CTW 2010

Scheduling, Dynamic Mode Switching and Dimension Bottlenecks in MU-MIMO Cellular Downlink

Giuseppe Caire, University of Southern California, Los Angeles, CA

Cancun, May 10-12 2010

Collaborators

- Nihar Jindal, Niranjay Ravindran (U-Minnesota);
- Mari Kobayashi (Supelec);
- Babis Papadopoulos, Sean Ramprashad (DoCoMo USA Labs);
- Hoon Huh, Hooman Shirani-Mehr, Krishna Raj Kumar, Mike Neely (USC);
- Antonia Tulino (Alcatel-Lucent, Bell Labs).

Outline of this talk

- A general framework for scheduling and resource allocation.
- Incorporating state information uncertainty: ergodic rates versus outage rates and HARQ. What should we care about?
- Examples: random ICI and channel prediction errors.
- Dimensionality bottleneck of cooperative MU-MIMO (block-fading model): Shall I coordinate multiple base-stations?

Single-cell downlink

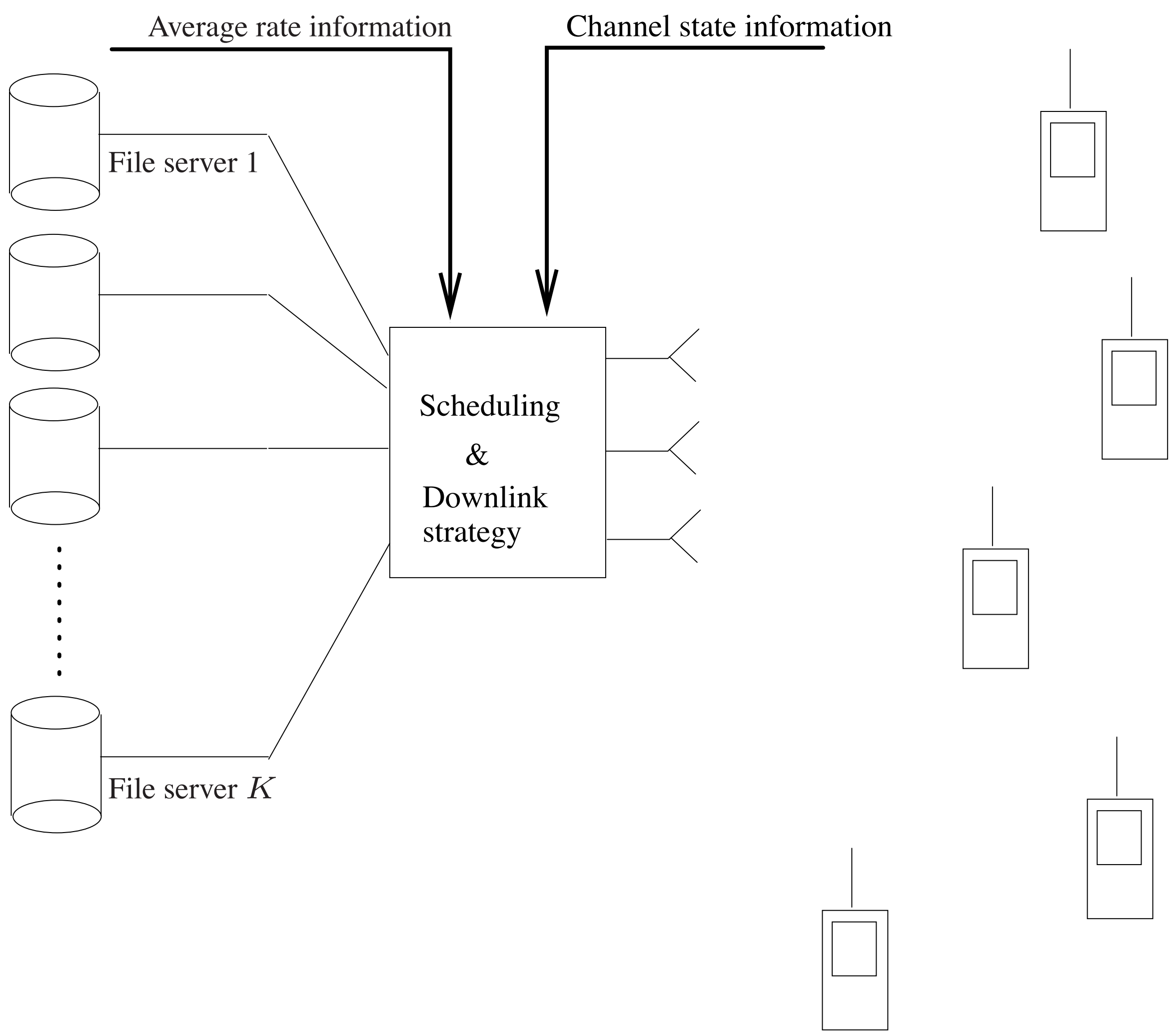

The scheduling and resource allocation problem

- Time is slotted.
- The system is characterized by a time-varying state $S(t)$.
- A given PHY layer (including signaling strategy and feasibility constraints) is characterized by the set of instantaneous rate functions:

$$R_k(s, \gamma(s)), \quad \gamma \in \Gamma(P)$$

- Under suitable ergodicity and stationarity assumptions, the long-term average throughput region is given by

$$\overline{\mathcal{R}} = \text{coh} \bigcup_{\gamma \in \Gamma(P)} \left\{ \overline{\mathbf{R}} \in \mathbb{R}_+^K : \overline{R}_k \leq \mathbb{E}[R_k(S, \gamma(S))] \right\}$$

- For a desired Network Utility Function $g(\cdot)$, we wish to operate at the point solution of:

$$\text{maximize} \quad g(\overline{\mathbf{R}}), \quad \text{subject to} \quad \overline{\mathbf{R}} \in \overline{\mathcal{R}}$$

- $\overline{\mathcal{R}}$ is convex and, for any meaningful setting, bounded.

- Typically, we choose $g(\cdot)$ to be concave and componentwise non-decreasing.

- Therefore, the above problem is convex .... but it is not easy to solve since $\overline{\mathcal{R}}$ is usually difficult to describe (e.g., given in terms of an uncountable number of linear constraints).

- In addition, we would like to have a dynamic policy that learns adaptively the system statistics, rather than a collection of static policies for each possible system statistics.

Network utility examples

- Example. Proportional Fair Scheduling:

$$g(\overline{\mathbf{R}}) = \sum_{k=1}^{K} \log \overline{R}_k$$

- Example. Max-Min Fair Scheduling:

$$g(\overline{\mathbf{R}}) = \min_{k \in \{1,\ldots,K\}} \overline{R}_k$$

- Example. Differentiated Rate Scheduling:

$$g(\overline{\mathbf{R}}) = \sum_{k=1}^{K} \omega_k \overline{R}_k$$

A general result (Georgiadis, Neely and Tassioulas, 2004)

- At every t, pick the instantaneous rates $R_k(t) = R_k(S(t), \gamma^\star(S(t)))$ such that $\gamma^\star$ is the solution to:

$$\text{maximize} \quad \sum_{k=1}^{K} Q_k(t) R_k(S(t), \gamma(S(t))), \quad \text{subject to} \quad \gamma \in \Gamma(P)$$

where $\mathbf{Q}(t)$ are scheduler weights.

- Let the arrival processes be given by $\mathbf{A}(t) = \mathbf{a}^\star$, solution of

$$\text{maximize} \quad V g(\mathbf{a}) - \sum_{k=1}^{K} Q_k(t) a_k, \quad \text{subject to} \;\; 0 \leq a_k \leq A$$

- Update the weights according to the "virtual queue" evolution equation:

$$Q_k(t+1) = [Q_k(t) - R_k(t)]_+ + A_k(t)$$

- Let

$$\overline{\mathbf{R}}(V, A) = \lim_{t \to \infty} \frac{1}{t} \sum_{\tau=0}^{t-1} \mathbf{R}(\tau)$$

denote the long-term average rate K-tuple achieved by the above dynamic policy with parameters V and A.

