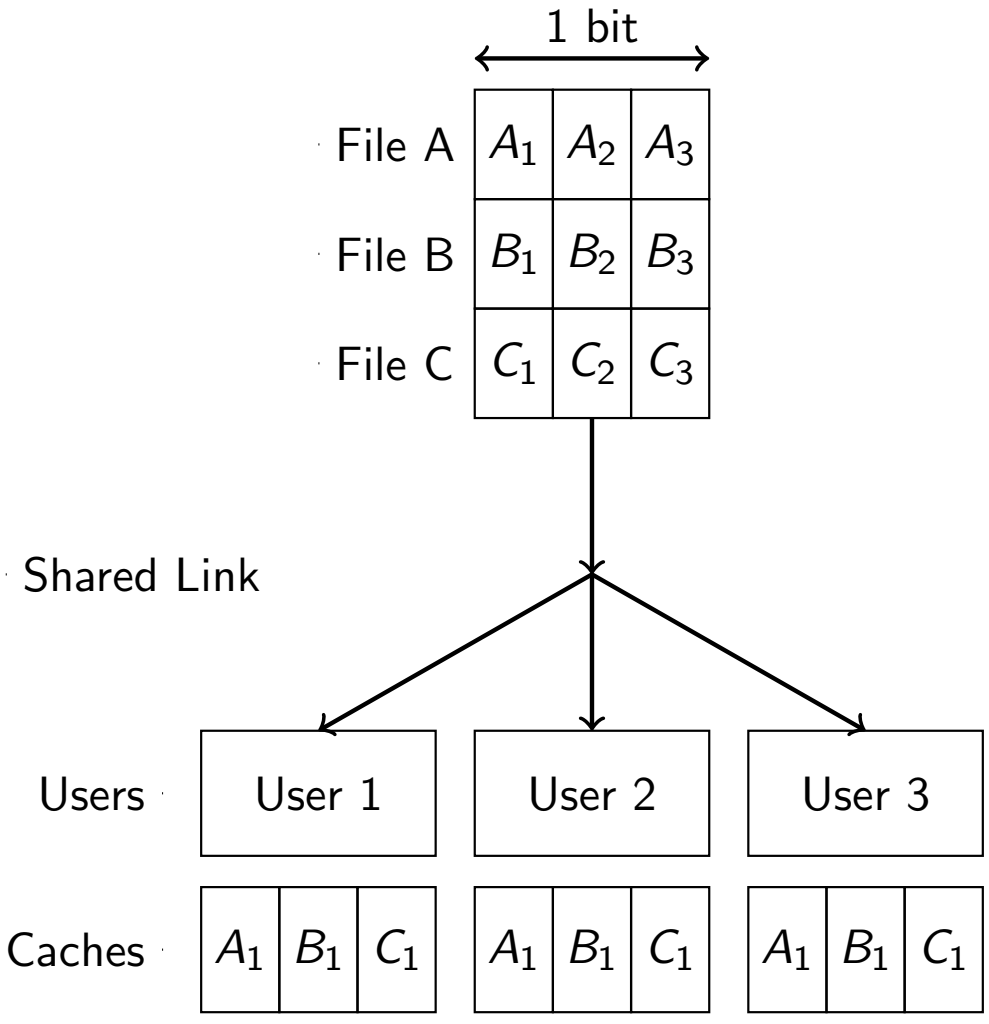

Traditional Uncoded Caching

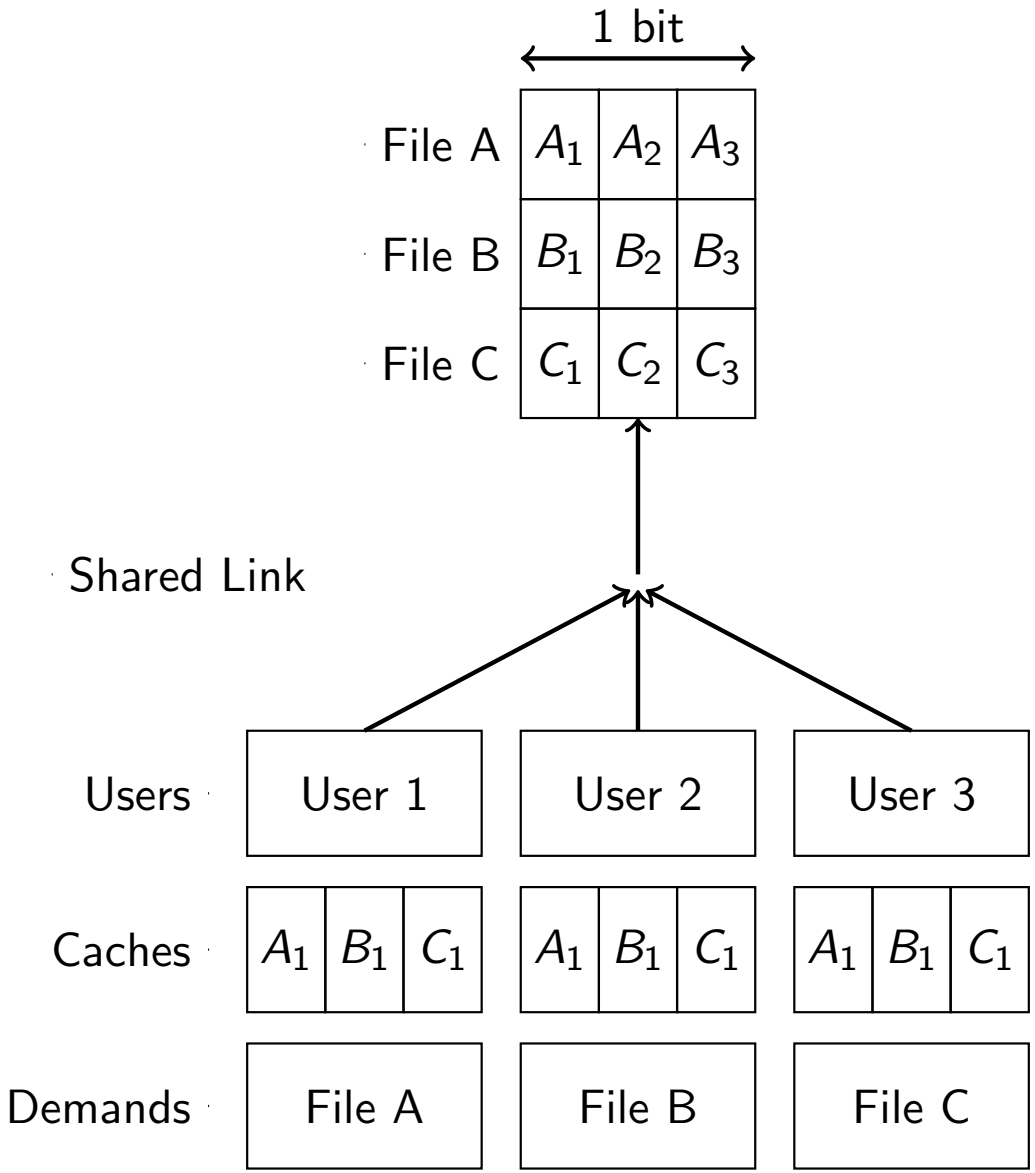

Traditional Uncoded Caching

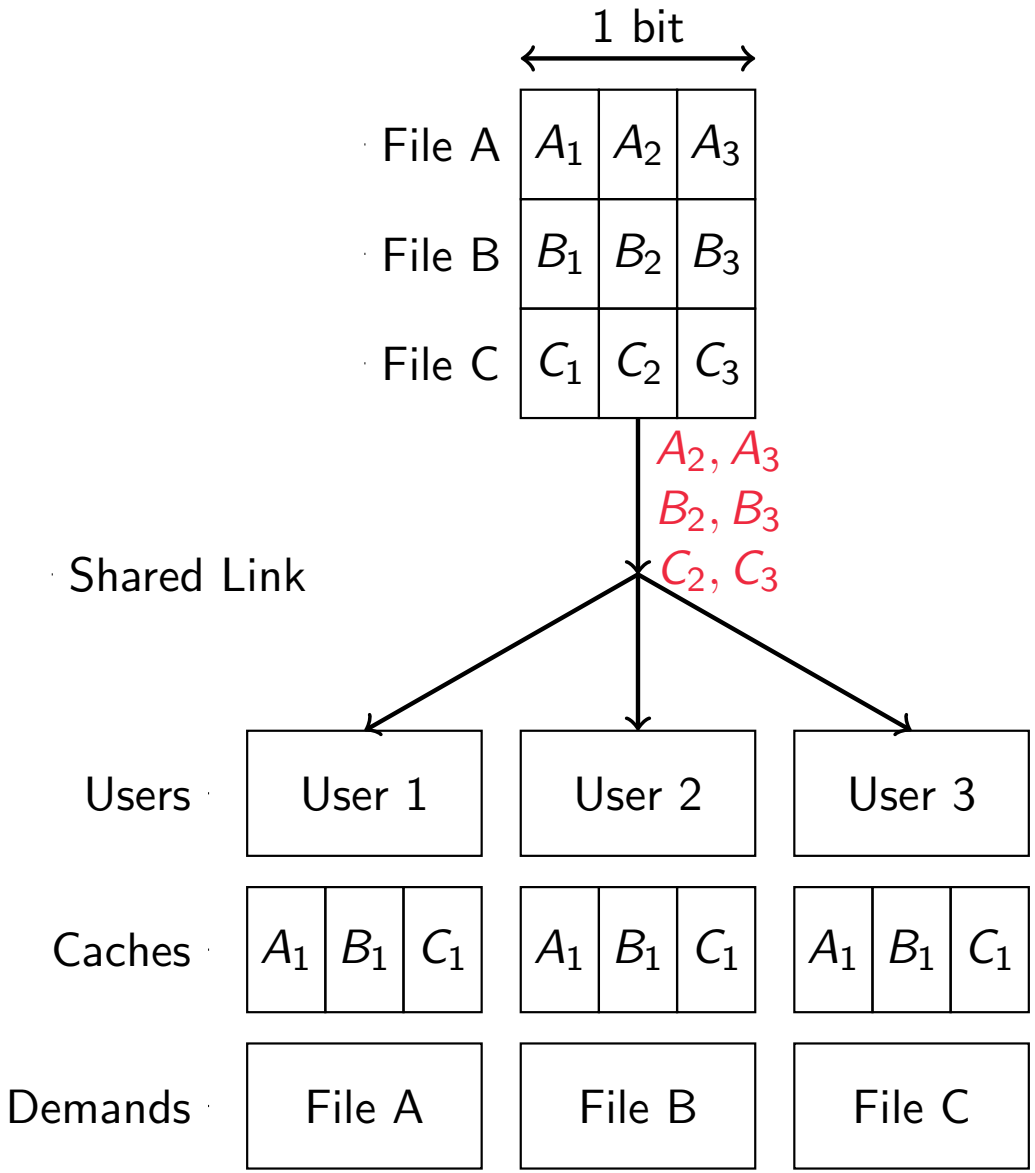

Traditional Uncoded Caching

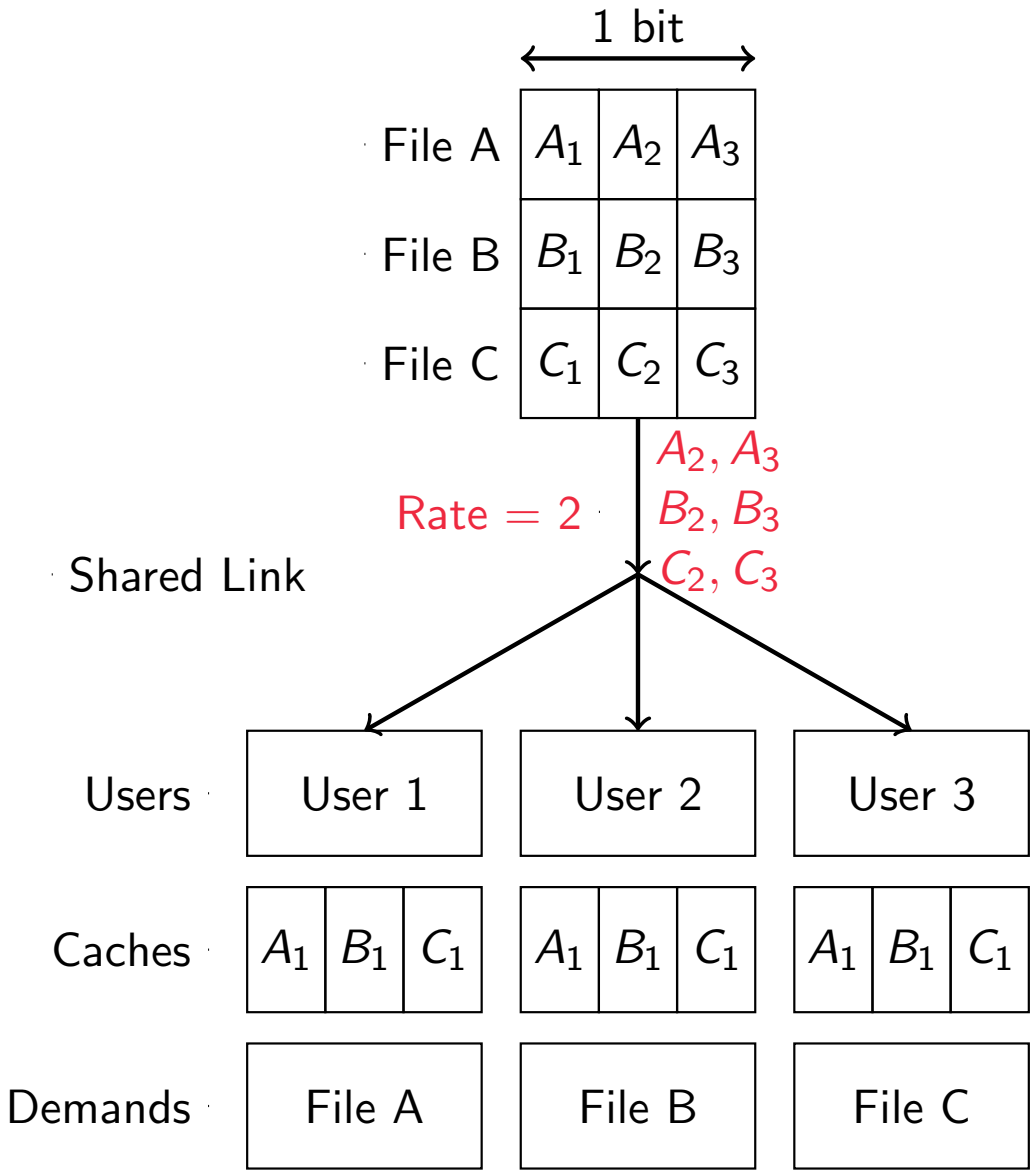

Coded Caching (Maddah-Ali and Niesen)

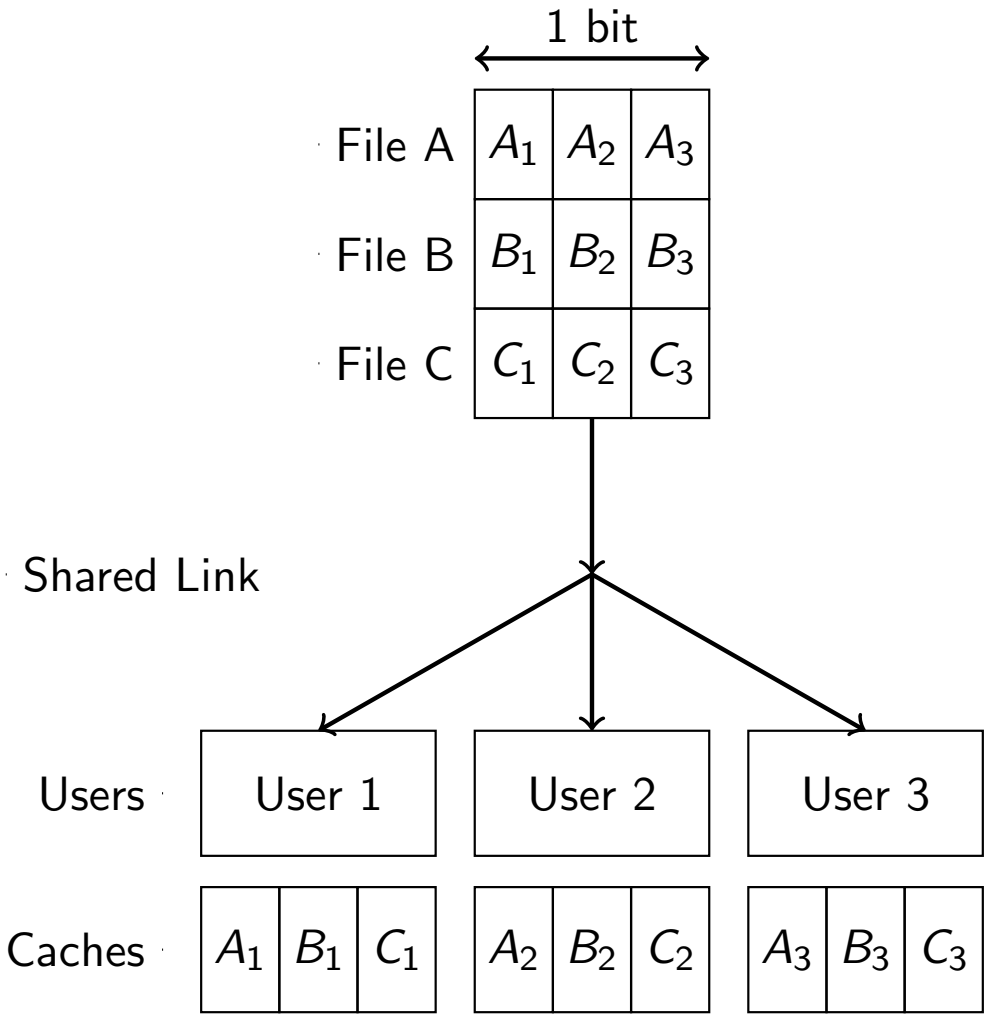

Coded Caching (Maddah-Ali and Niesen)