- It is possible to choose V and A such that

$$g(\overline{\mathbf{R}}(V, A)) \geq g(\overline{\mathbf{R}}^{\star}) - O(1/V)$$

with average virtual buffer length (a measure of the convergence delay)

$$\lim_{t \to \infty} \frac{1}{t} \sum_{\tau=0}^{t-1} \mathbb{E}[\mathbf{Q}(\tau)] = O(V)$$

Imperfect channel state information

- In the above policy, the instantaneous weighted sum-rate maximization requires perfect knowledge of $S(t)$.
- This is usually not the case!
- Let $S(t)$ and $\widehat{S}(t)$ denote the true channel state and the CSIT.
- The PHY instantaneous rate functions are $R_k(s, \gamma(\widehat{s}))$ (generally different from before!).
- The system long-term average throughput region becomes

$$\overline{\mathcal{R}} = \text{coh} \bigcup_{\gamma \in \Gamma(P)} \left\{ \overline{\mathbf{R}} \in \mathbb{R}_+^K : \overline{R}_k \leq \mathbb{E}[R_k(S, \gamma(\widehat{S}))] \right\}$$

Using retransmissions

- In the presence of uncertainty, the scheduled instantaneous rates $\mathbf{R}(t)$ may not be achievable and an information outage occurs with some probability.

- In this case, some form of ARQ is implemented: how should this be done?

- Classical (Hybrid) ARQ:

  1. Rate $R_k(t)$ is scheduled and a codeword is produced;
  2. An error occurs and NACK is sent back;
  3. The same codeword, or additional parity symbols for the same codeword are sent (Chase combining .... and various classical techniques).

- The problem with classical (H)ARQ at the PHY layer is that it does not work well with scheduling!

ARQ at the Logical Link Control (LLC) layer

1. Rate $R_k(t)$ is scheduled and a codeword is produced;
2. An error occurs and NACK is sent back;
3. The corresponding information bits are left in the transmission buffer and will be transmitted (possibly at different rate) later on.

- In this case, the system rate functions are the outage rates

$$R_k(s, \gamma(\widehat{s})) = r_k \times 1 \{ r_k < I_k(s, \gamma(\widehat{s})) \}$$

for some suitable rate threshold $I_k(s, \gamma(\widehat{s}))$ related to achievability.

- The scheduled rate r_k is part of the policy γ.

- We can show that the optimal dynamic scheduling policy in this case solves the problem:

$$\text{maximize} \quad \sum_{k=1}^{K} Q_k(t)\mathbb{E}\left[R_k(S(t), \gamma(\widehat{S}(t))\middle|\, \widehat{S}(t)\right], \quad \text{subject to} \quad \gamma \in \Gamma(P)$$

- It follows that the scheduled rate is given by

$$r_k^\star = \arg\ \max \quad r_k \times \left[1 - \mathbb{P}\left(I_k(S(t), \gamma(\widehat{S}(t))) \le r_k \middle|\, \widehat{S}(t)\right)\right]$$

Incremental redundancy ARQ at the PHY layer

1. When ACK is received for user k, a new packet of B_k bits (notice: independent of t) is encoded using a rateless code;

2. each time t user k is scheduled, if the last feedback was NACK, then the next block of parity symbols from the current packet is transmitted. Otherwise, if the last feedback was ACK, the first block of the newly encoded packet is transmitted.

- We can show that, for B_k sufficiently large, a scheduling policy that updates the virtual buffers $\mathbf{Q}(t)$ using the virtual service rates $R_k(t) = I_k(S(t), \gamma(\widehat{S}(t)))$ operates as close as desired to a virtual system with long-term average throughput region

$$\overline{\mathcal{R}} = \text{coh} \bigcup_{\gamma \in \Gamma(P)} \left\{ \overline{\mathbf{R}} \in \mathbb{R}_+^K : \overline{R}_k \leq \mathbb{E}[I_k(S, \gamma(\widehat{S}))] \right\}$$

- The optimal dynamic scheduling policy in this case solves the problem:

$$\text{maximize} \quad \sum_{k=1}^{K} Q_k(t)\mathbb{E}\left[I_k(S(t), \gamma(\widehat{S}(t)) \,\middle|\, \widehat{S}(t) \right], \quad \text{subject to} \quad \gamma \in \Gamma(P)$$

- Notice that with this scheme there is no explicit rate allocation.

- As a matter of fact, for each user, we can choose the parameter B_k in order to hit a desired throughput/delay tradeoff.

- The parameters B_k can be chosen independently for each user, without affecting the operations of the scheduler.

ARQ and mutual information accumulation

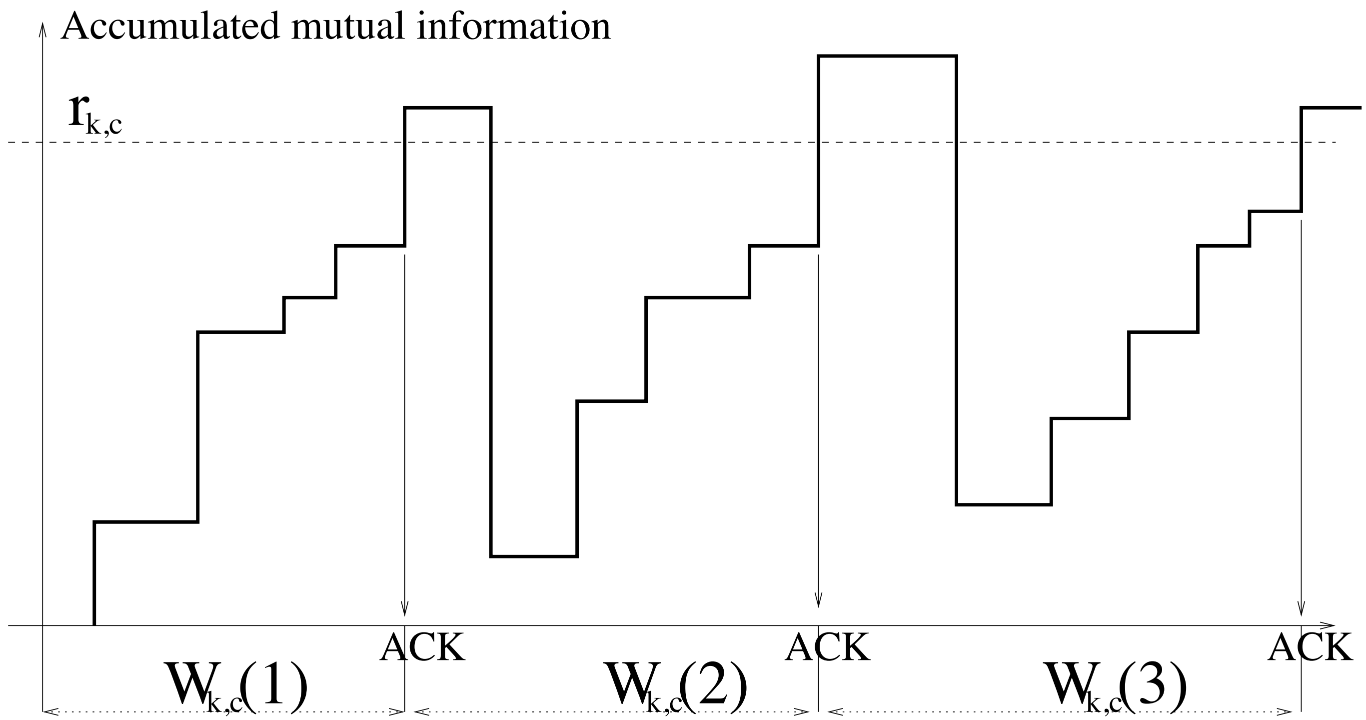

Approaching the virtual genie-aided throughput region

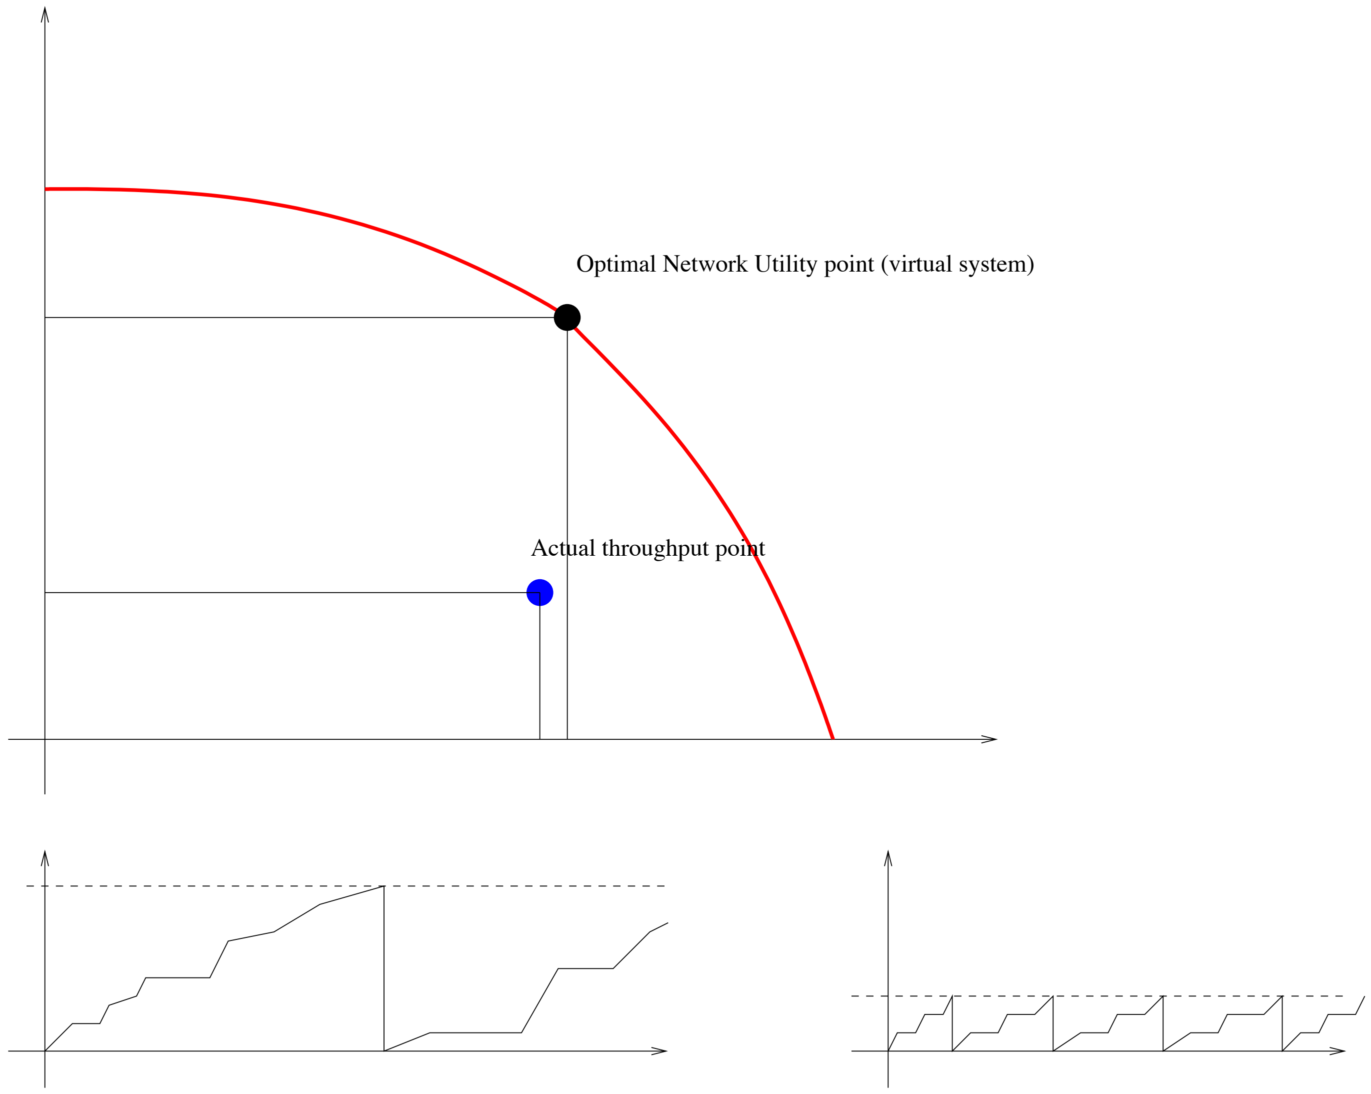

Application to multicell MU-MIMO downlink

- We assume that each cell has perfect CSIT/CSIR for its own users, and knows the pdf of the interference from the other cells to each of its user locations.

- The channel state $S(t)$ and the CSIT for the reference cell c are:

$$S_c(t) = (\mathbf{H}_c(t), \chi_{1,c}(t), \ldots, \chi_{K,c}(t)), \quad \widehat{S}_c(t) = \mathbf{H}_c(t).$$

- The instantaneous WRSM problem that the scheduler has to solve (assuming LZFB downlink precoding) is:

$$\max \quad \sum_{k \in \mathcal{U}_c(t)} Q_{k,c}(t) \int_0^\infty \log\left(1 + \frac{|\mathbf{h}_{k,c}^{\mathsf{H}}(t)\mathbf{v}_{k,c}(t)|^2 p_{k,c}(t)}{1+z}\right) dF_{\chi_{k,c}}(z)$$