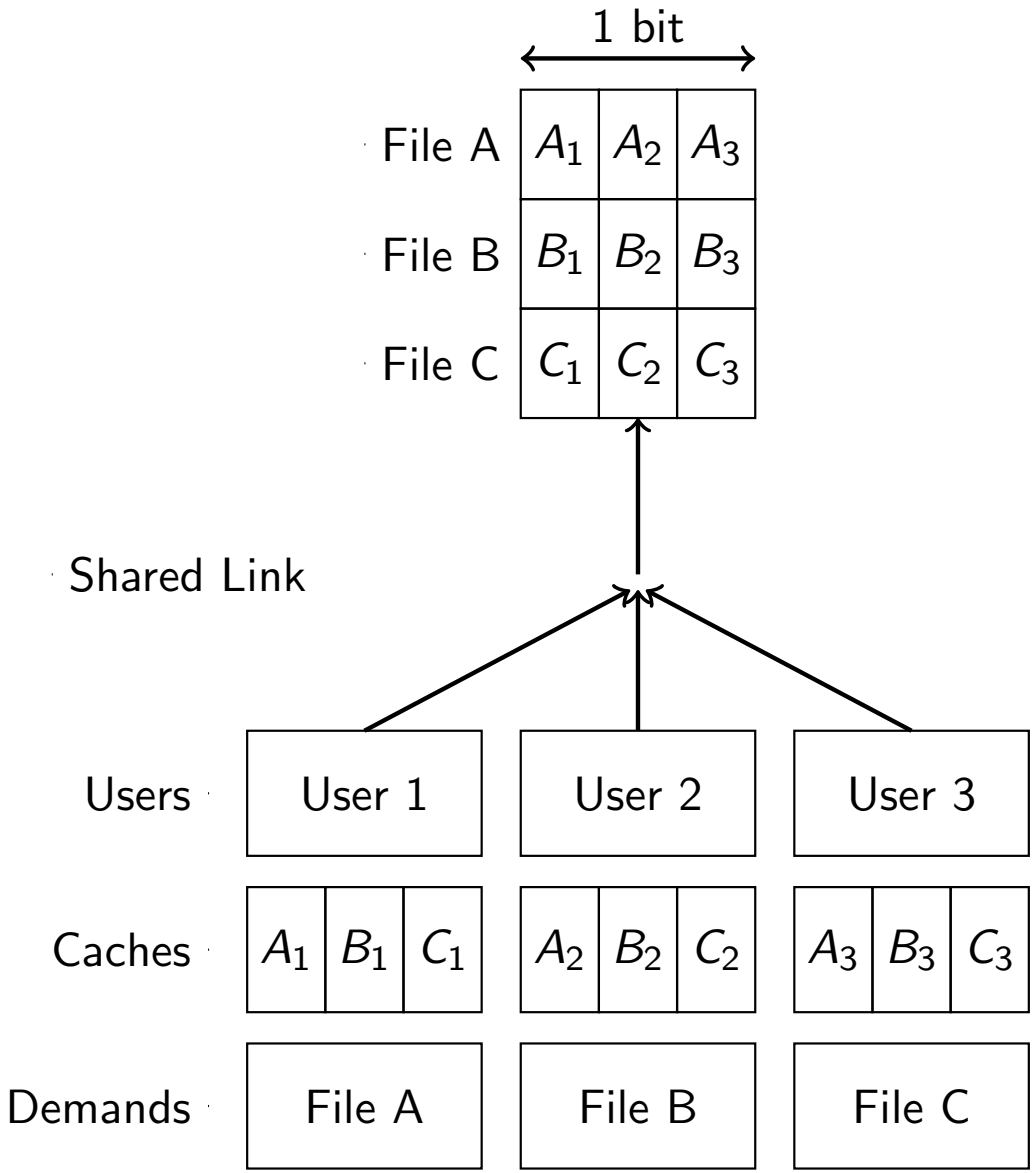

Coded Caching (Maddah-Ali and Niesen)

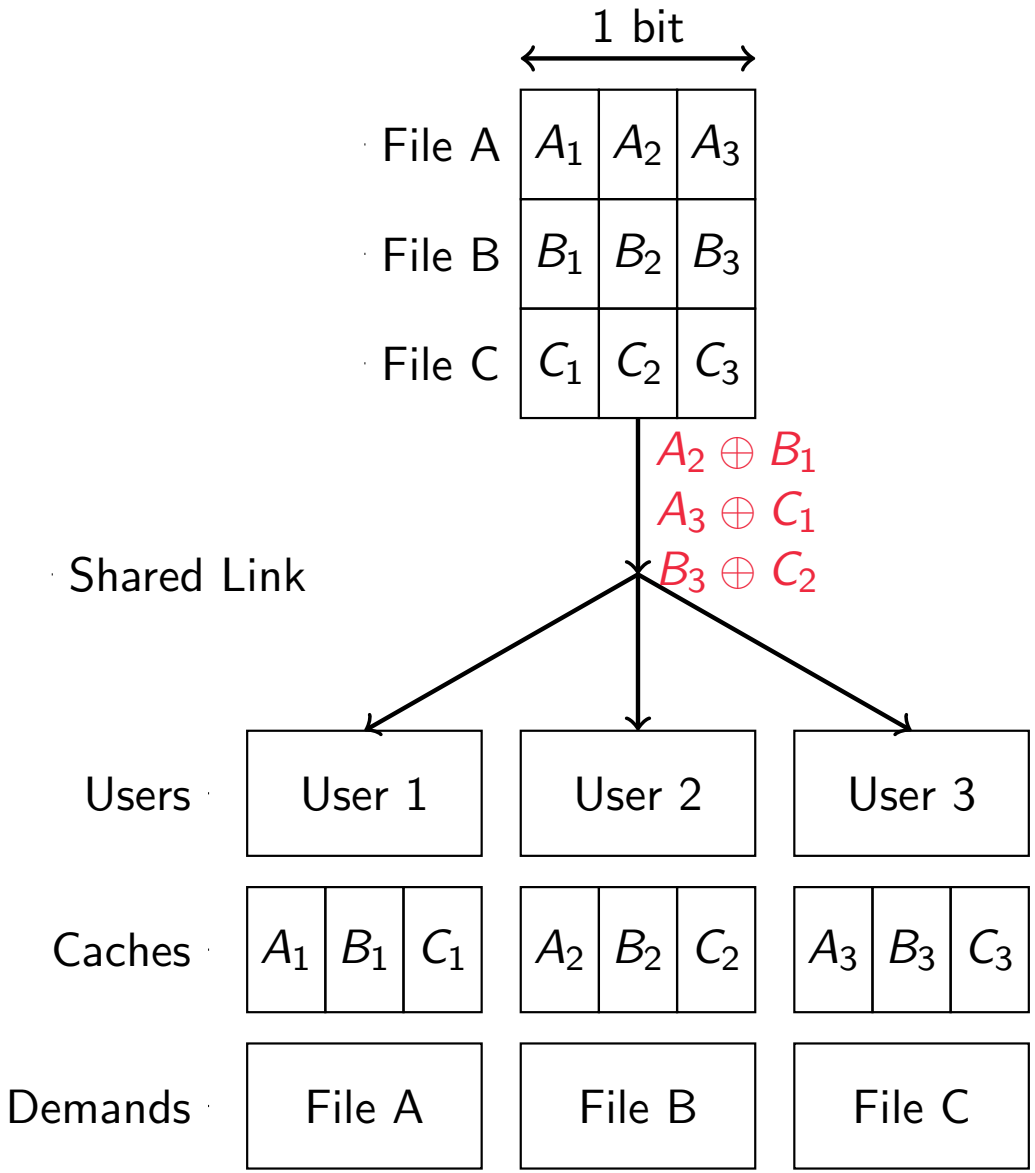

Coded Caching (Maddah-Ali and Niesen)

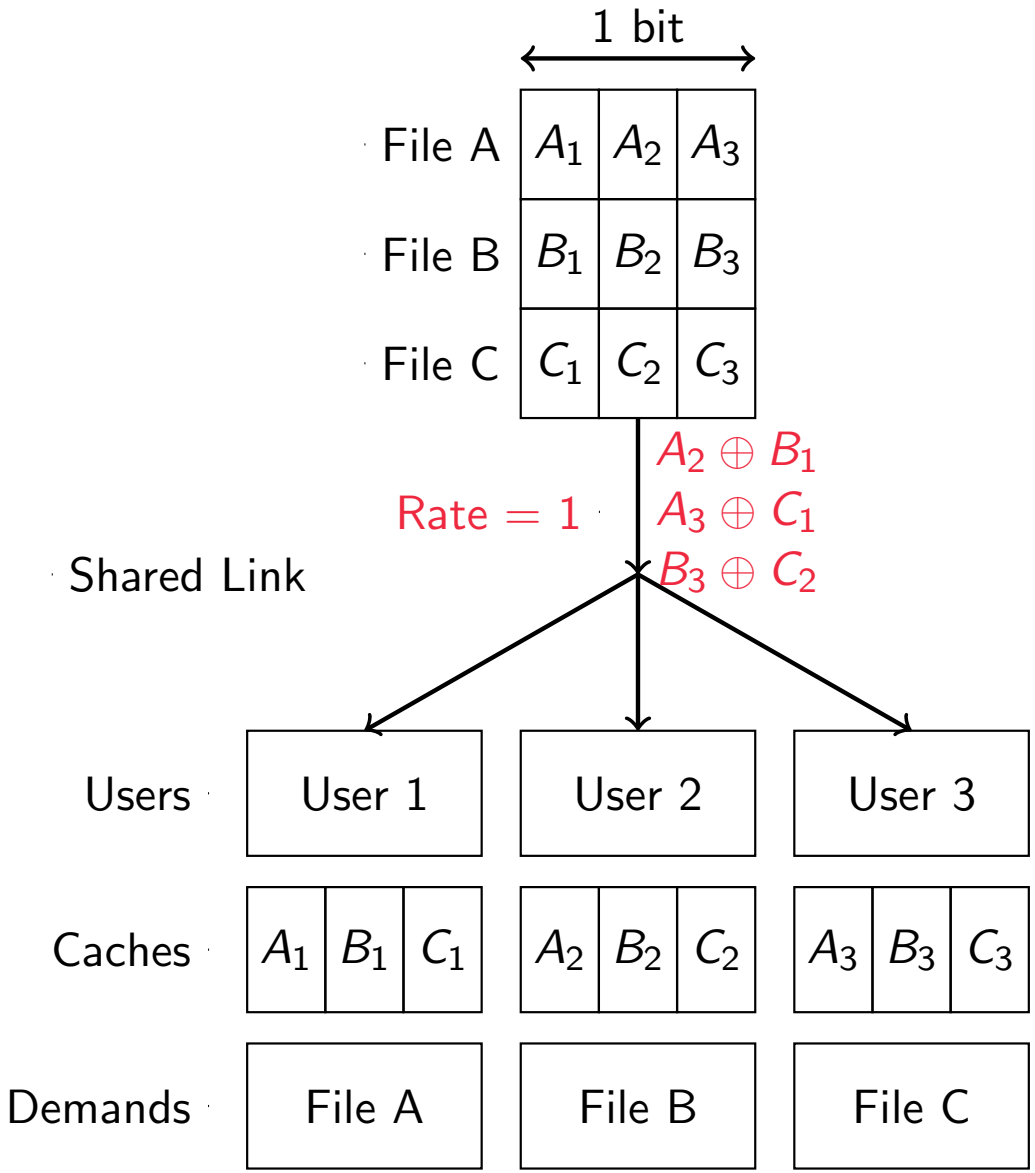

Different setups

Setup	File popularity	Across Users

Different setups

Setup	File popularity	Across Users
Decentralized	Uniform	Same

Different setups

Setup	File popularity	Across Users
Decentralized	Uniform	Same
Non-uniform file popularity	Non-uniform	Same

Different setups

Setup	File popularity	Across Users
Decentralized	Uniform	Same
Non-uniform file popularity	Non-uniform	Same
Global vs Local popularity	Non-uniform	Different