Linear layout, 18 cells, $M = 2$ antennas, $K = 36$ users/cell

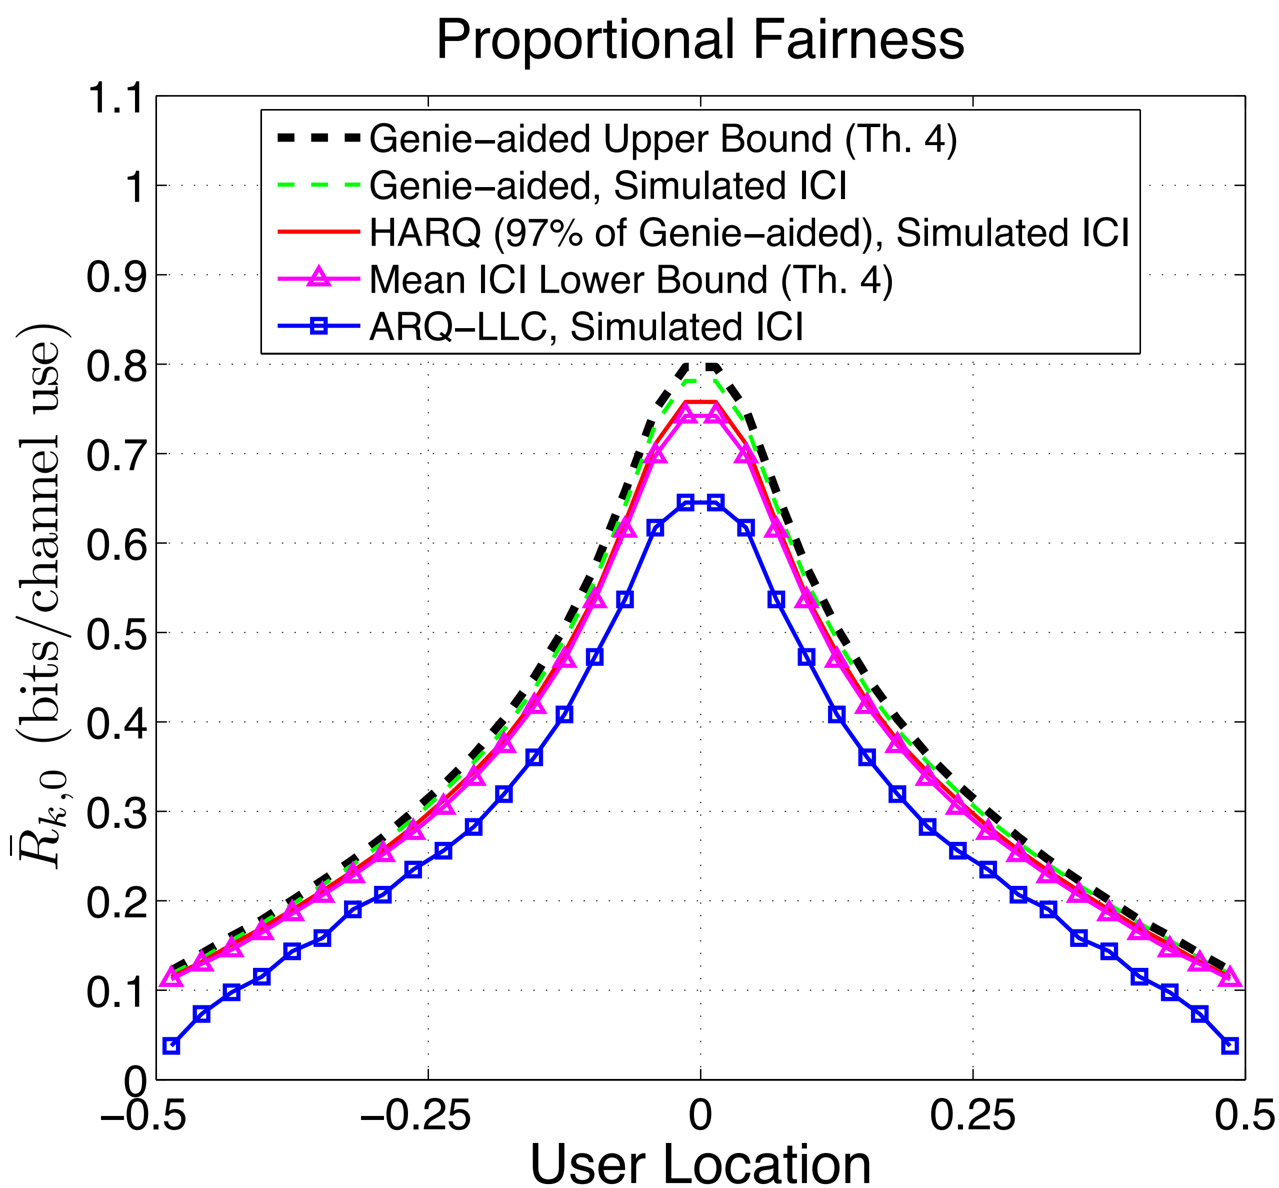

Linear layout, 18 cells, $M = 2$ antennas, $K = 36$ users/cell

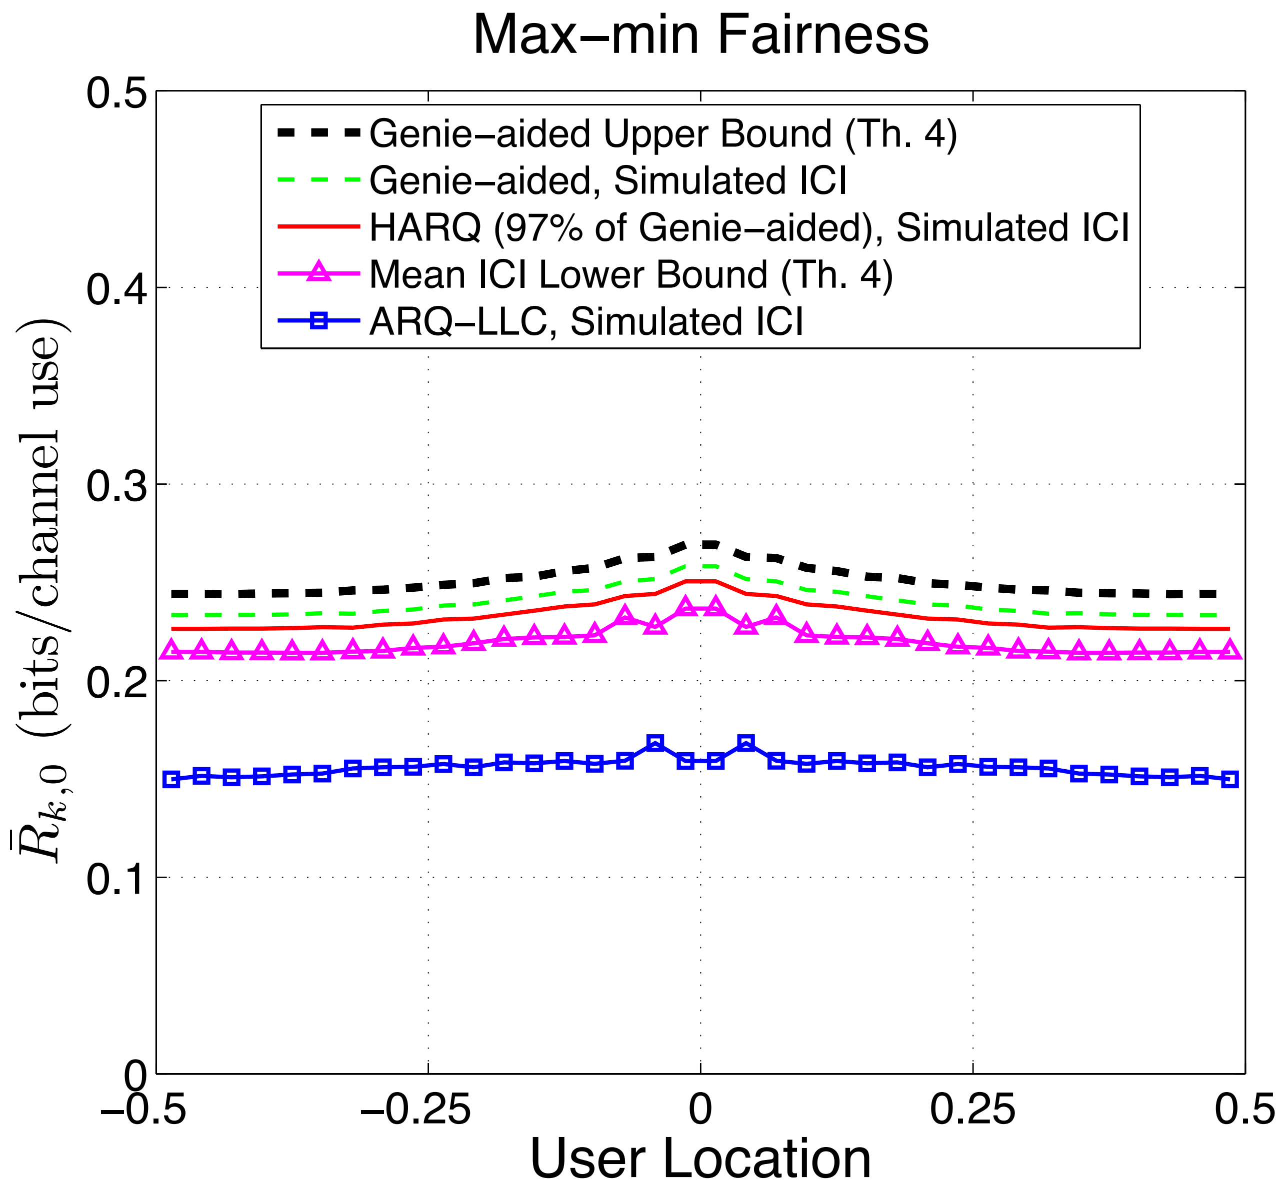

Application to MU-MIMO with channel prediction

- The scheduler should compute:

$$\max \sum_{k \in \mathcal{U}} Q_k(t) \mathbb{E}\left[\log\left(1 + \frac{|\mathbf{h}_k^{\mathsf{H}}(t)\mathbf{v}_k(t)|^2 p_k(t)}{1 + \sum_{j \in \mathcal{U}: j \neq k} |\mathbf{h}_k^{\mathsf{H}}(t)\mathbf{v}_j(t)|^2 p_j(t)} \right) \middle| \widehat{\mathbf{H}} \right]$$

- This is generally very hard. Nevertheless, we observed a sort of threshold behavior in the prediction MMSE.

- Users are partitioned in “predictable” and “non-predictable” according to their prediction MMSE.

- Eventually, the scheduler chooses opportunistically if serving a single non-predictable user with space-time coding or up to M predictable users with multiuser precoding.