Location wise popularity

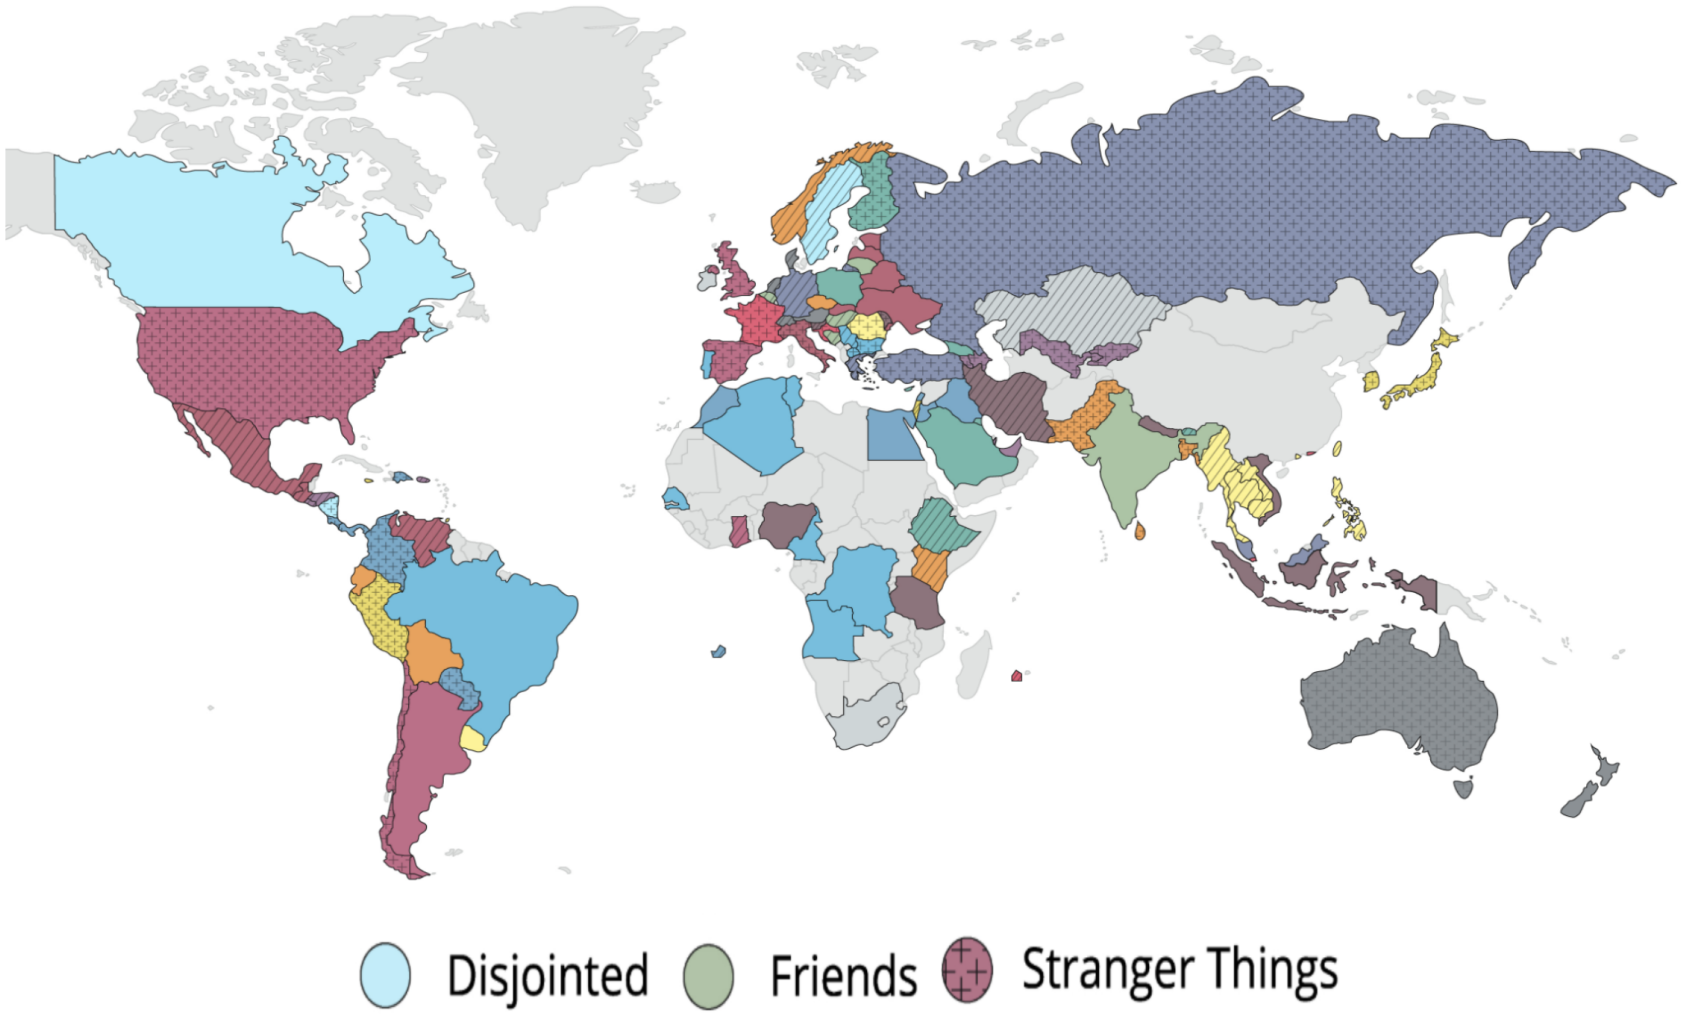

File demand model

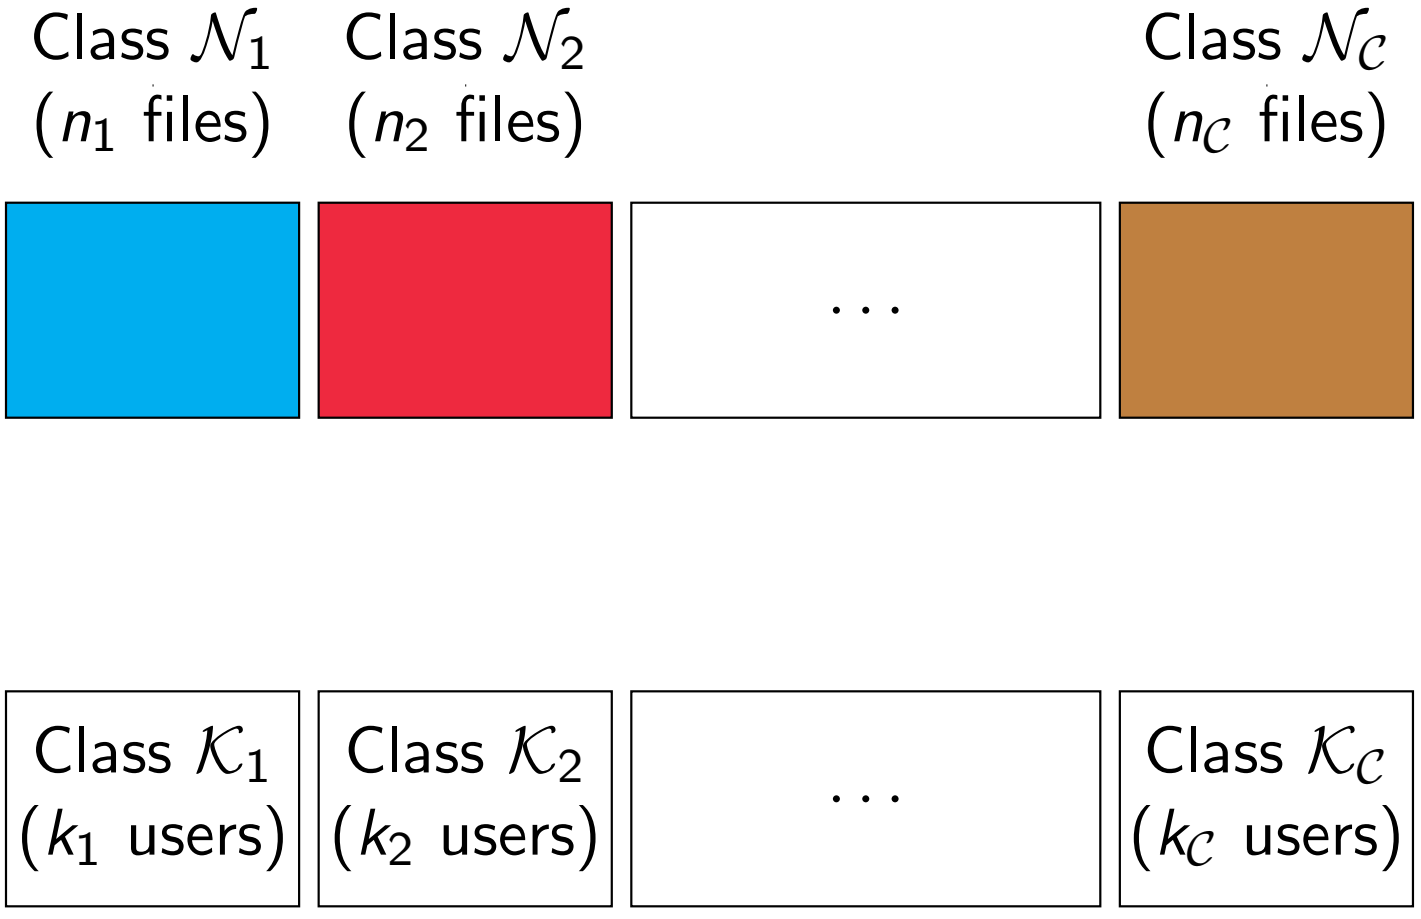

File demand model

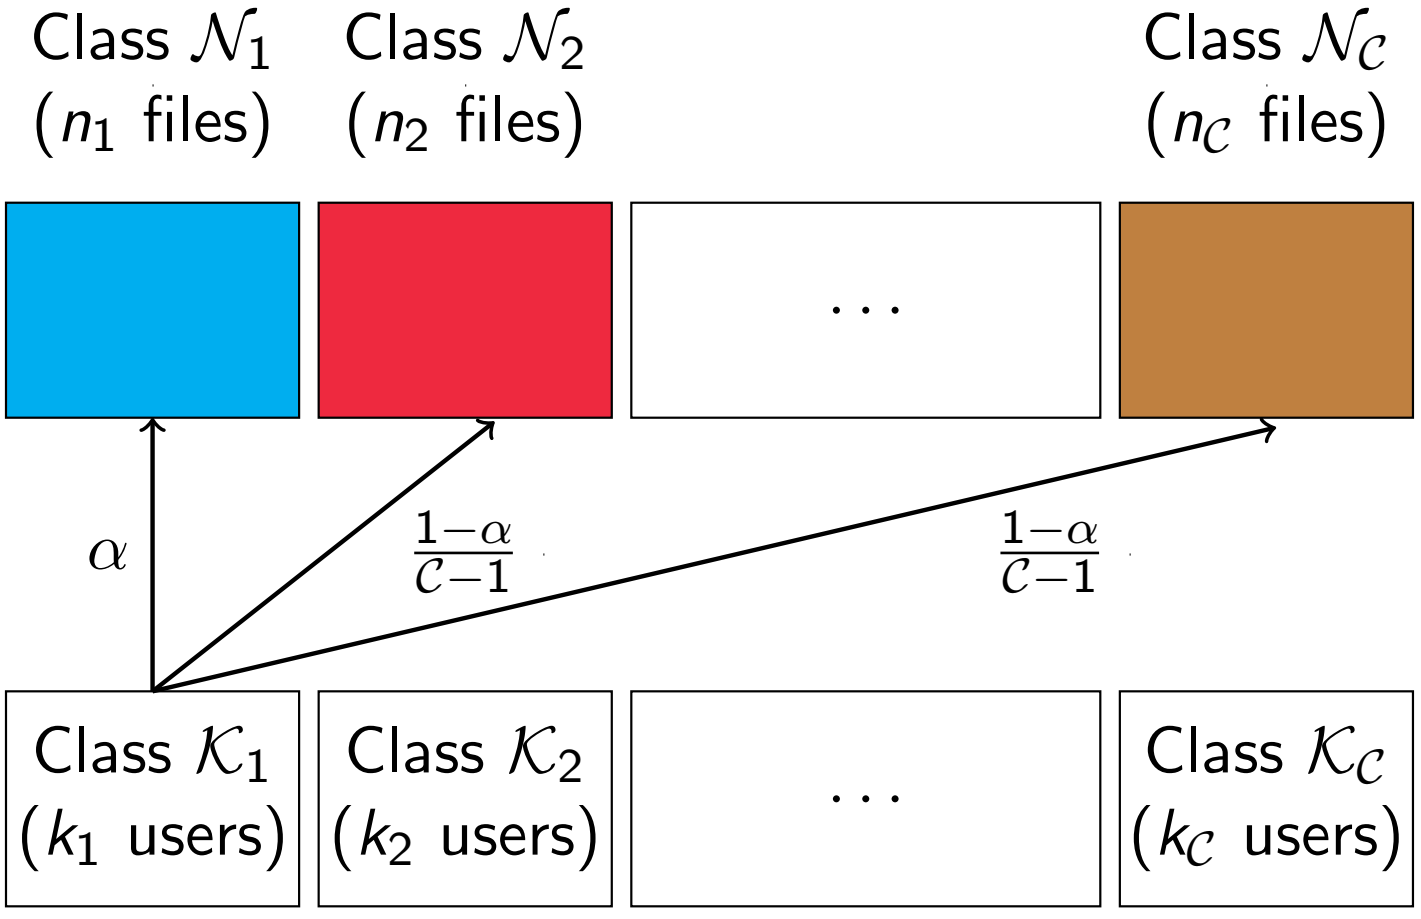

File demand model

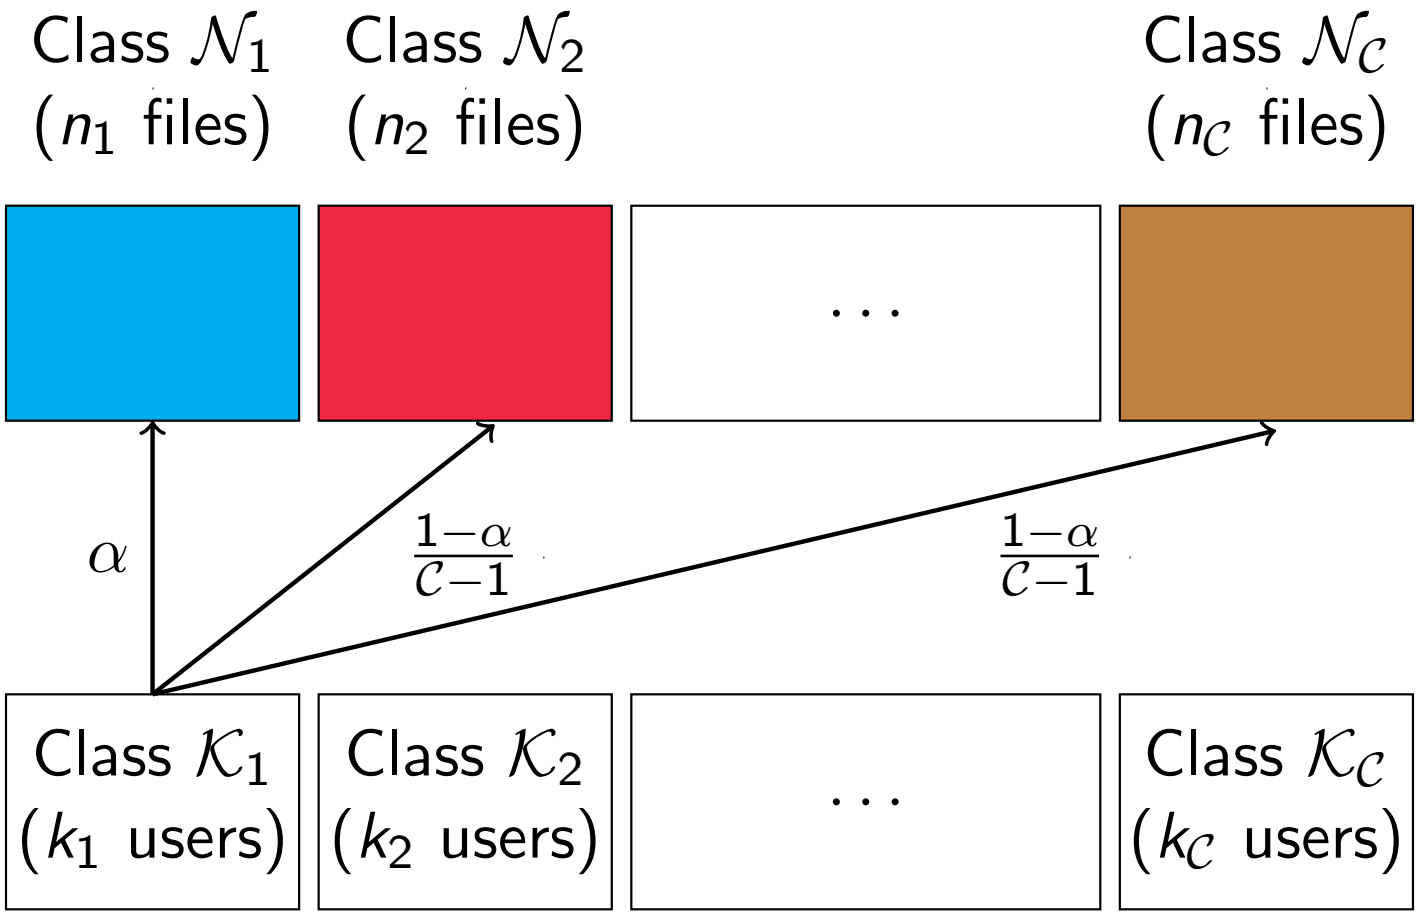

$n_1 = n_2 = \cdots n_C$ and $k_1 = k_2 = \cdots k_C$

Take away

There is a threshold α^* such that

Take away

There is a threshold α^* such that

- $\forall \alpha > \alpha^* \longrightarrow$ 'Local Preference' scheme outperforms 'Global Preference' scheme in some memory regime.

Take away

There is a threshold α^* such that

- $\forall \alpha > \alpha^* \longrightarrow$ 'Local Preference' scheme outperforms 'Global Preference' scheme in some memory regime.
- $\forall \alpha \leq \alpha^* \longrightarrow$ 'Global Preference' scheme outperforms 'Local Preference' scheme in (almost) all memory regimes.