Example with 8 users and 4 antennas

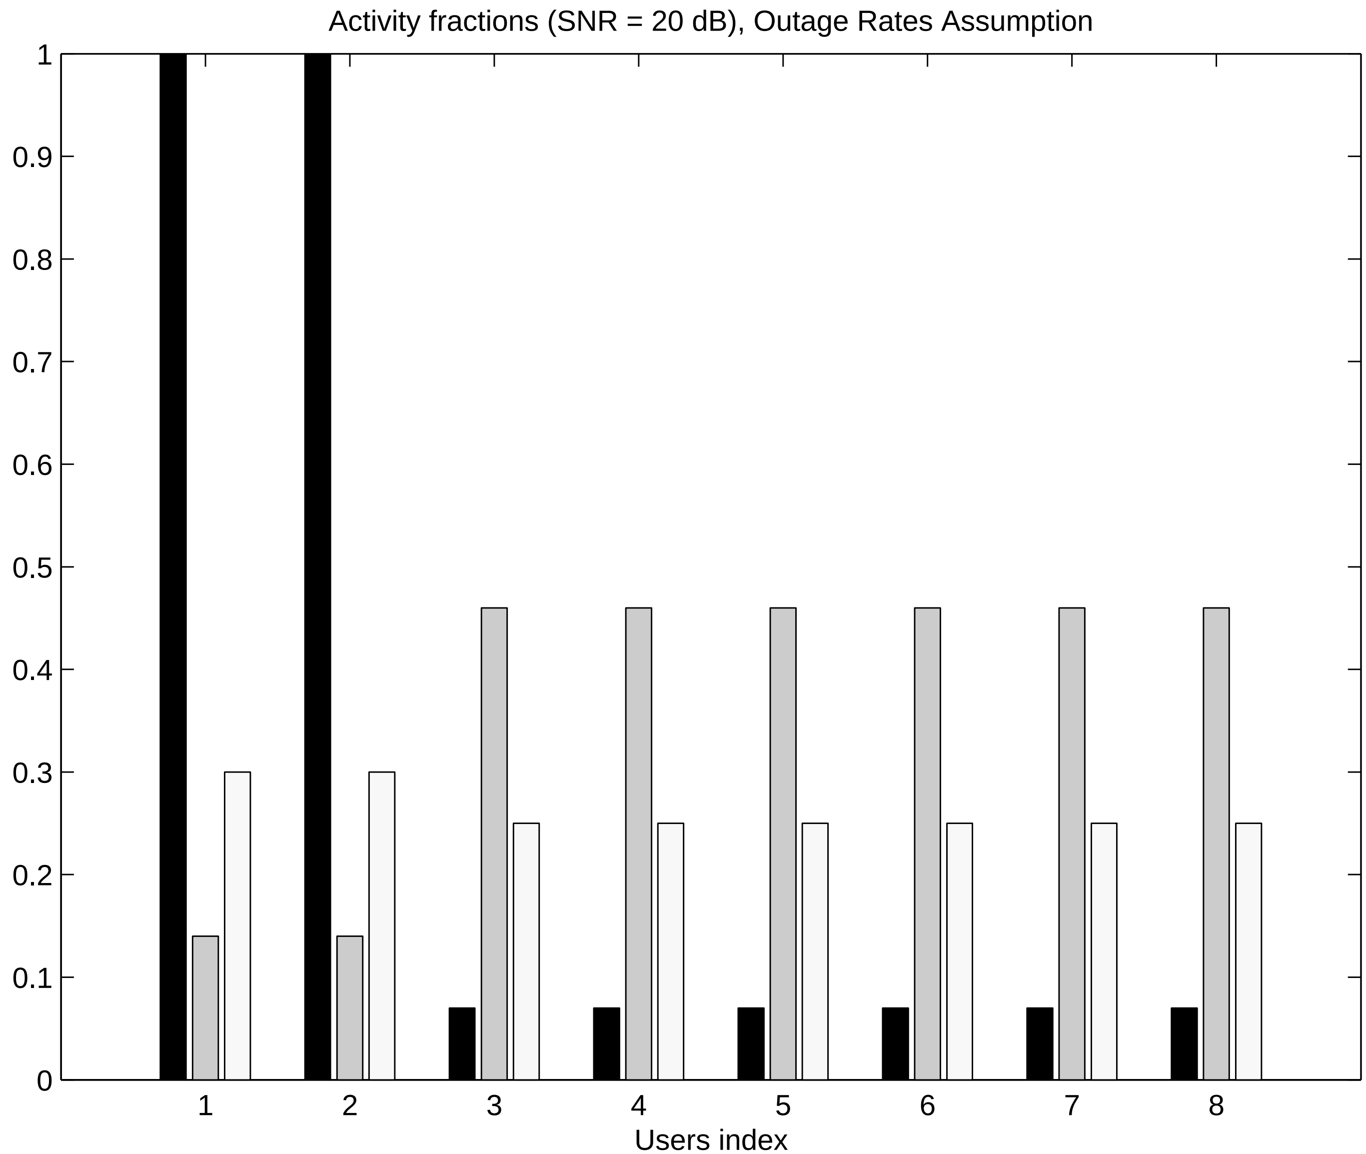

Dimensionality bottleneck: cost of CSIT

- We have established before that long-term average throughput rates (or ergodic rates) matter, if the system is well designed.

- Now, we can use bounds on ergodic rates in order to quantify the cost of CSIT.

- A "quick fix" to ICI is to let base stations to cooperate (radio over fiber ... distributed MIMO systems).

- Since we do not have a clear measure of the complexity of cooperation, let's restrict to the basics:

  fixed maximum power per cell $P_{\max}$ and channel $W \times T$ product

- However, CSIT does not come for free.

FDD, closed-loop training and feedback

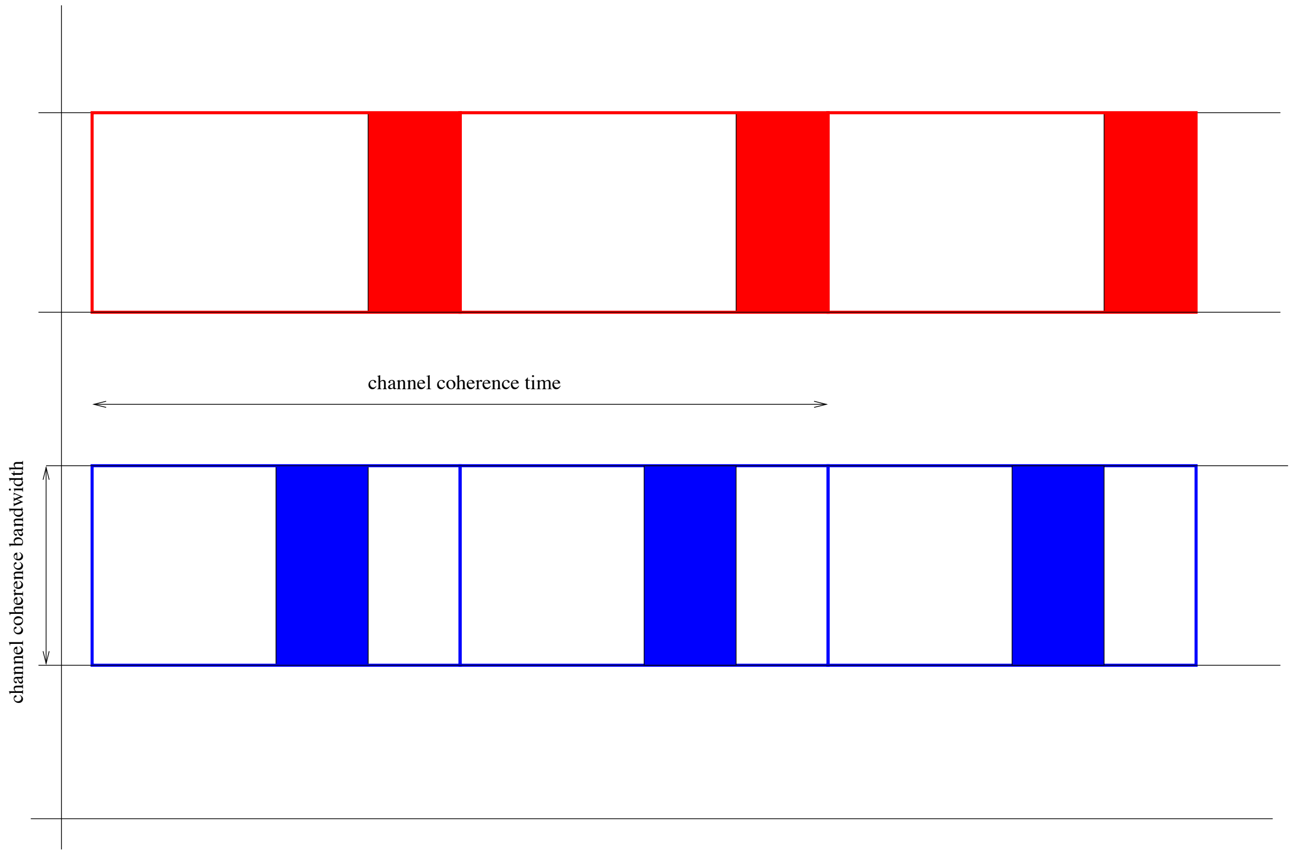

Downlink training and CSIT feedback

- The BS must broadcast at least M symbols (common downlink pilot signal), per channel coherence block of $\approx WT$ dimensions.

- In order to perform scheduling, at least M users (generally more than M) must feedback their CSIT.

- Cooperative MIMO bound: focusing only on downlink training, we know from Marzetta-Hochwald and Zheng-Tse that the high-SNR capacity behaves like:

$$C(\mathsf{SNR}) = M^* \left(1 - \frac{M^\star}{WT}\right) \log \mathsf{SNR} + c + o(1)$$

where $M^* = \min\{M, K, WT/2\}$.