Global popularity

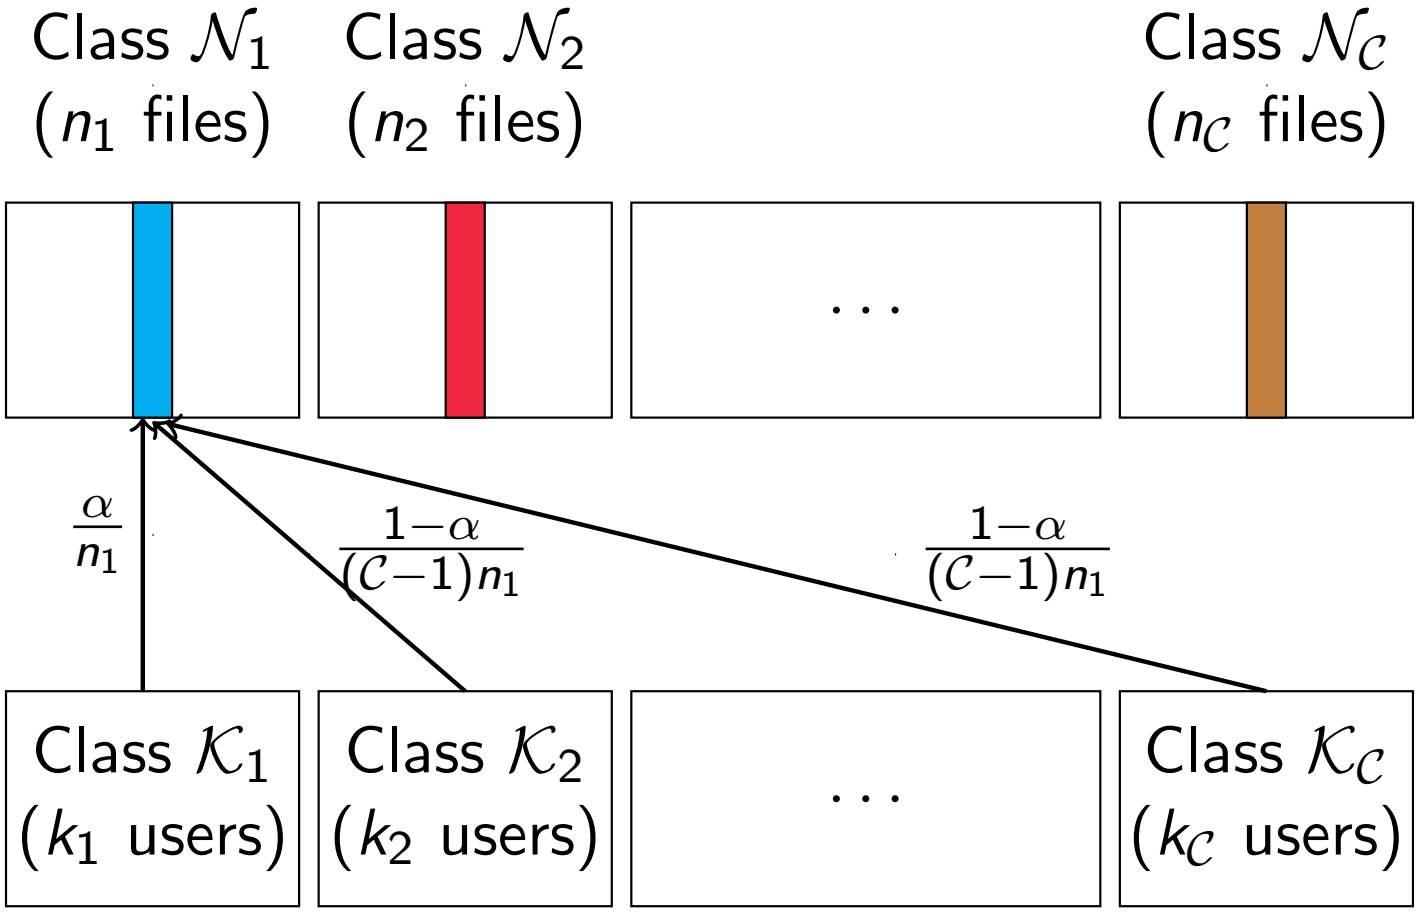

Global popularity

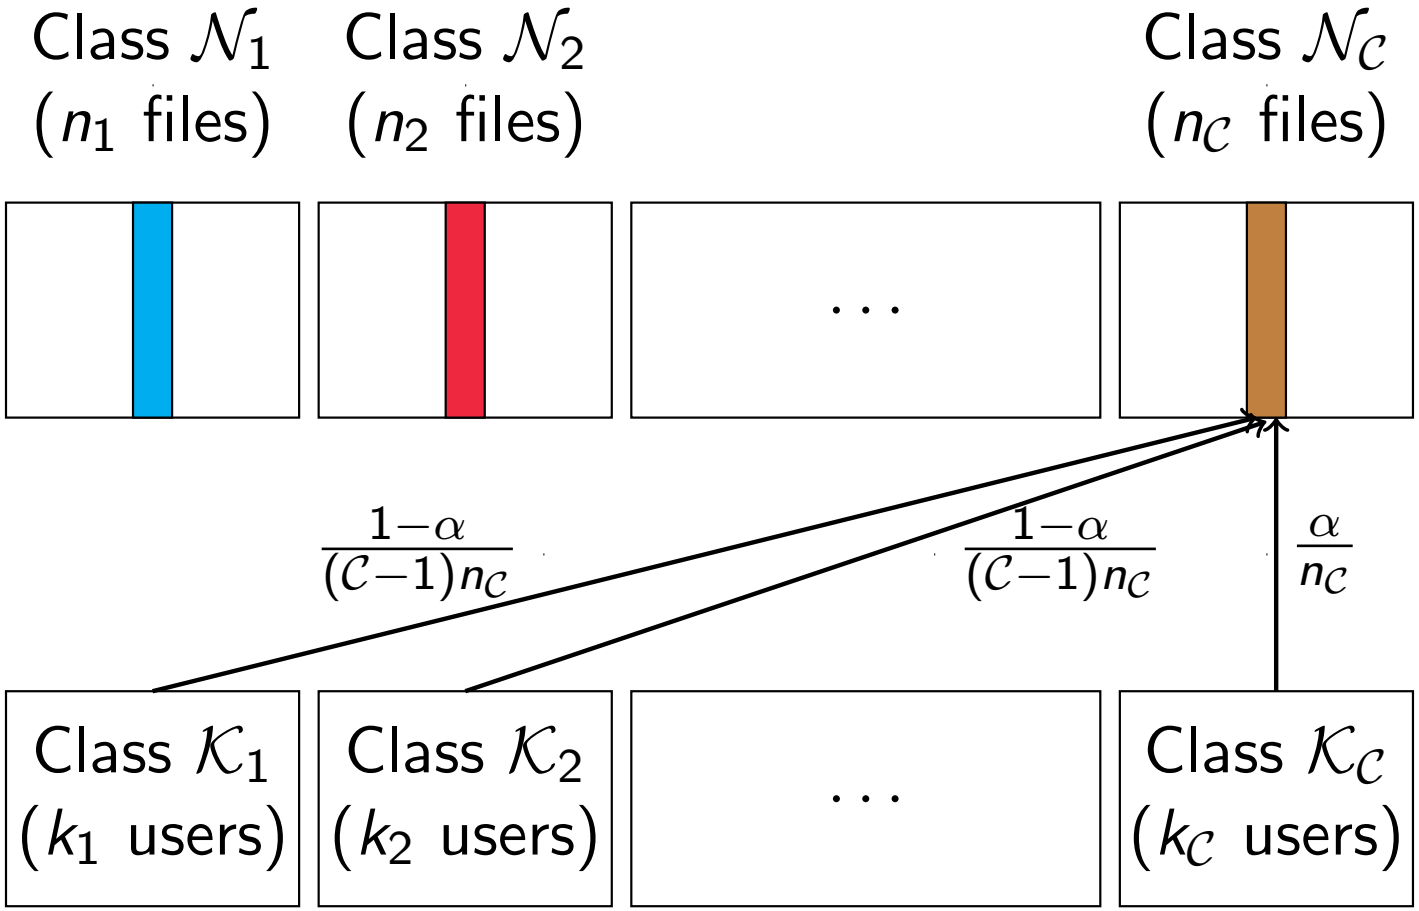

Global preference - placement

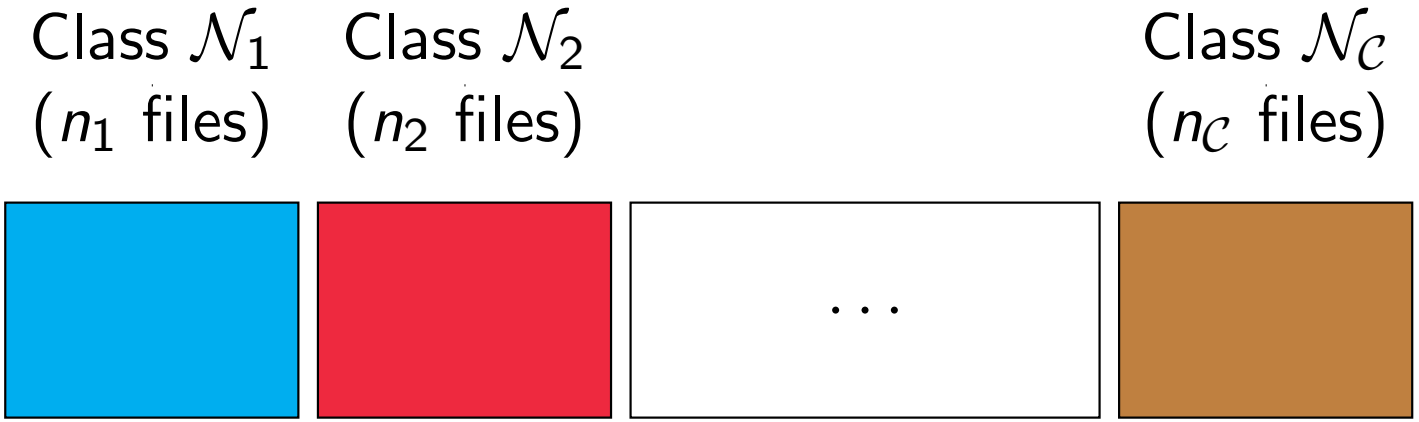

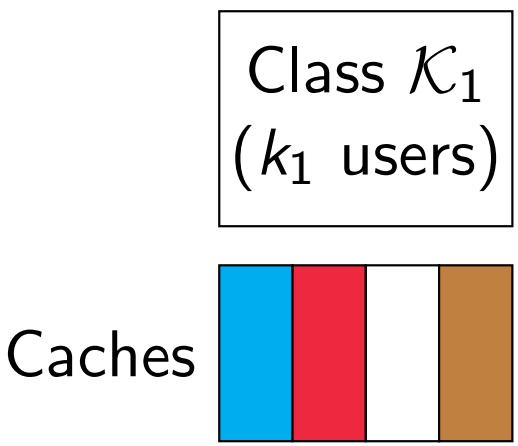

Global preference - placement

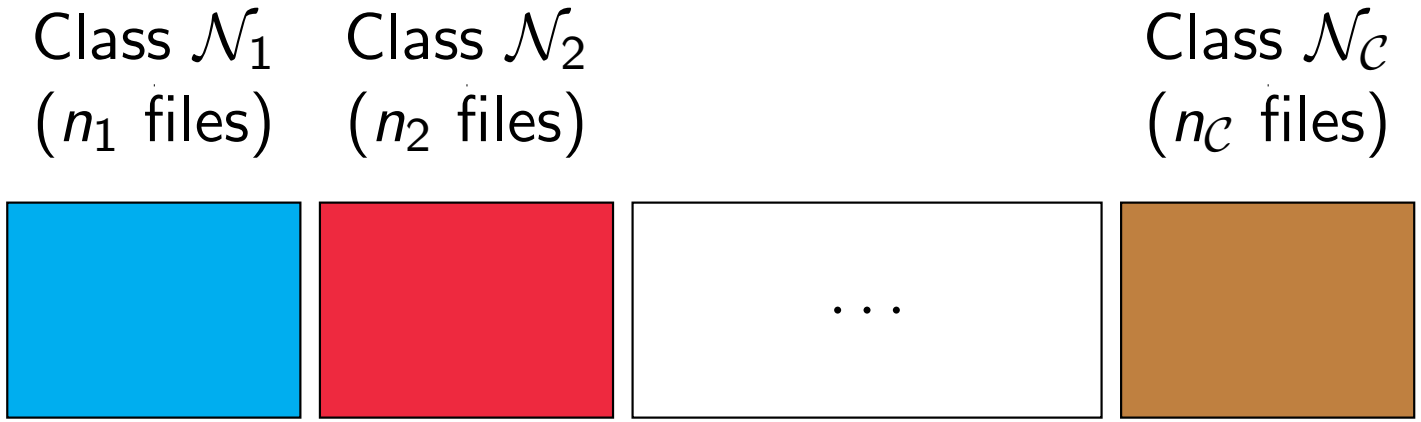

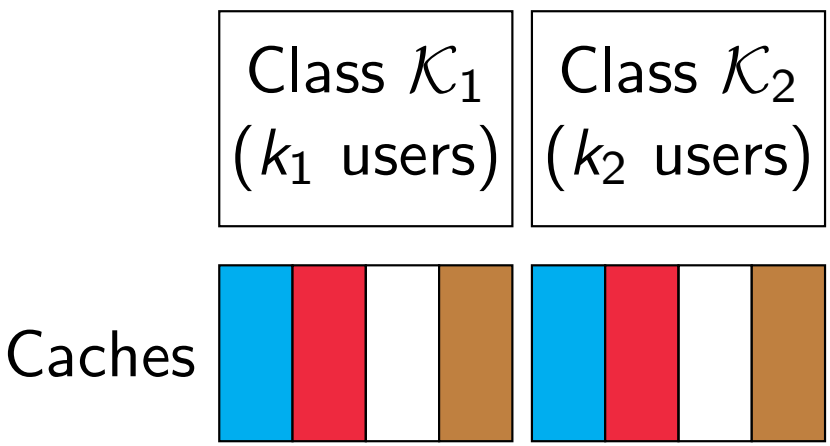

Global preference - placement

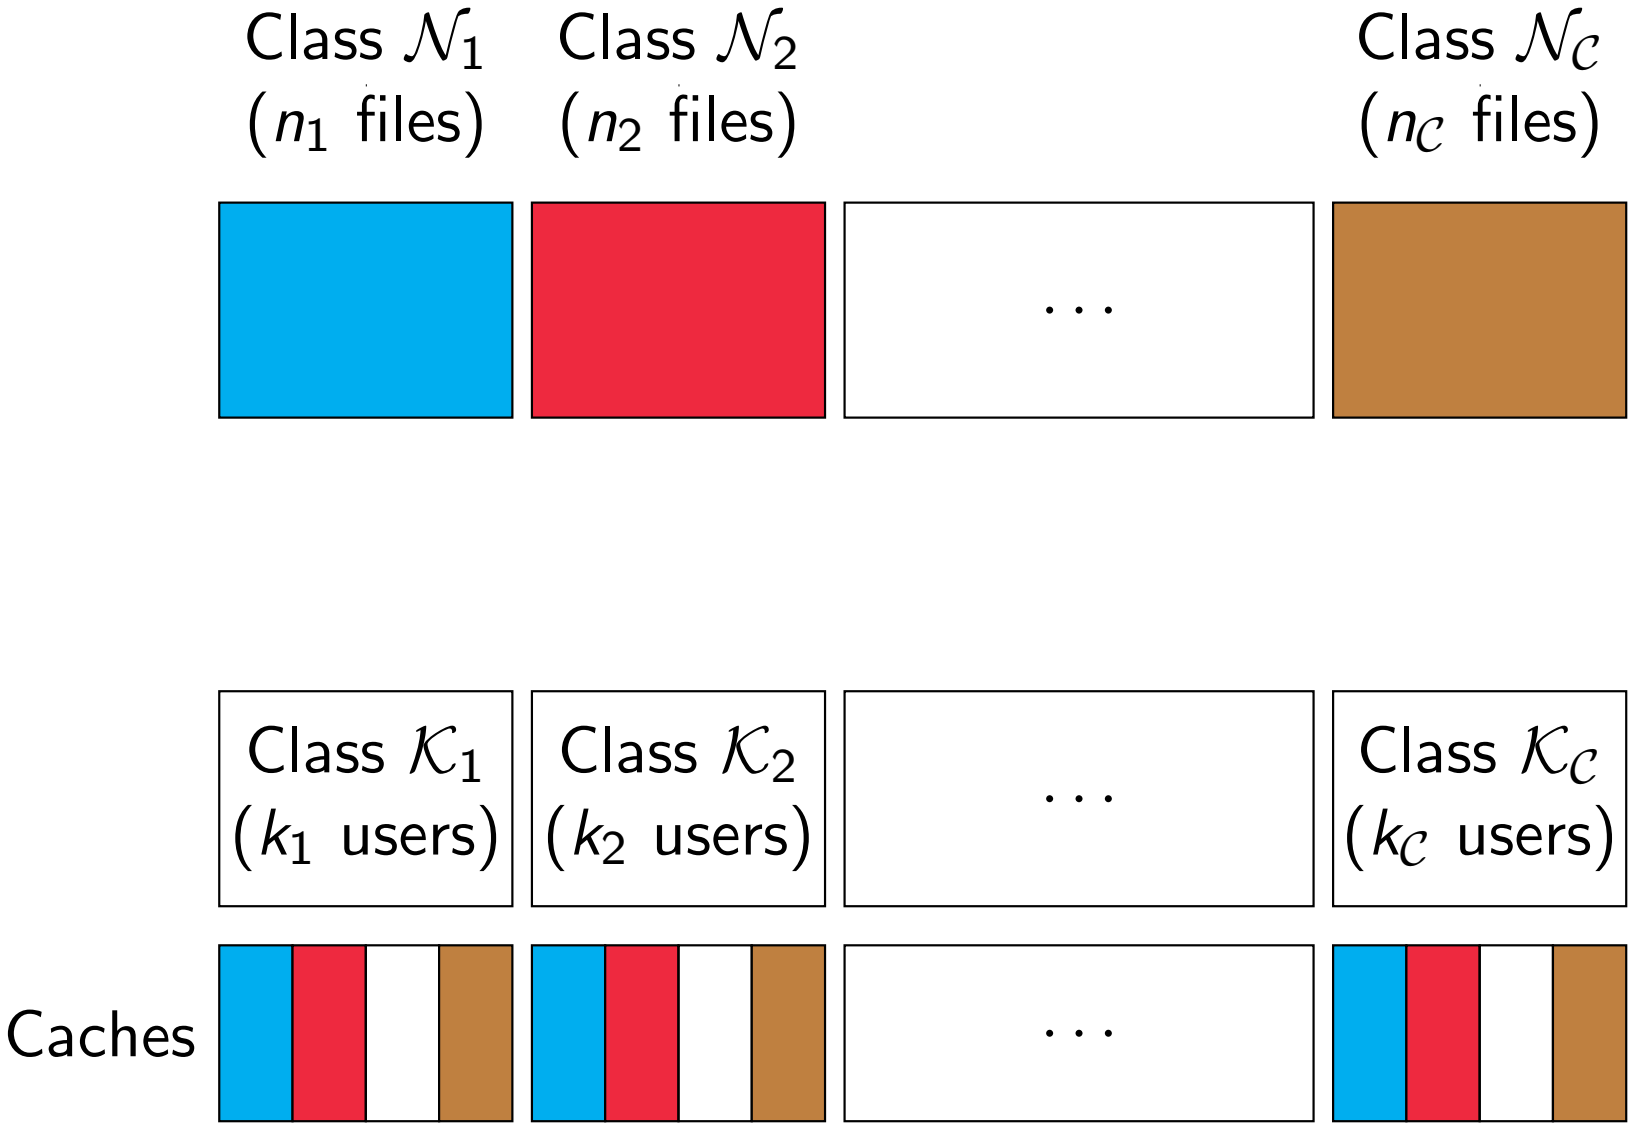

Global preference + File perspective delivery

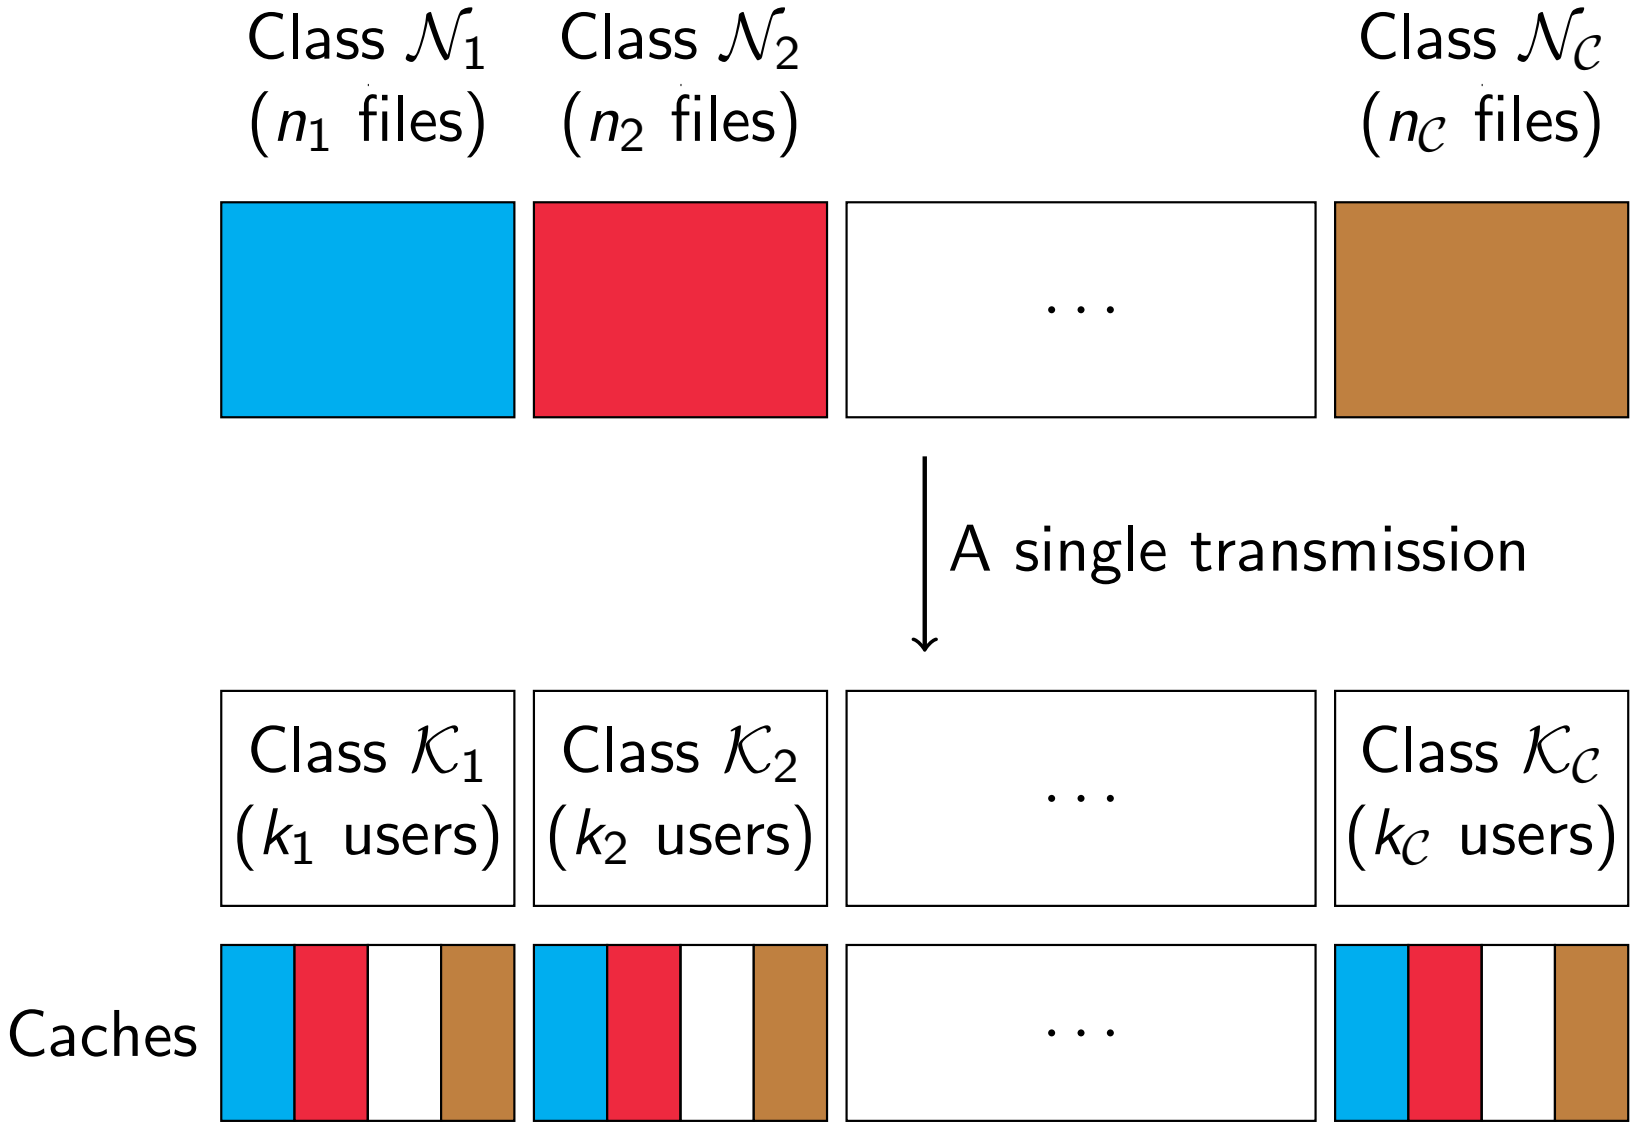

Global preference + File perspective delivery

$$R_T = \left[\frac{N}{M} - 1\right]\left(1 - \left[1 - \frac{M}{N}\right]^K\right)$$