- In FDD, we loose a factor of 2, i.e., for a properly designed training and feedback scheme the following sum-rate is achievable:

$$R_{\text{sum}}(\mathsf{SNR}) = M^* \left(1 - \frac{2M^\star}{WT}\right) \log \mathsf{SNR} + c' + o(1)$$

where $M^* = \min\{M, K, WT/4\}$.

- In a cellular system characterized by distance-dependent pathloss, we have a phenomenon of diminishing return: as we coordinate more and more basestations, the benefit in terms of ICI become negligible with respect to the loss of degree of freedom due to training and feedback.

Quantifying the cost of CSIT in MU-MIMO

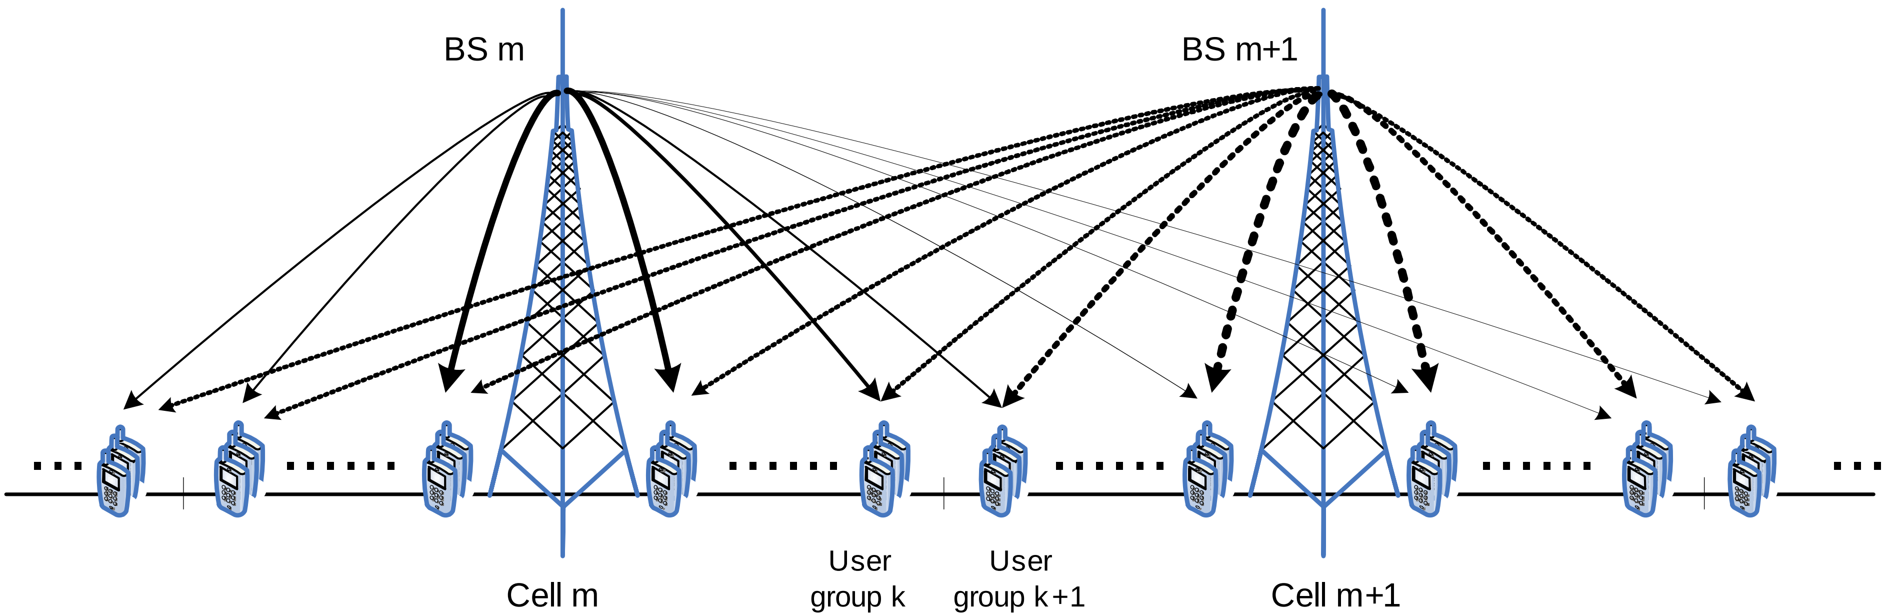

Multi-cell cooperation and inter-cluster interference

- We defined arbitrary cooperation clusters. ICI is treated as noise:

$$\sigma_k^2 = 1 + \sum_{m \notin \mathcal{M}_\ell} \alpha_{m,k}^2 P_m$$

- After a convenient re-normalization, the reference cluster channel is given by

$$\mathbf{y} = \mathbf{H}^{\mathsf{H}}\mathbf{x} + \mathbf{z}$$

with $\mathbf{y} \in \mathbb{C}^{NA}$, $\mathbf{x} \in \mathbb{C}^{\gamma NB}$, $\mathbf{z} \sim \mathcal{CN}(\mathbf{0}, \mathbf{I}_{NA})$ and the channel matrix $\mathbf{H} \in \mathbb{C}^{\gamma NB \times NA}$ given by

$$\mathbf{H} = \begin{bmatrix} \beta_{1,1}\mathbf{H}_{1,1} & \cdots & \beta_{1,A}\mathbf{H}_{1,A} \\ \vdots & & \vdots \\ \beta_{B,1}\mathbf{H}_{B,1} & \cdots & \beta_{B,A}\mathbf{H}_{B,A} \end{bmatrix},$$

where we define $\beta_{m,k} = \frac{\alpha_{m,k}}{\sigma_k}$.

Network utility maximization and fairness

- It is fundamental to consider the system under some well-defined fairness criterion.
- Otherwise, the sum-rate may be maximized by allocating zero rate and power to users in critical conditions edge users.
- The network utility maximization framework said before applies here:

$$\begin{aligned} \text{maximize} \quad & g(\overline{\mathbf{R}}) \\ \text{subject to} \quad & \overline{\mathbf{R}} \in \overline{\mathcal{R}} \end{aligned}$$

where $\overline{\mathcal{R}}$ is the achievable ergodic rate region for the reference cluster.

Large system limit

- In order to overcome the analytical difficulties, we consider the limit for $N \to \infty$ and fixed γ, A, B.

- It turns out that the matrix $\mathbf{H}_{\boldsymbol{\mu}}$ has the "Girko-type" structure

$$\mathbf{H}_{\boldsymbol{\mu}} = \mathbf{H}_{\mathrm{iid}} \odot \mathbf{P}$$

  where masking matrix $\mathbf{P}$ imposes a certain variance profile given by the path-gains $\beta_{m,k}^2$.

- We exploit some known results in order to arrive at a closed-form solution.

Non-perfect CSIT

- For each such block, $\gamma_p NB$ dimensions are dedicated to common downlink training: each user in the cluster estimates its the corresponding column of $\mathbf{H}$, of length γNB coefficients (it must be $\gamma_p \geq \gamma$).

- The ratio $\gamma_p/\gamma > 1$ denotes the "pilot dimensionality overhead", relative to the minimum number of pilots dimensions for MMSE estimation.

- After some rather standard algebra, given the jointly Gaussian model, we arrive at

$$\mathbf{H} = \widehat{\mathbf{H}} + \mathbf{E},$$

where

$$\widehat{\mathbf{H}} = \left[\begin{array}{ccc} \widehat{\beta}_{1,1}\mathbf{H}_{1,1} & \cdots & \widehat{\beta}_{1,A}\mathbf{H}_{1,A} \\ \vdots & & \vdots \\ \widehat{\beta}_{B,1}\mathbf{H}_{B,1} & \cdots & \widehat{\beta}_{B,A}\mathbf{H}_{B,A} \end{array} \right],$$

and

$$\mathbf{E} = \left[\begin{array}{ccc} \bar{\beta}_{1,1}\mathbf{E}_{1,1} & \cdots & \bar{\beta}_{1,A}\mathbf{E}_{1,A} \\ \vdots & & \vdots \\ \bar{\beta}_{B,1}\mathbf{E}_{B,1} & \cdots & \bar{\beta}_{B,A}\mathbf{E}_{B,A} \end{array} \right],$$

with

$$\widehat{\beta}_{m,k} = \frac{\beta_{m,k}^2}{\sqrt{1/p + \beta_{m,k}^2}}, \quad \bar{\beta}_{m,k} = \frac{\beta_{m,k}}{\sqrt{1 + p\beta_{m,k}^2}}$$

where $p = \frac{\gamma_p}{\gamma} PB$, and where the blocks $\mathbf{H}_{m,k}$, $\mathbf{E}_{m,k}$ are independent with iid $\sim \mathcal{CN}(0, 1)$ elements.

- A standard technique to lower bound the mutual information $I(u_k^{(j)}; y_k^{(j)} | \widehat{\mathbf{H}})$ yields:

$$I(u_k; y_k | \widehat{\mathbf{H}}) \geq \log \left(1 + \frac{\widehat{\Lambda}_k(\boldsymbol{\mu}) q_k}{1 + \sum_{m=1}^{B} \bar{\beta}_{m,k}^2 P} \right)$$

where we have removed the expectation since in the limit for $N \to \infty$ the term inside the expectation converges to a deterministic limit (computed from the large-random matrix theory results).

- SNR loss: due to non-perfect CSIT, intra-cluster multiuser interference appears:

$$\Lambda_k(\boldsymbol{\mu}) \rightarrow \frac{\widehat{\Lambda}_k(\boldsymbol{\mu})}{1+\sum_{m=1}^{B}\bar{\beta}_{m,k}^2 P}$$

- Dimensionality loss: the system spectral efficiency is scaled by a factor $\left[1-\frac{\gamma_p NB}{WT}\right]_+$, taking into account the training overhead.

- In the large system limit, we let $\tau = \frac{N}{WT}$ to be the ratio between the number of users per group and the dimensions in a time-frequency slot.

- We investigated the system spectral efficiency for fixed τ, in the limit of $N \rightarrow \infty$.

Results: non-perfect CSIT, 8-cell linear topology, $\tau = 1/128$

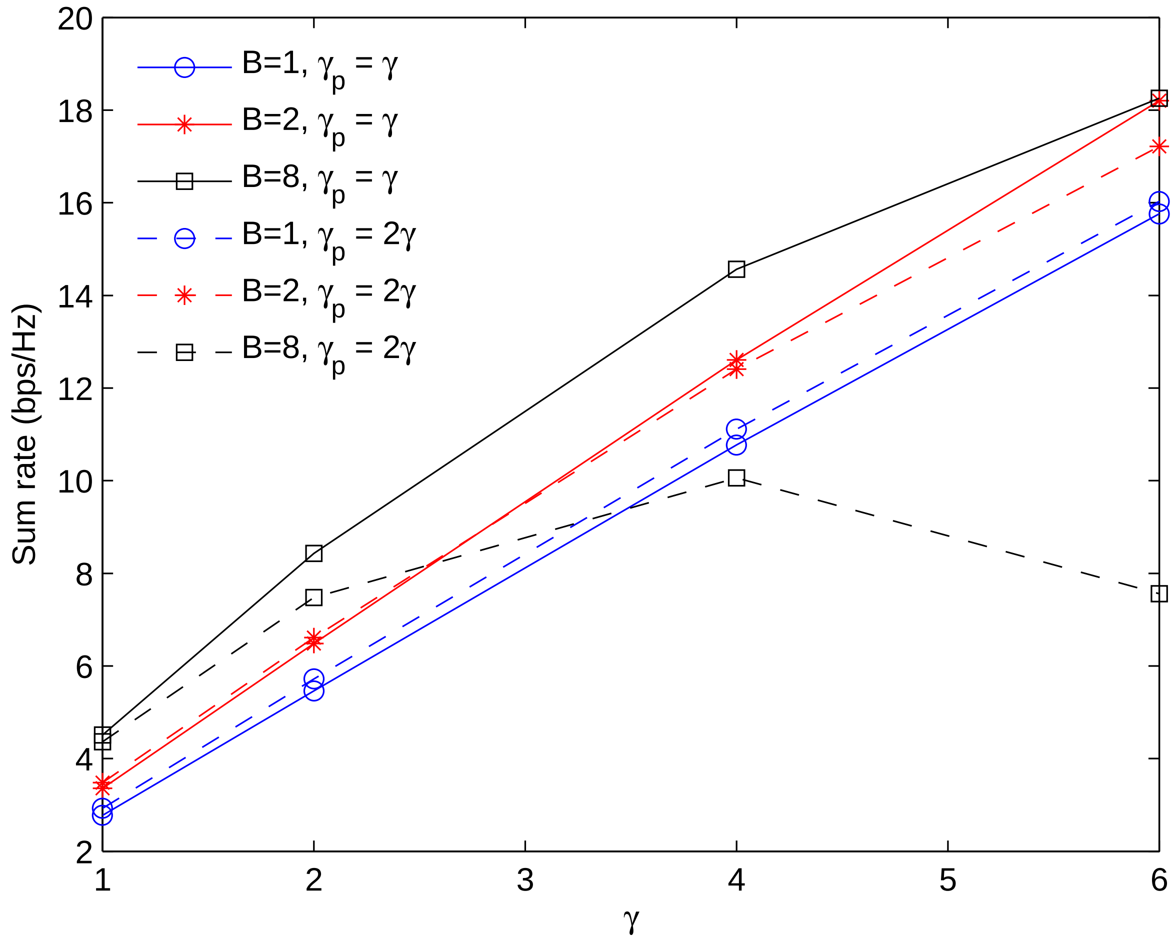

Results: non-perfect CSIT, 8-cell linear topology, $\tau = 1/64$

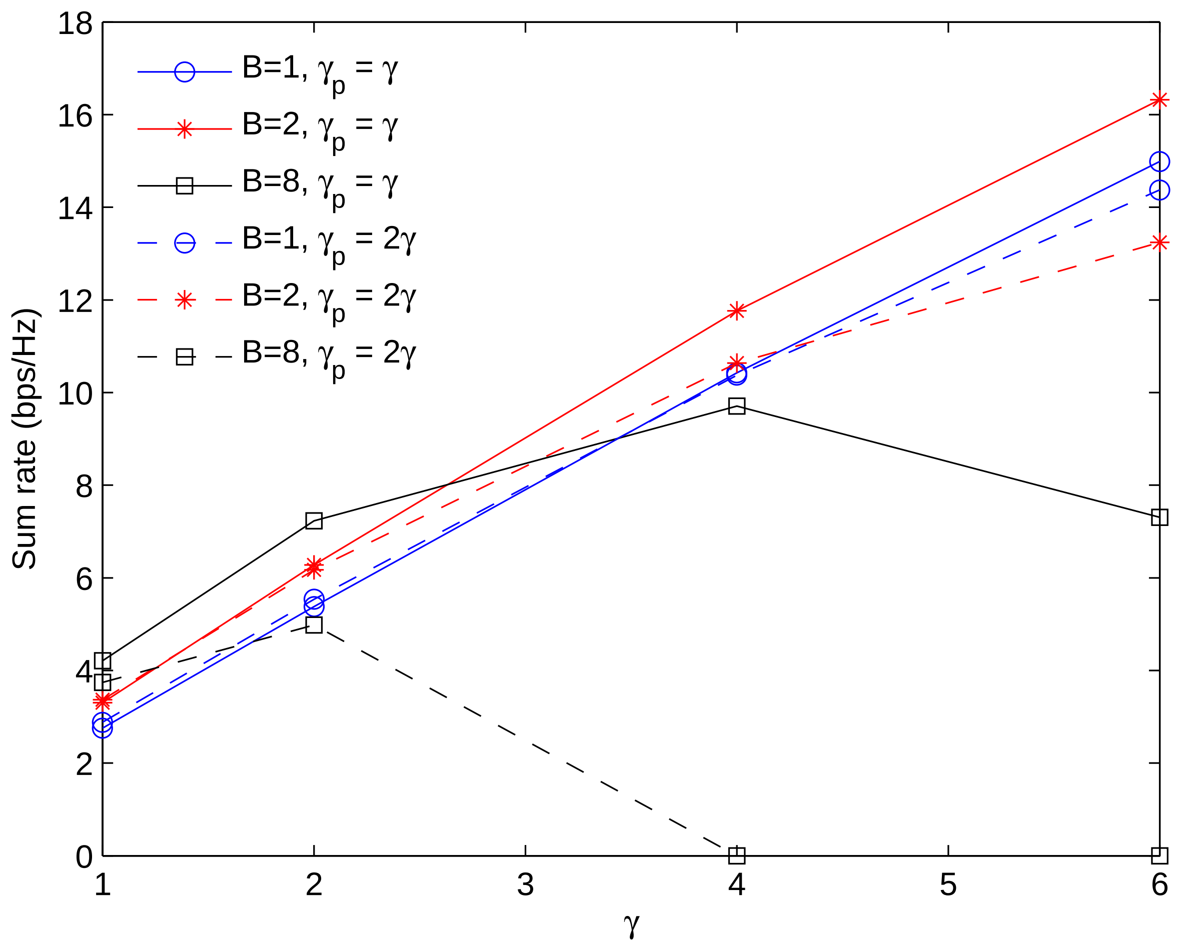

How shall I group my antennas?