Local preference - placement

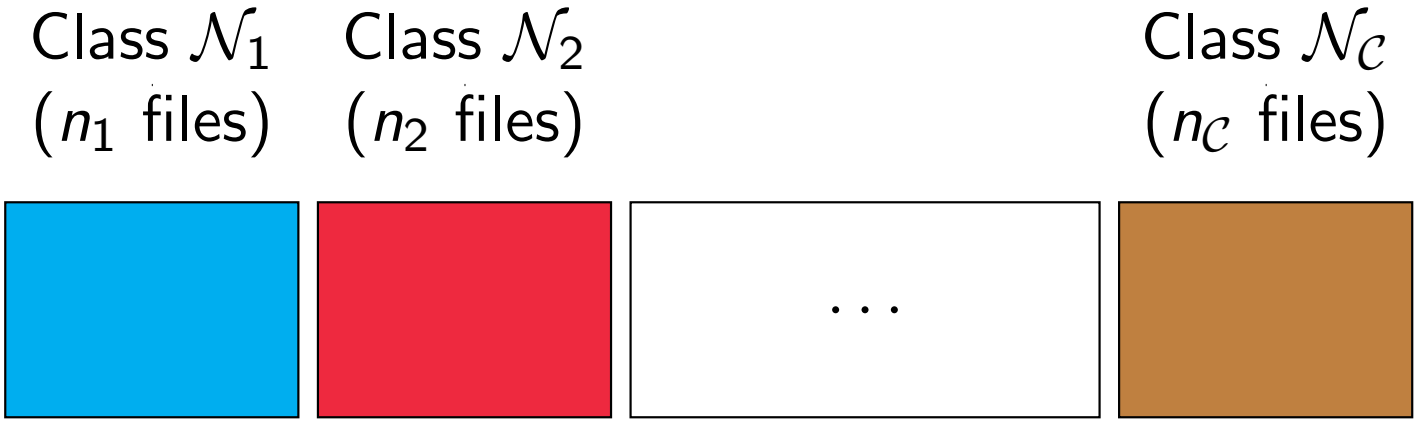

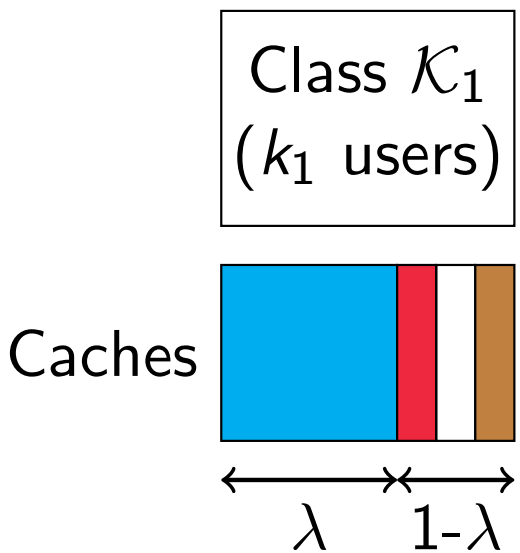

Local preference - placement

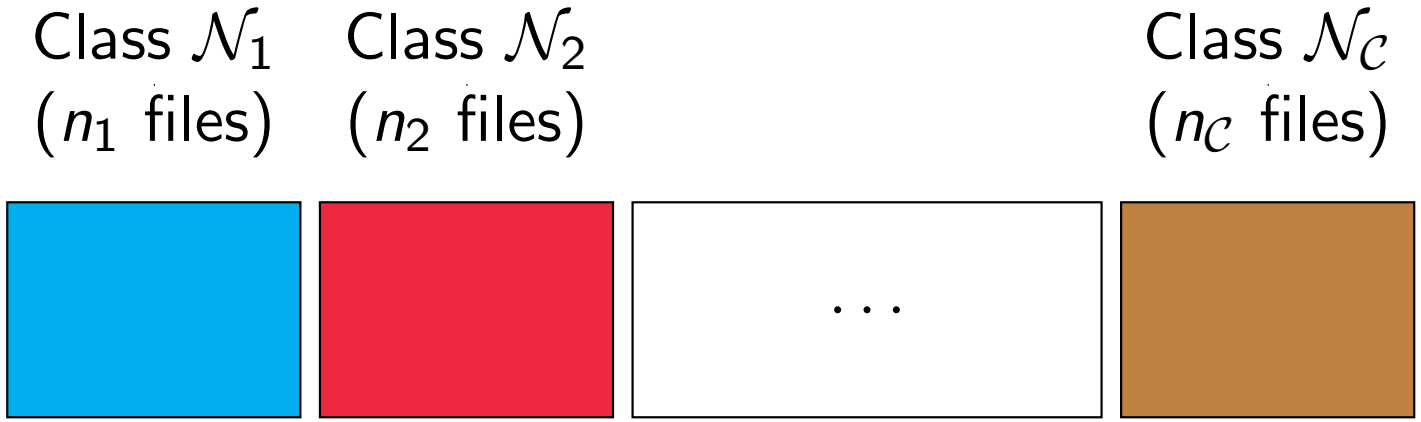

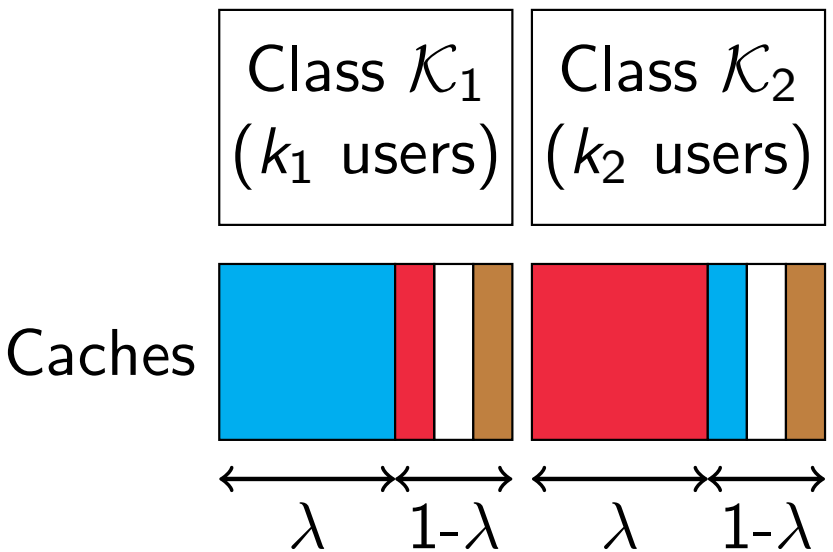

Local preference - placement

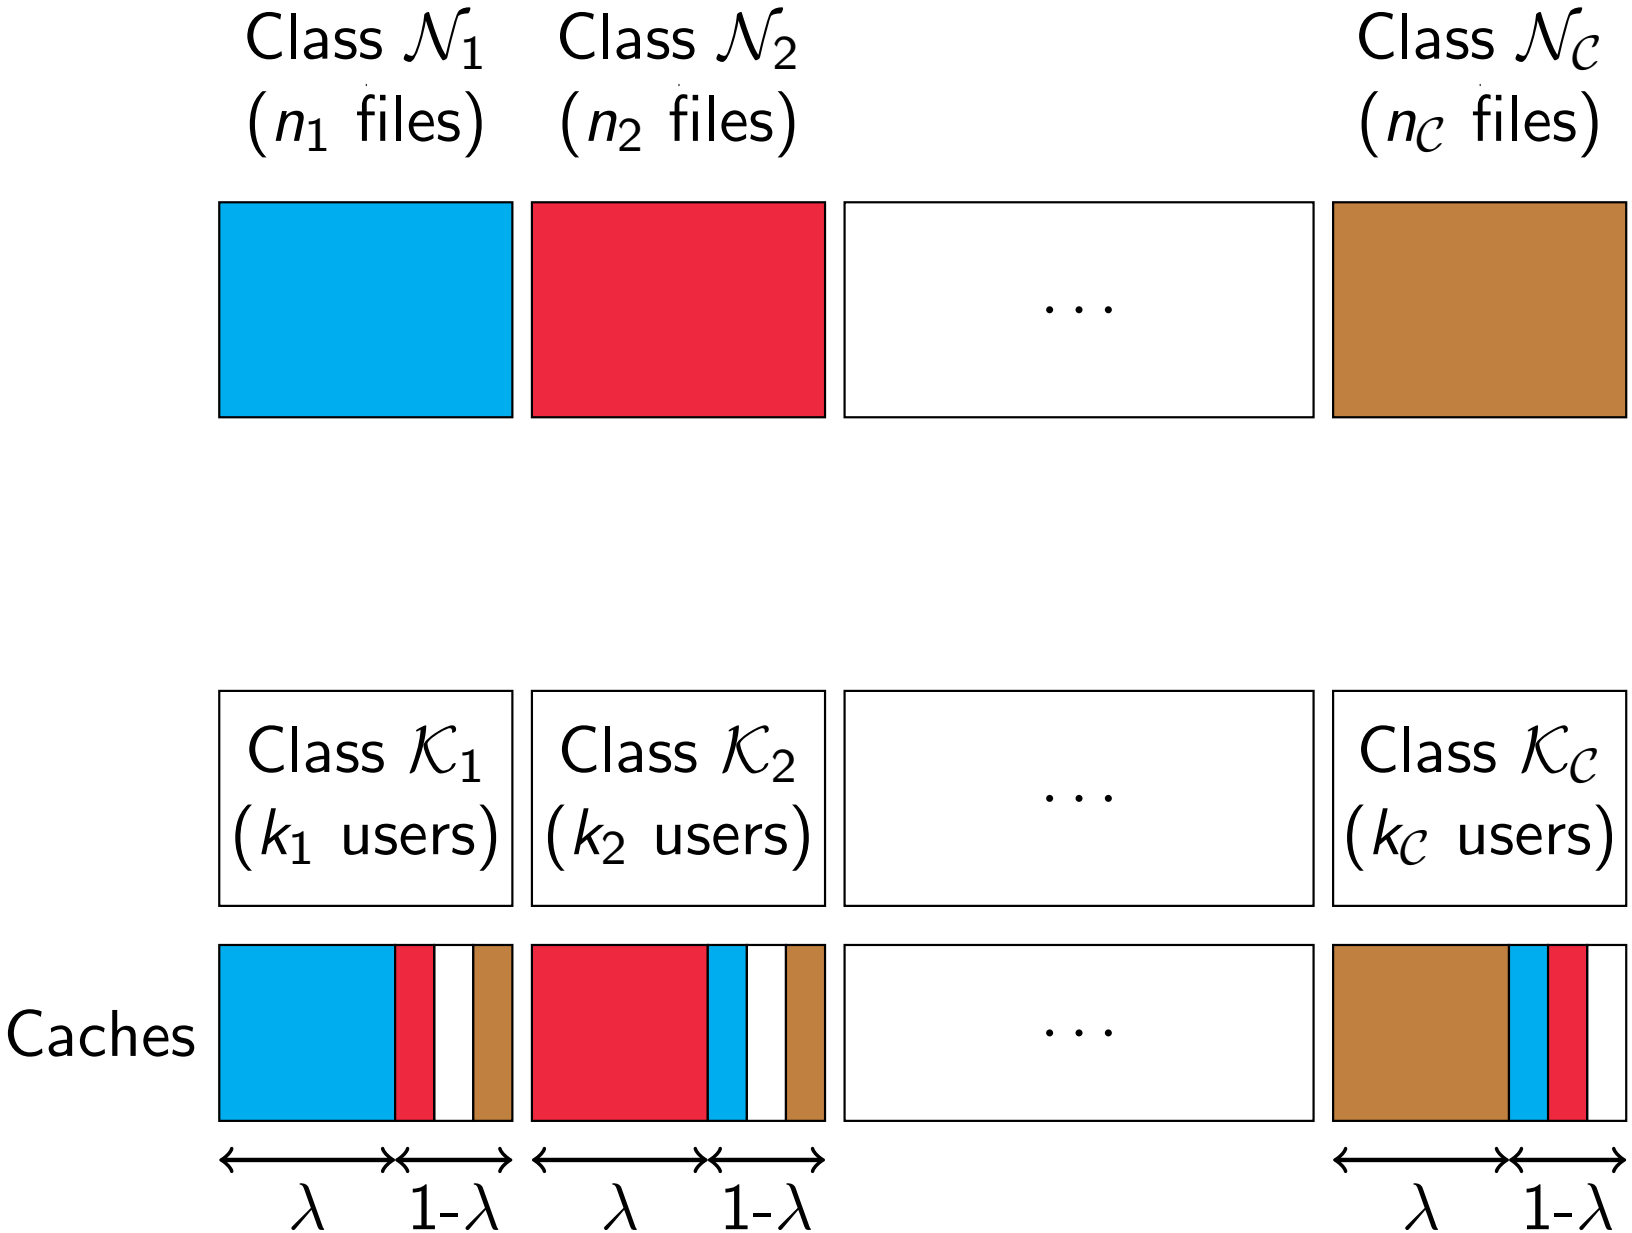

Local preference + User perspective delivery

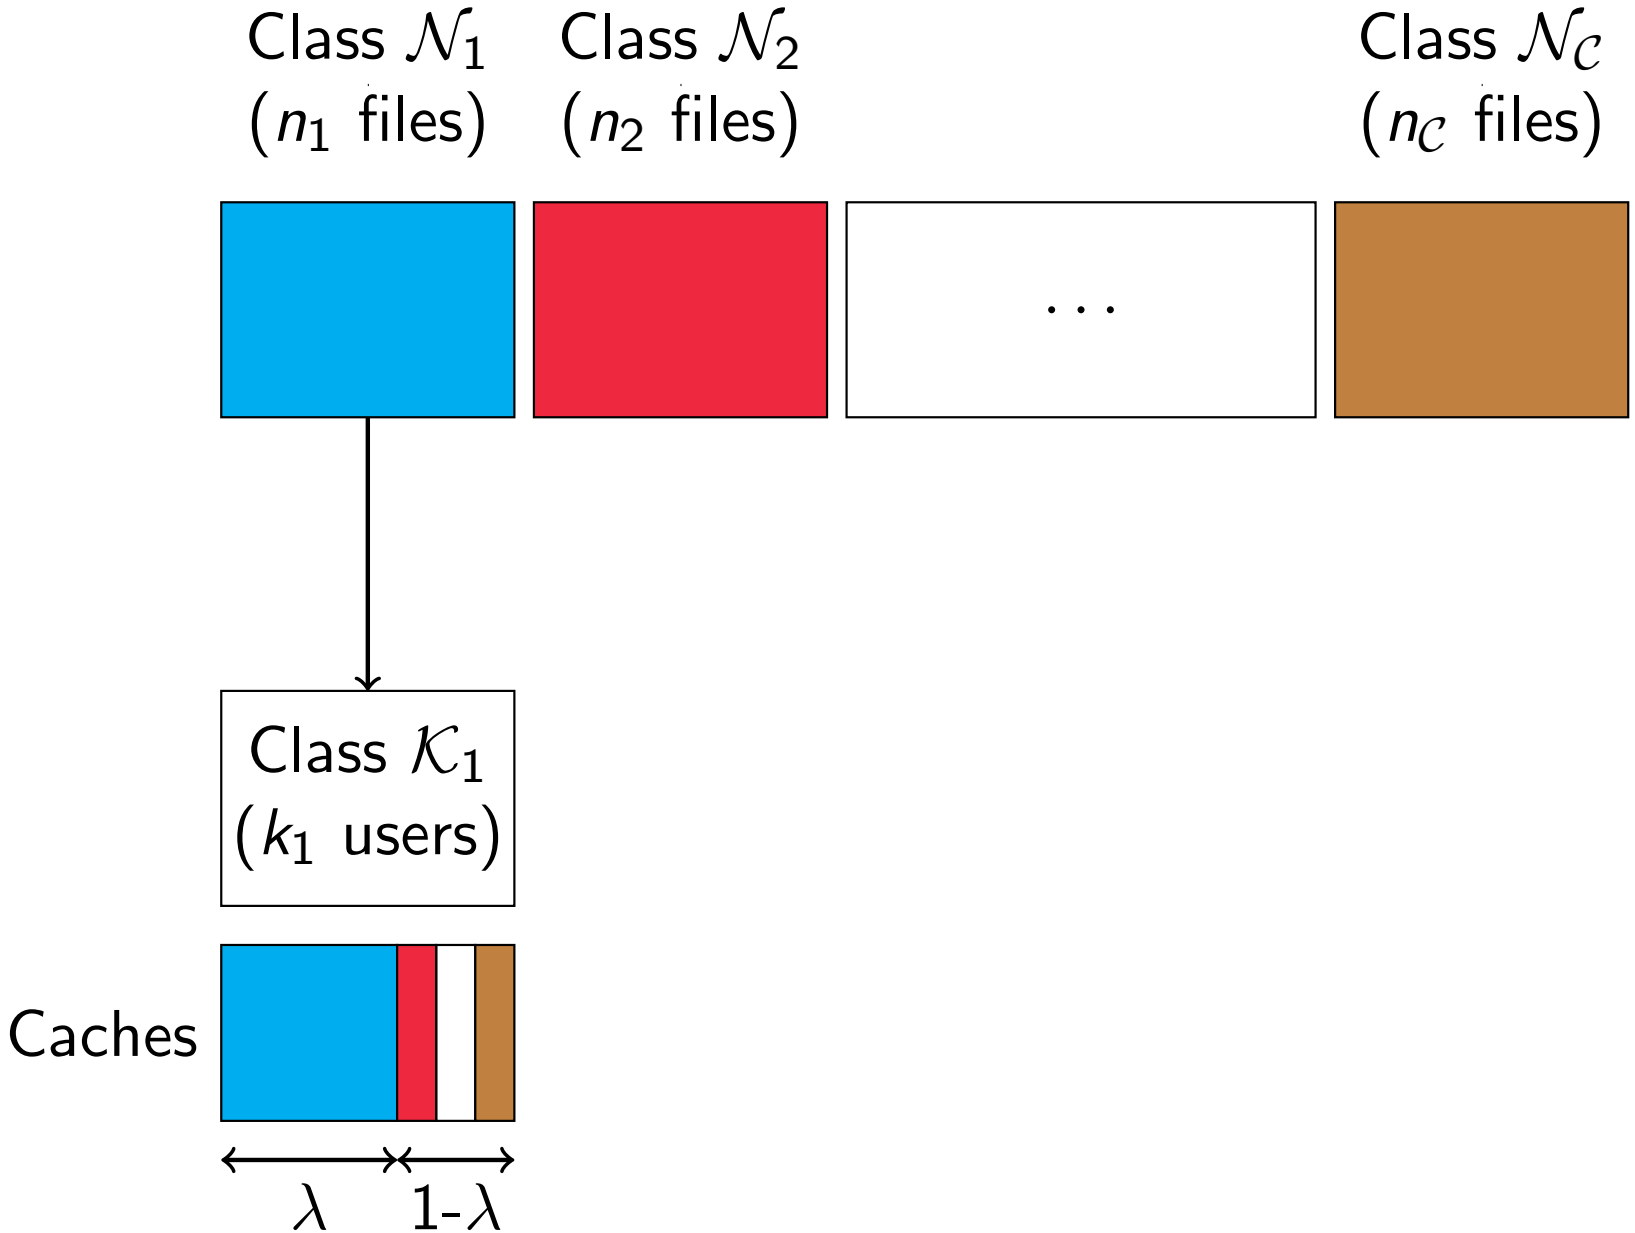

Local preference + User perspective delivery

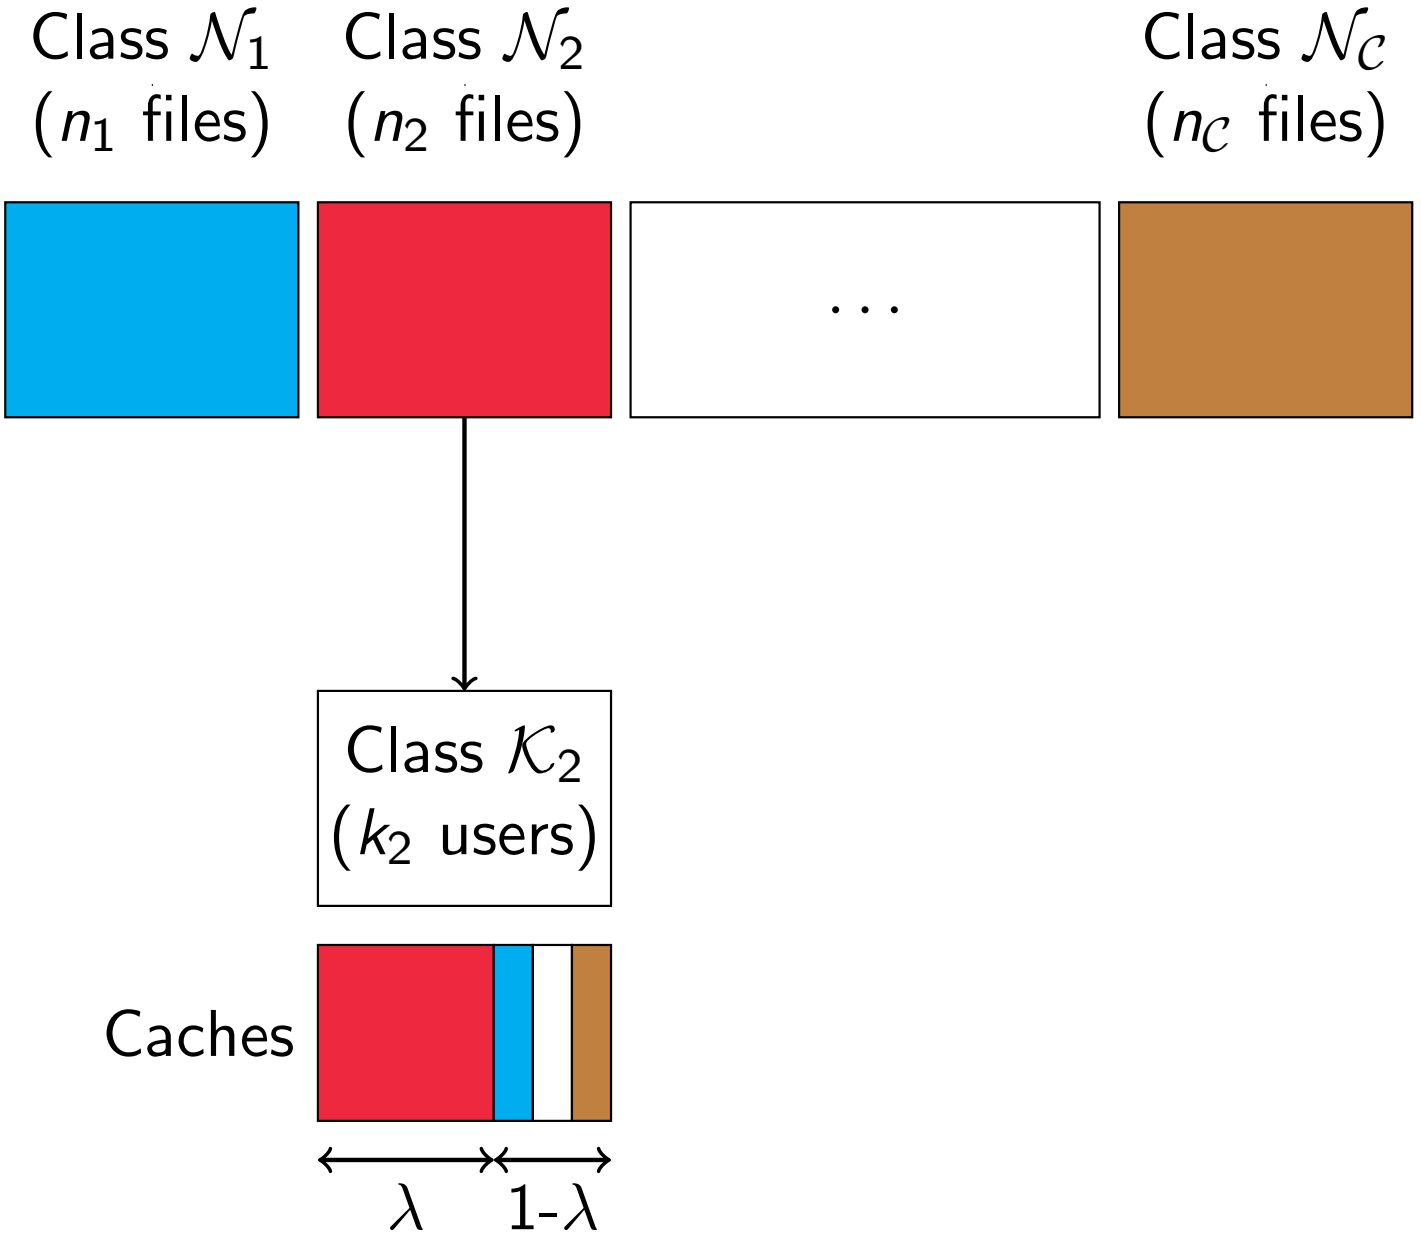

Local preference + User perspective delivery

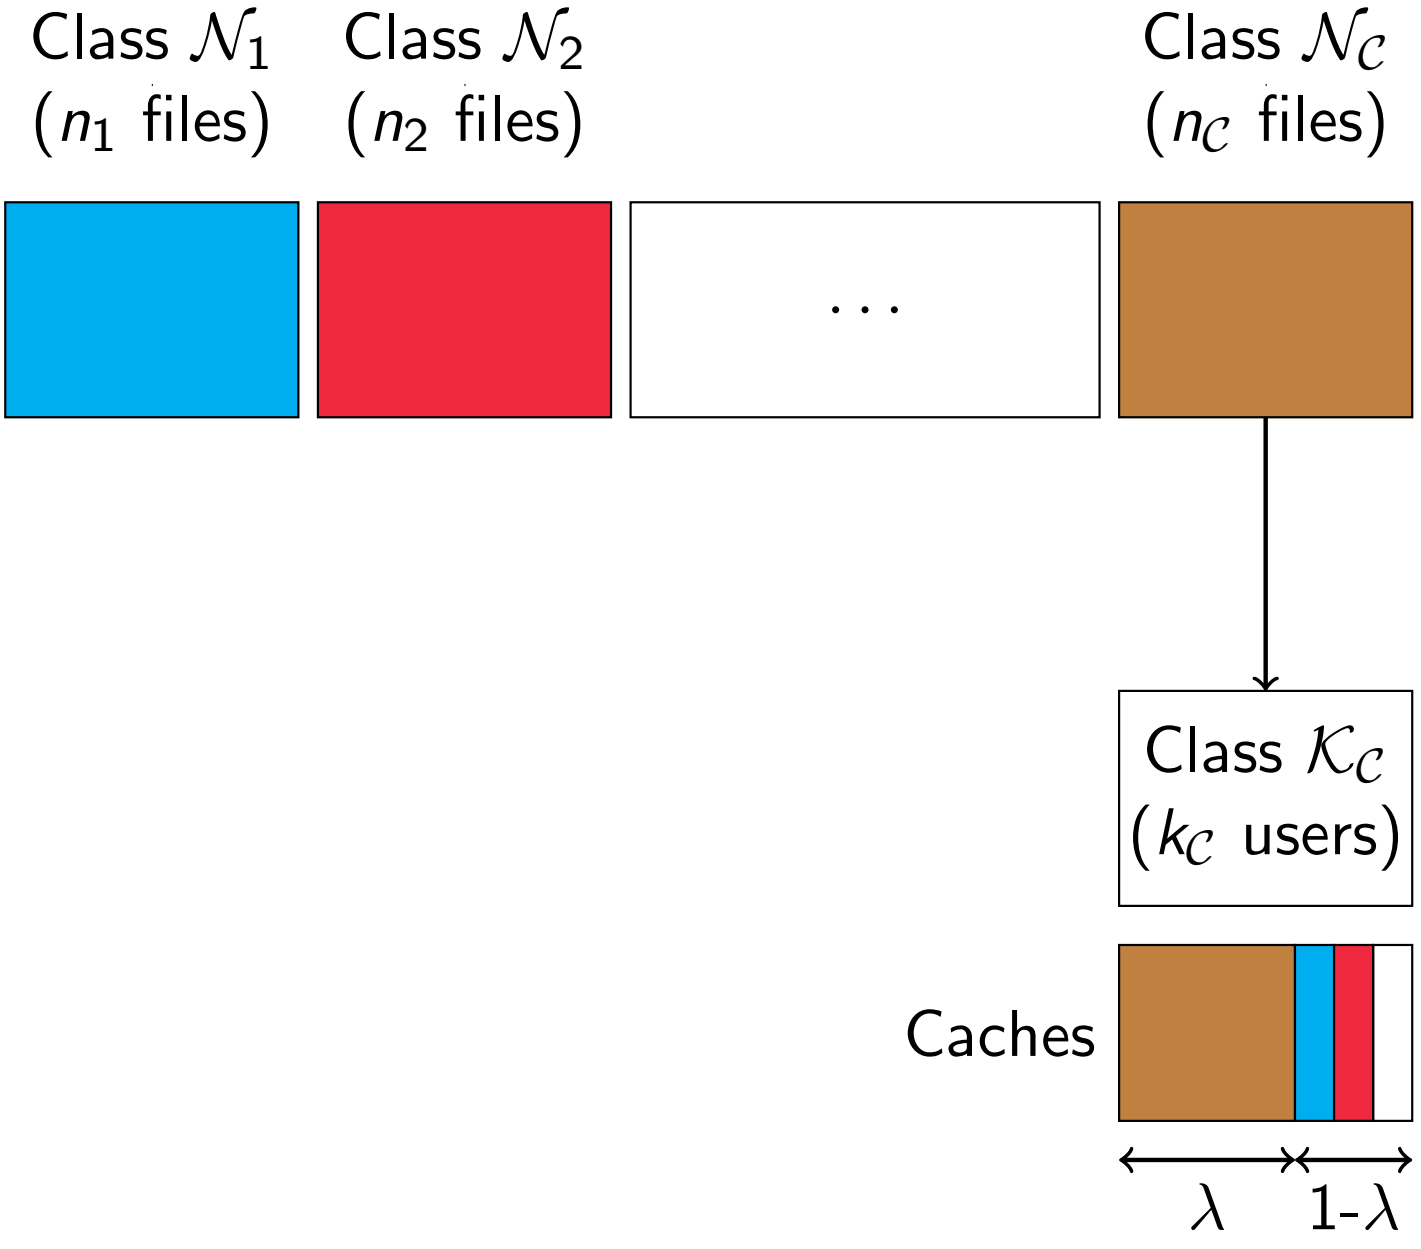

Local preference + User perspective delivery

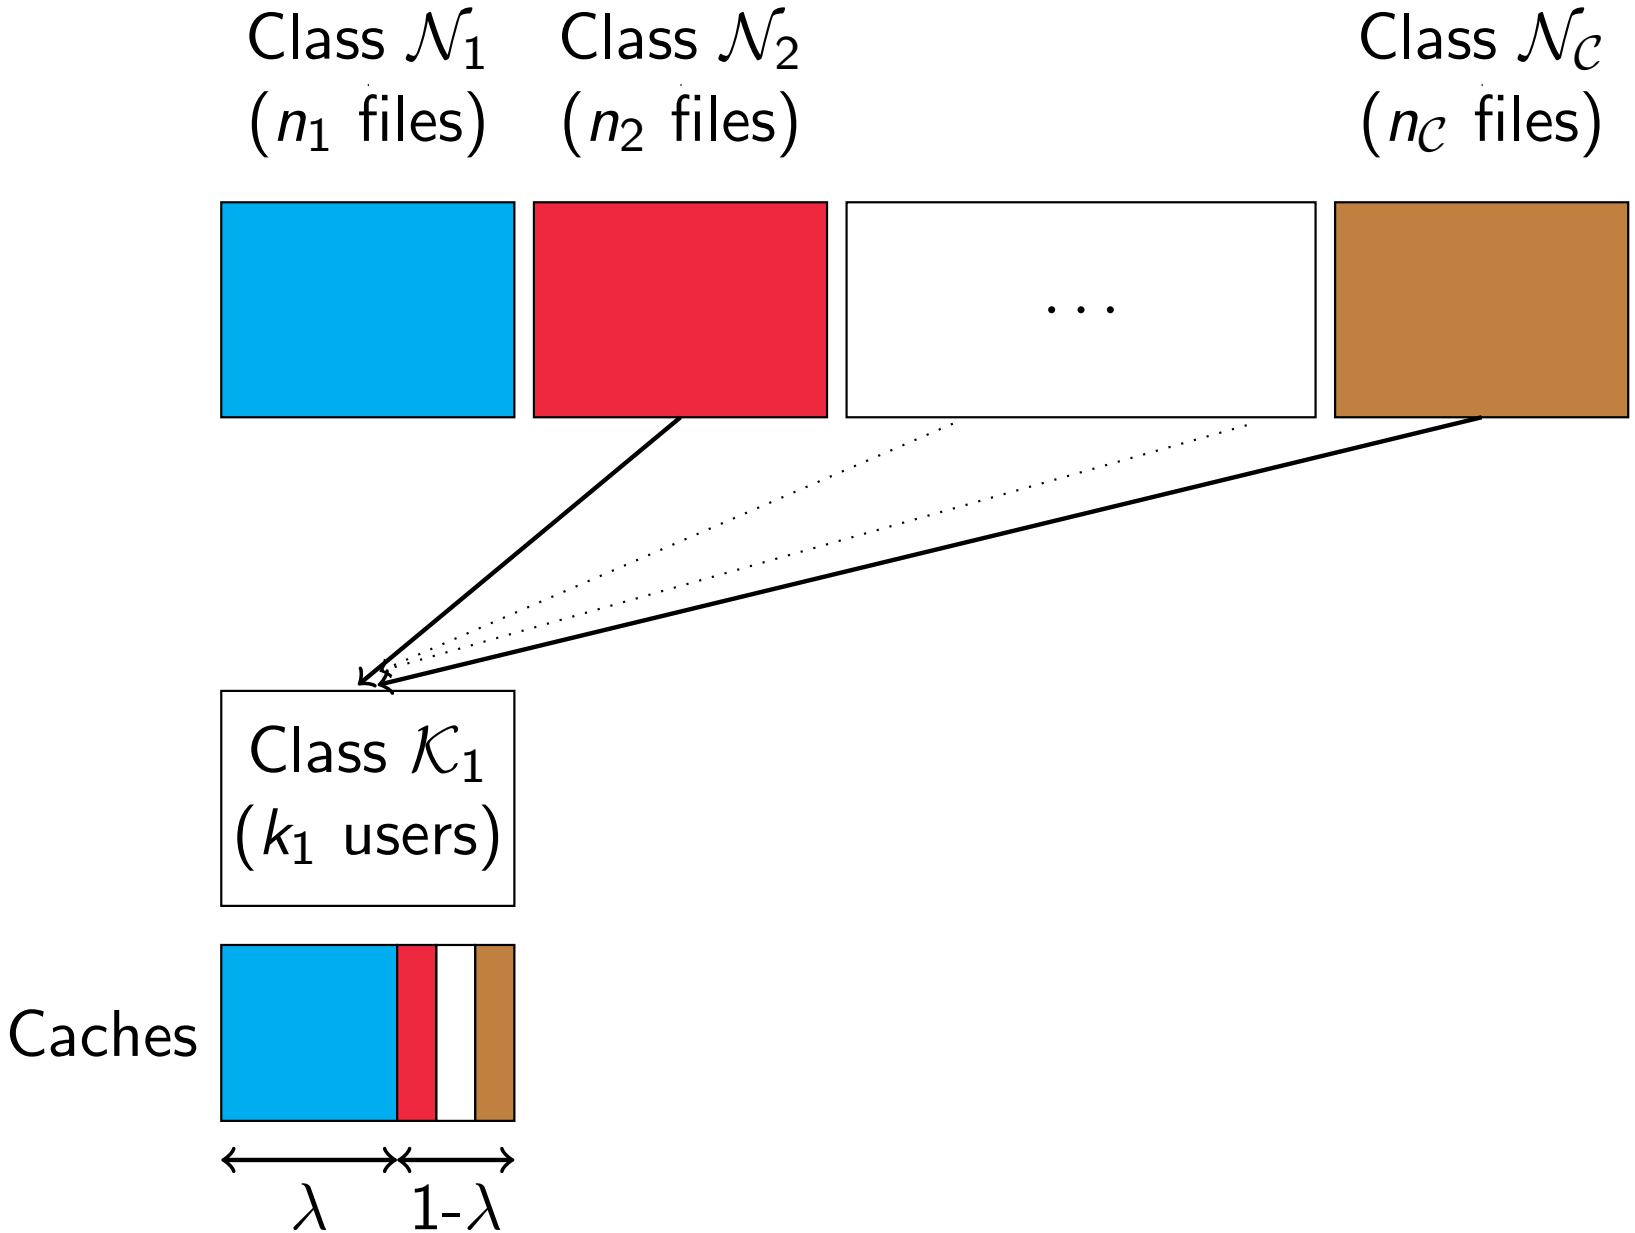

Local preference + User perspective delivery

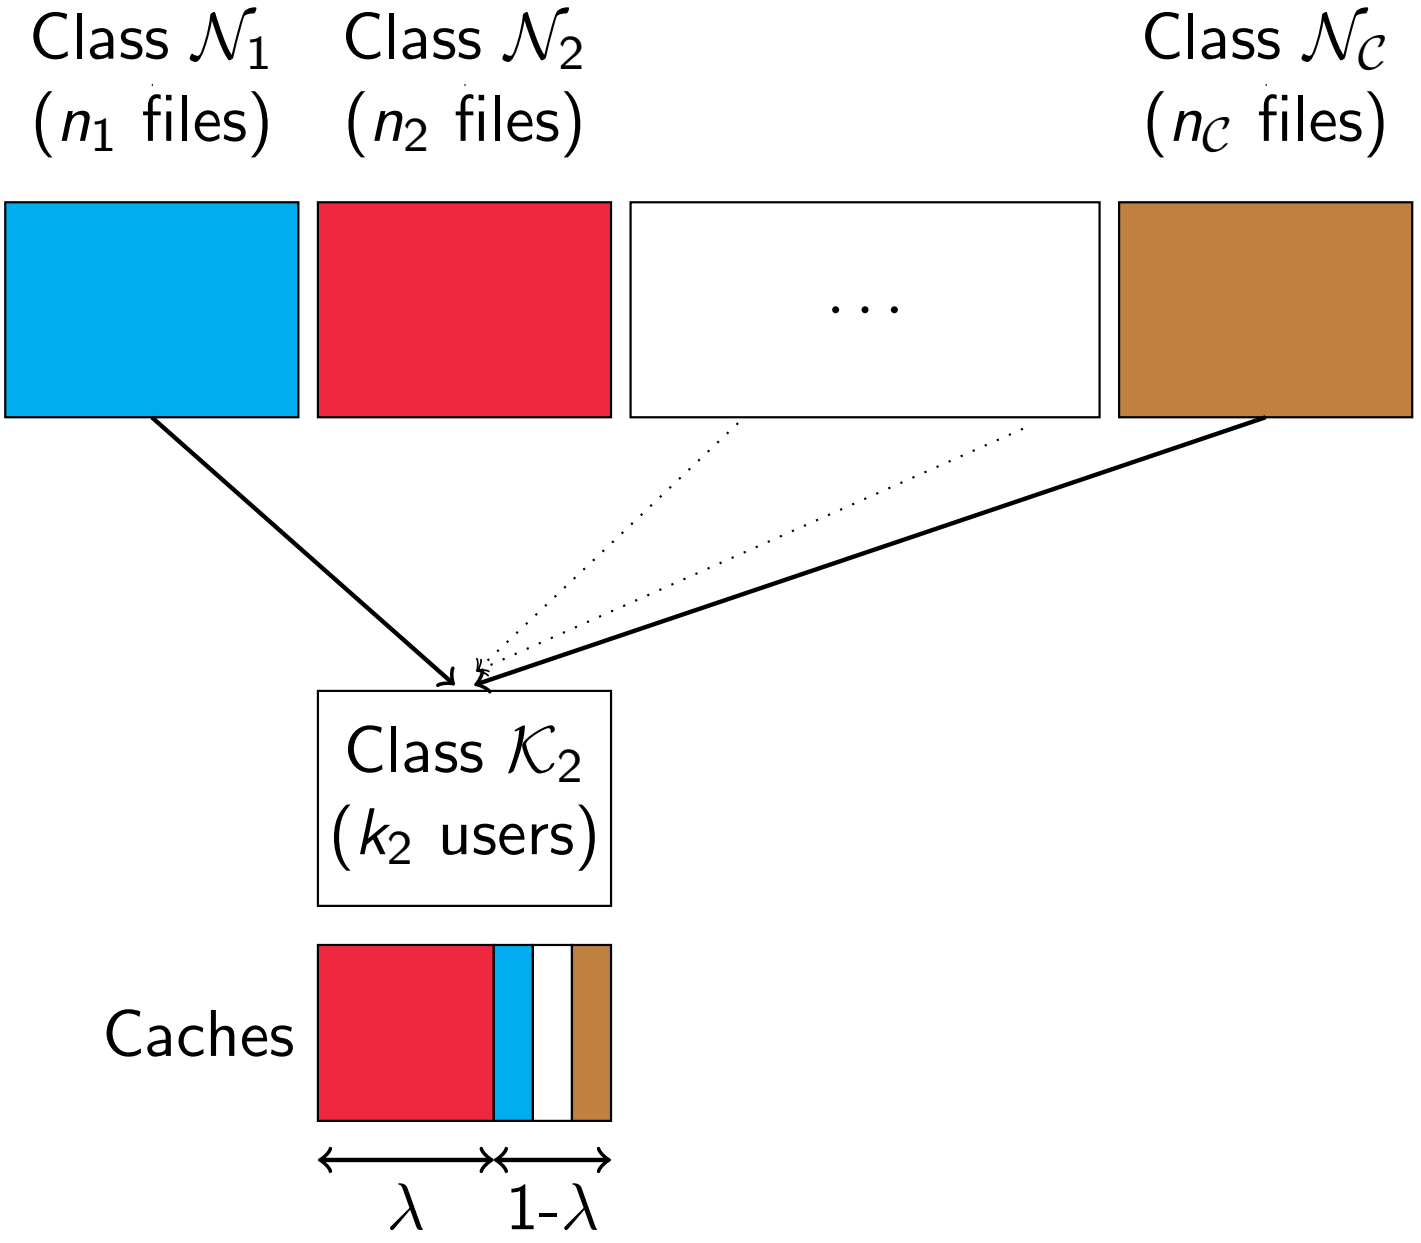

Local preference + User perspective delivery

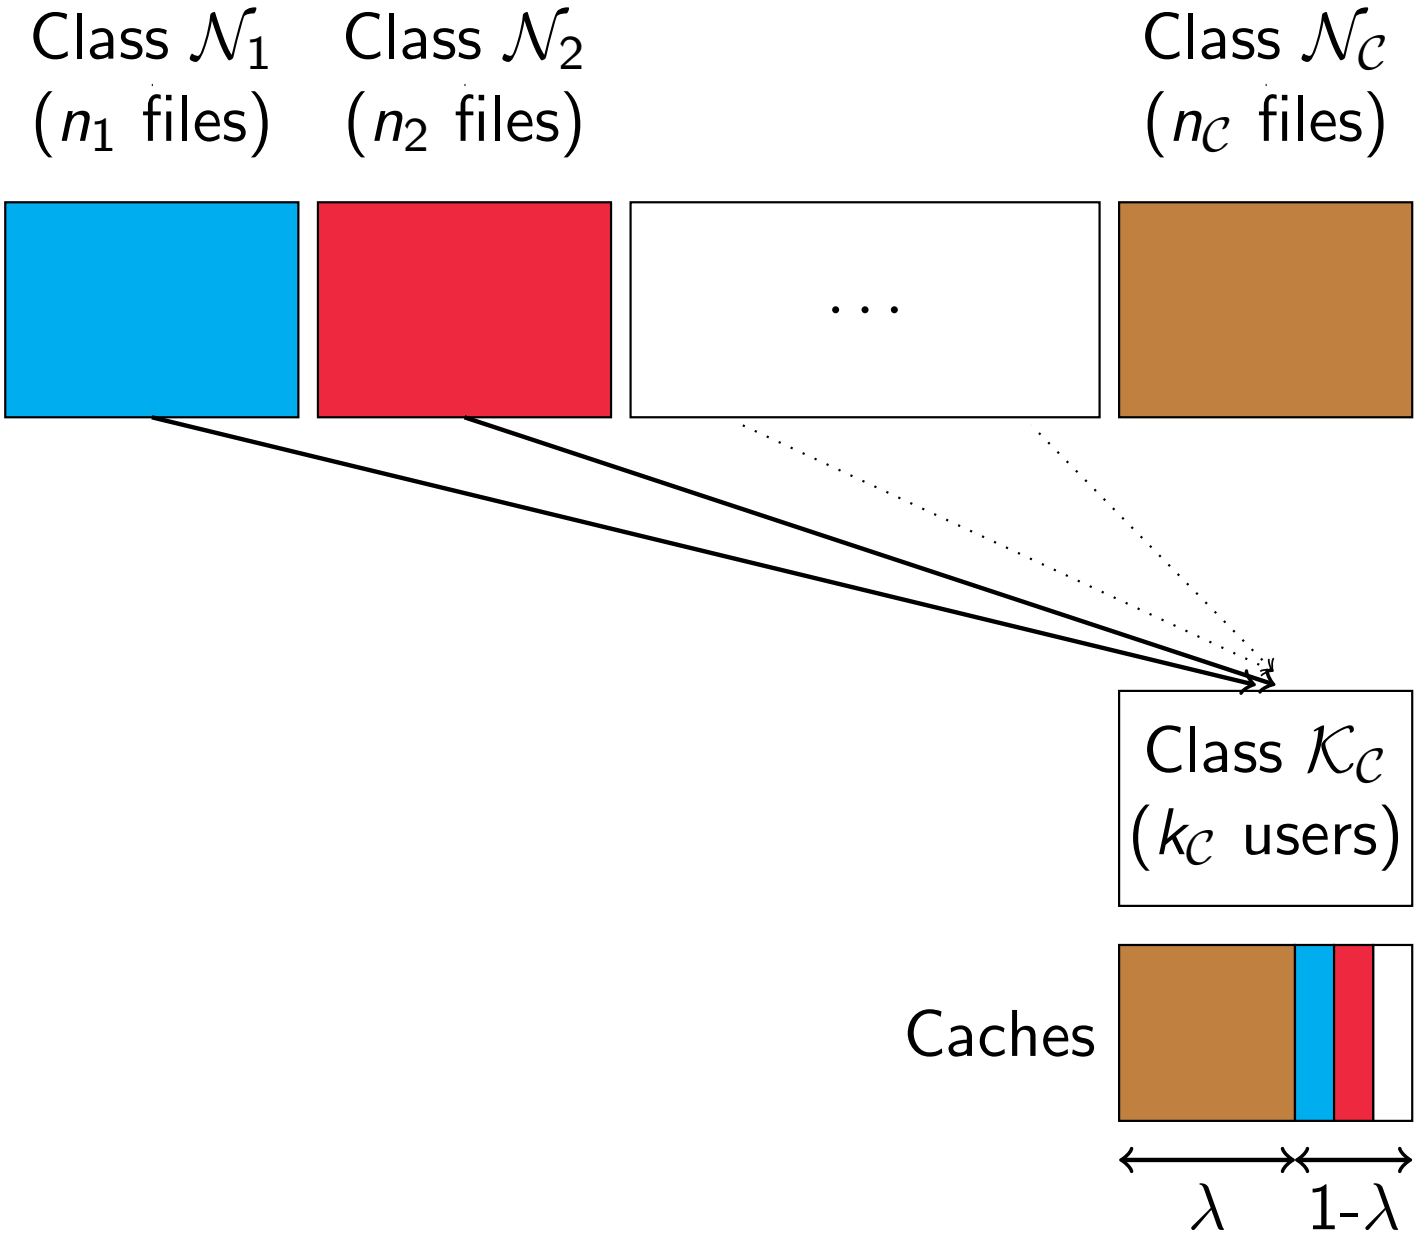

Local preference + User perspective delivery

$$R_U = \min_{\lambda \in [0,1]} \left(\left[1 - \frac{\lambda M \mathcal{C}}{N} \right]_+ \min \left\{ \frac{N}{\lambda M}, K\alpha \right\} + \left[1 - \frac{(1-\lambda) M \mathcal{C}}{N(\mathcal{C}-1)} \right]_+ \min \left\{ \frac{N(\mathcal{C}-1)}{(1-\lambda)M}, K(1-\alpha) \right\} \right)$$

where $[x]_+ = \max\{x, 0\}$

Local preference + File perspective delivery

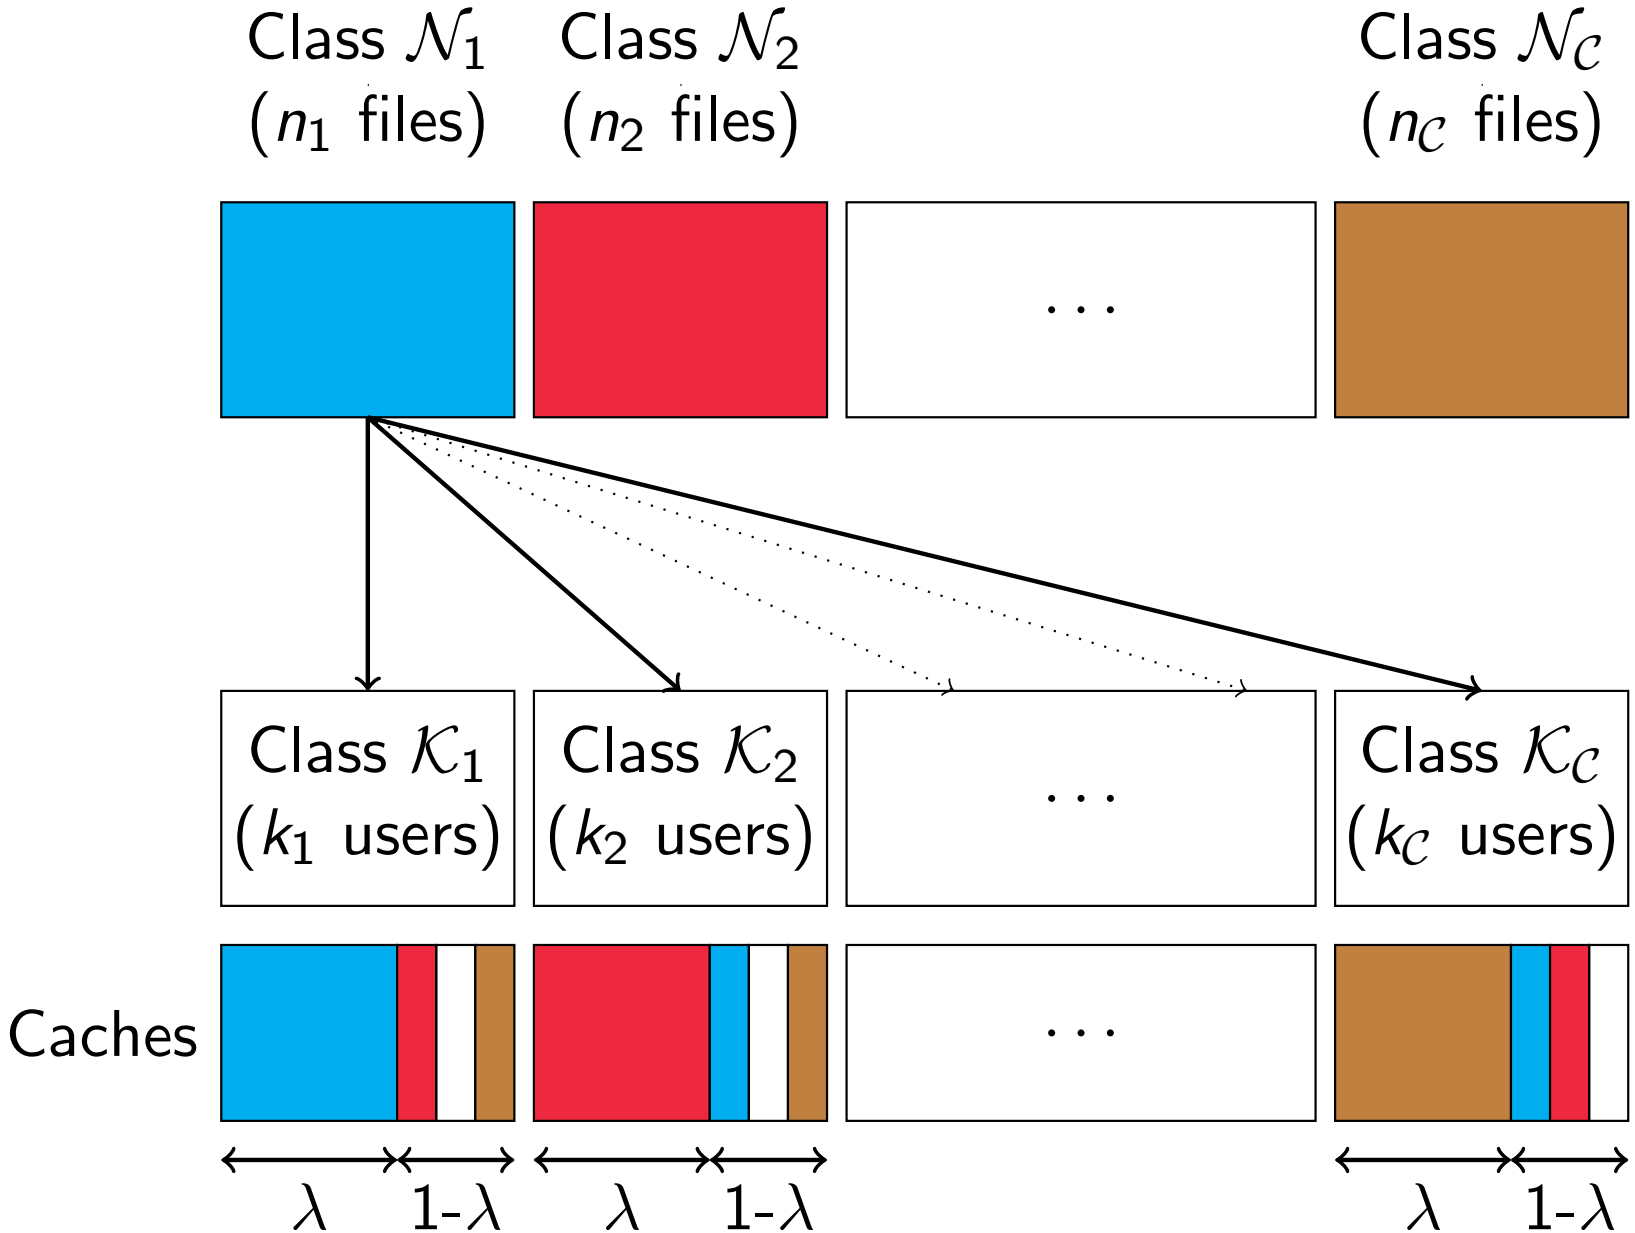

Local preference + File perspective delivery

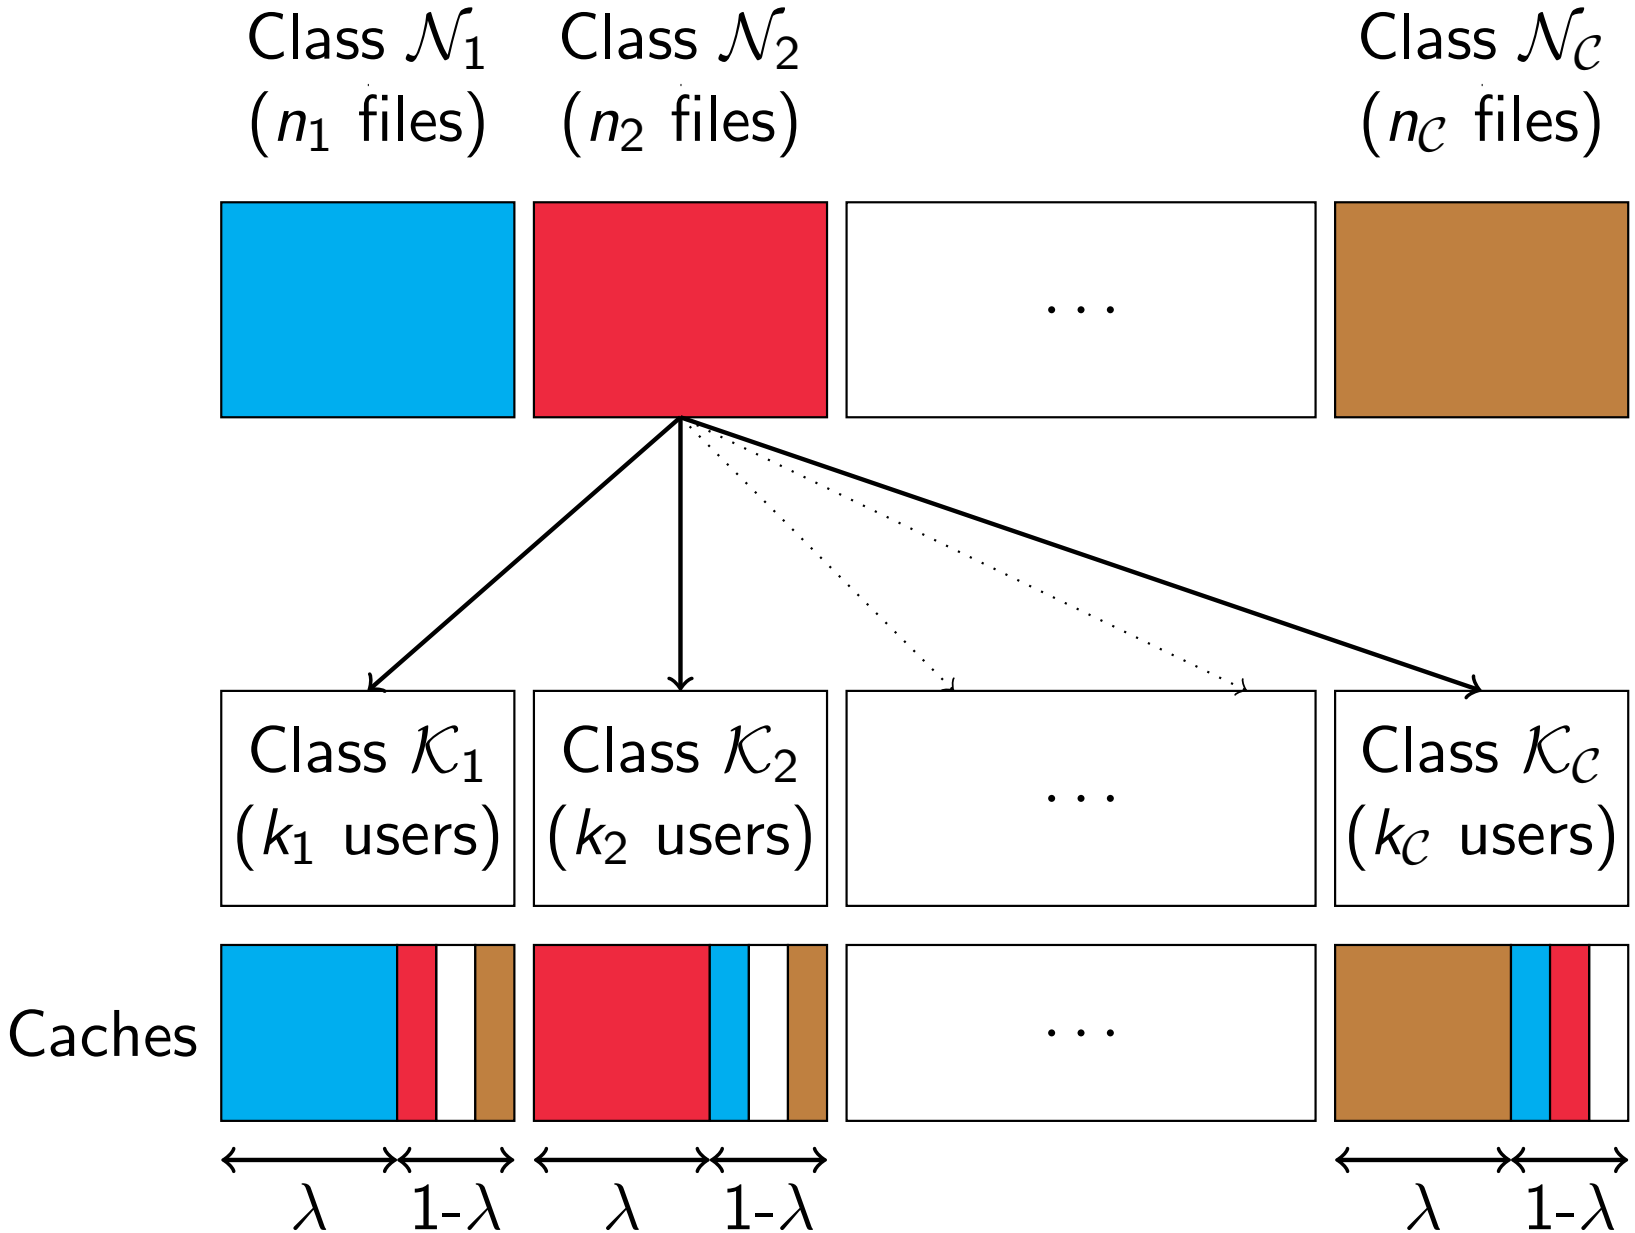

Local preference + File perspective delivery

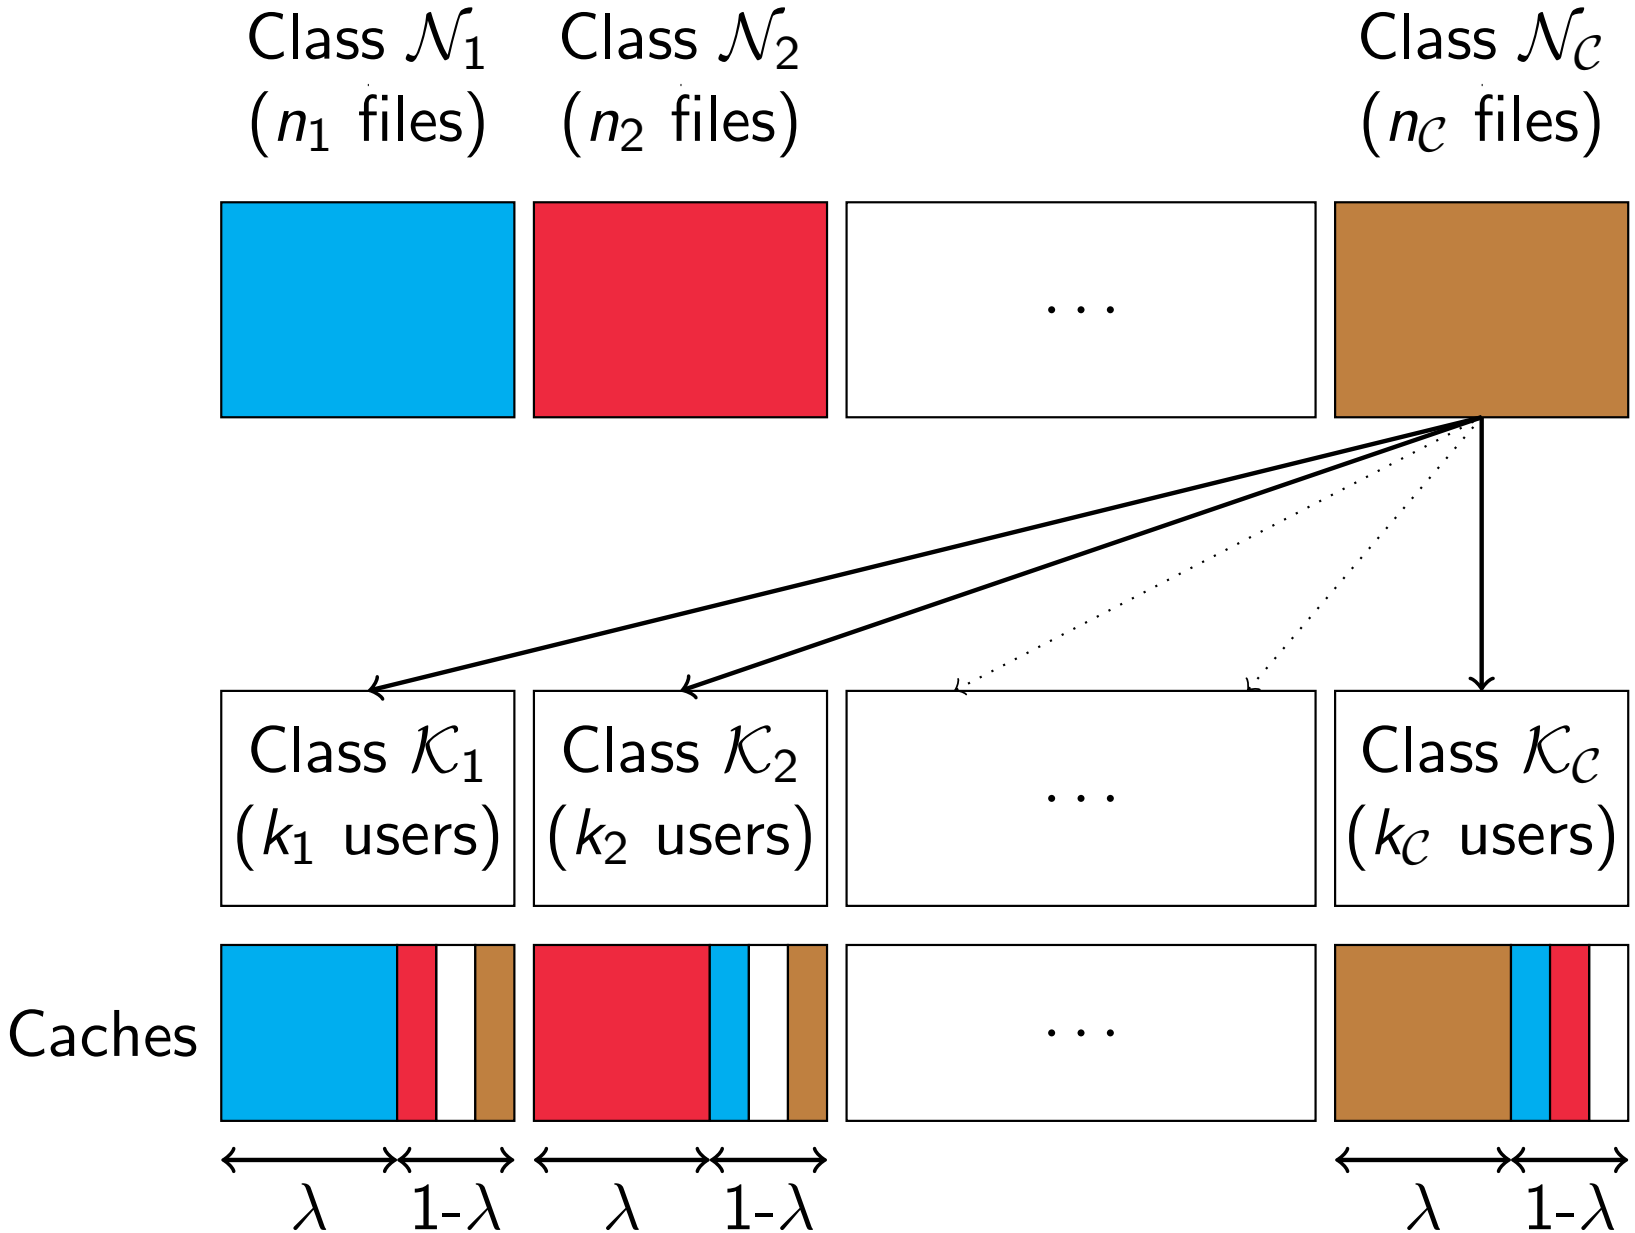

Local preference + File perspective delivery

$$R_H = \min_{\lambda\in[0,1]} \left(\left[1 - \frac{(1-\lambda)M\mathcal{C}}{N(\mathcal{C}-1)}\right]_+ \min\left\{\frac{N(\mathcal{C}-1)}{(1-\lambda)M}, K(1-\alpha)\right\} + \left[1 - \frac{\lambda M\mathcal{C}}{N}\right]_+ \left[1 - \frac{(1-\alpha)(1-\lambda)M\mathcal{C}}{N(\mathcal{C}-1)^2}\right]^{\frac{K(\mathcal{C}-1)}{\mathcal{C}}} \min\left\{\frac{N}{\lambda M}, K\alpha\right\}\right)$$

where $[x]_+ = \max\{x, 0\}$

Simulations

- $N = 1000$, $K = 200$, $\mathcal{C} = 10$

a) $\alpha = 0.7$

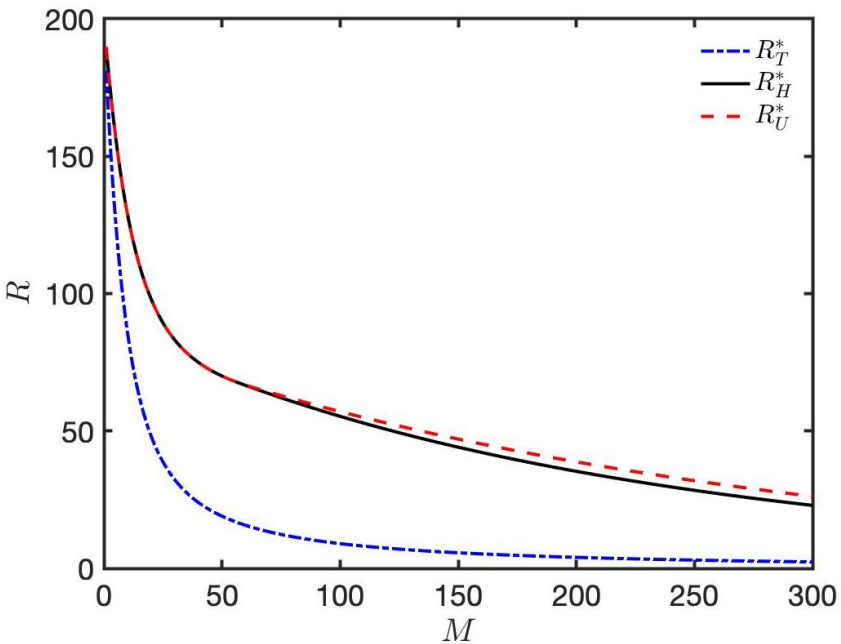

b) $\alpha = 0.975$

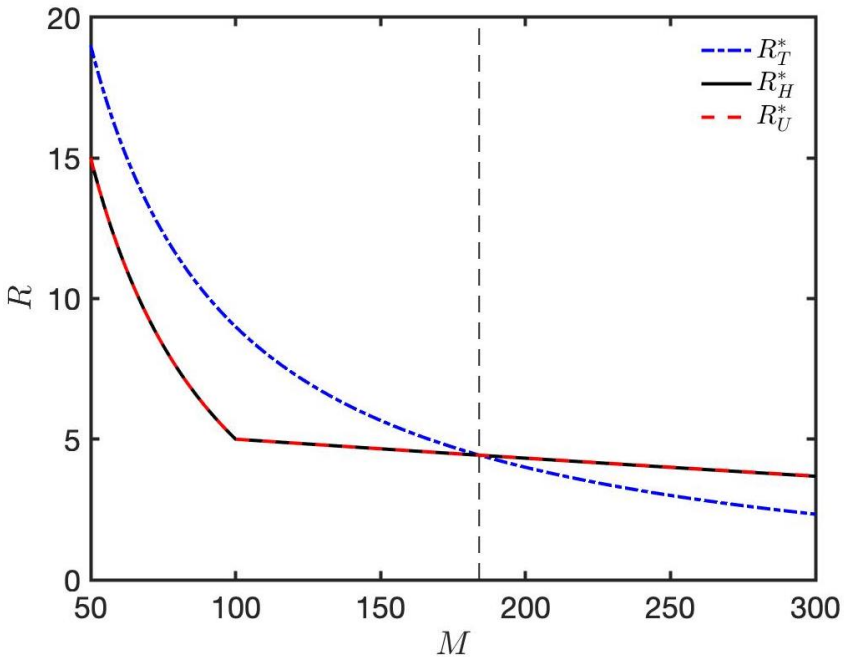

Conclusion

Quantify the level of skew (α) beyond which it is preferable to utilize local content popularity in making placement and delivery decisions.

Conclusion

Quantify the level of skew (α) beyond which it is preferable to utilize local content popularity in making placement and delivery decisions.

- $\alpha > 1 - \frac{C-1}{K}$
  $\longrightarrow$ 'Local Preference' scheme outperforms 'Global Preference' scheme in some memory regime.

Conclusion

Quantify the level of skew (α) beyond which it is preferable to utilize local content popularity in making placement and delivery decisions.

- $\alpha > 1 - \frac{\mathcal{C}-1}{K}$
  $\longrightarrow$ 'Local Preference' scheme outperforms 'Global Preference' scheme in some memory regime.
- $\alpha \leq 1 - \frac{\mathcal{C}-1}{K}$
  $\longrightarrow$ 'Global Preference' scheme outperforms 'Local Preference' scheme in (almost) all memory regimes.