- A related but different problem is the following: given that the CSIT overhead depends on the number of jointly processed antennas, for a given overhead what is the best antenna placement?

- We investigated three options:

1. Conventional cellular (BM antennas per basestation, no cooperation);
2. Expanded cellular (BM antennas per basestation, basestations active on different subbands);
3. Overlapped clustering (M antennas per basestation, clusters of size B basestation);

Conventional cellular (reuse 1)

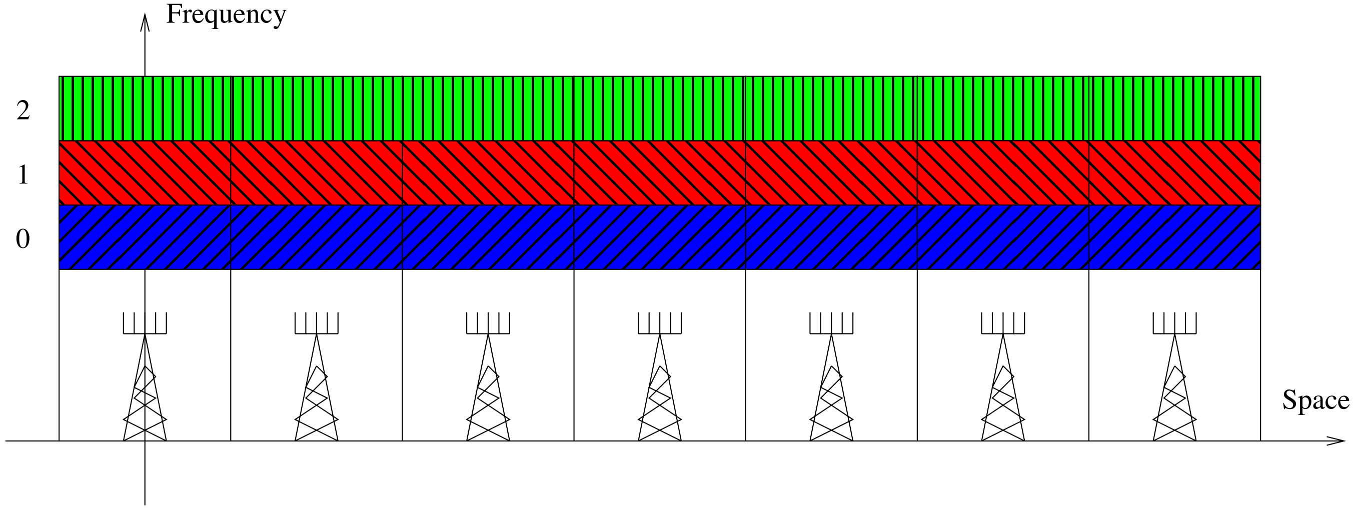

Expanded cellular (reuse 1)

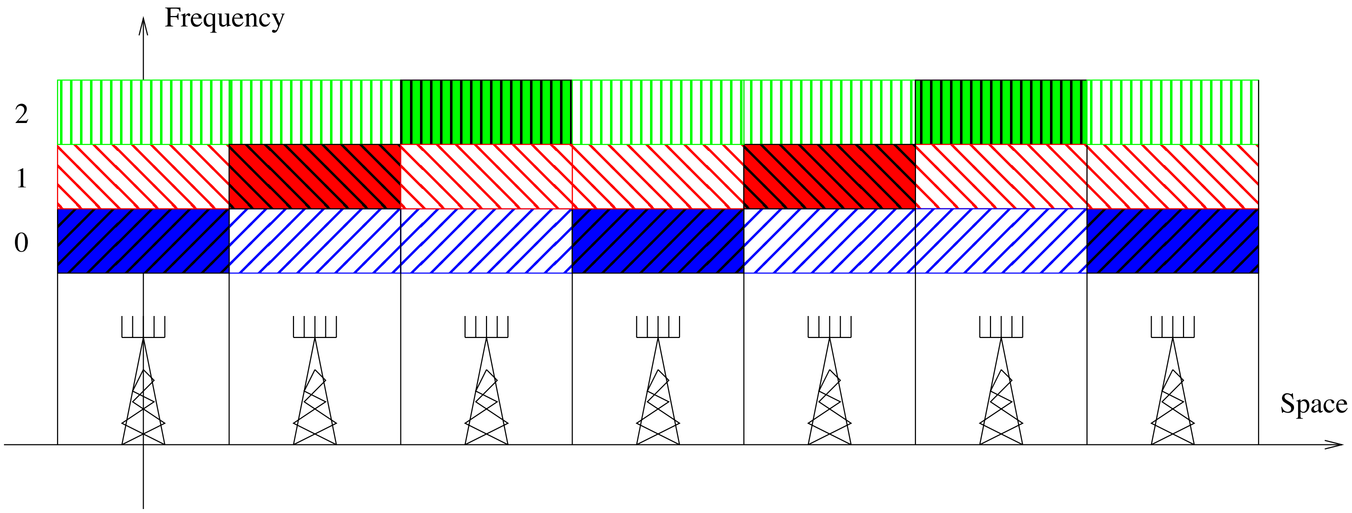

Overlapping clusters (reuse 1)

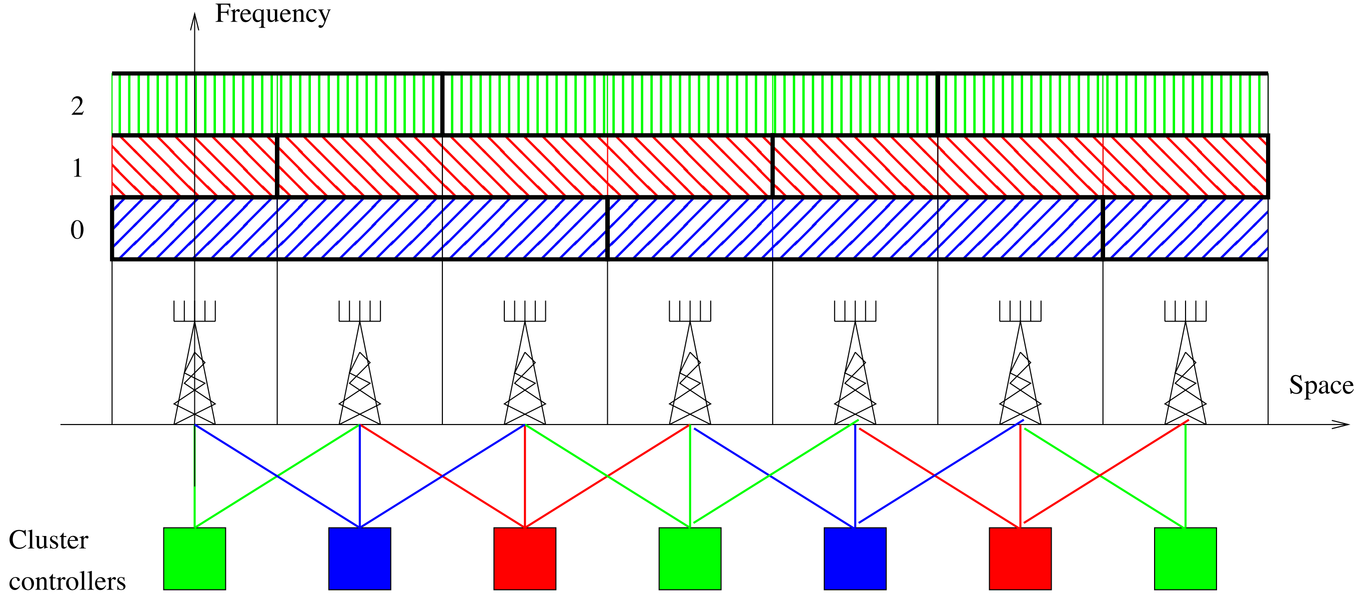

Results are analogous ....

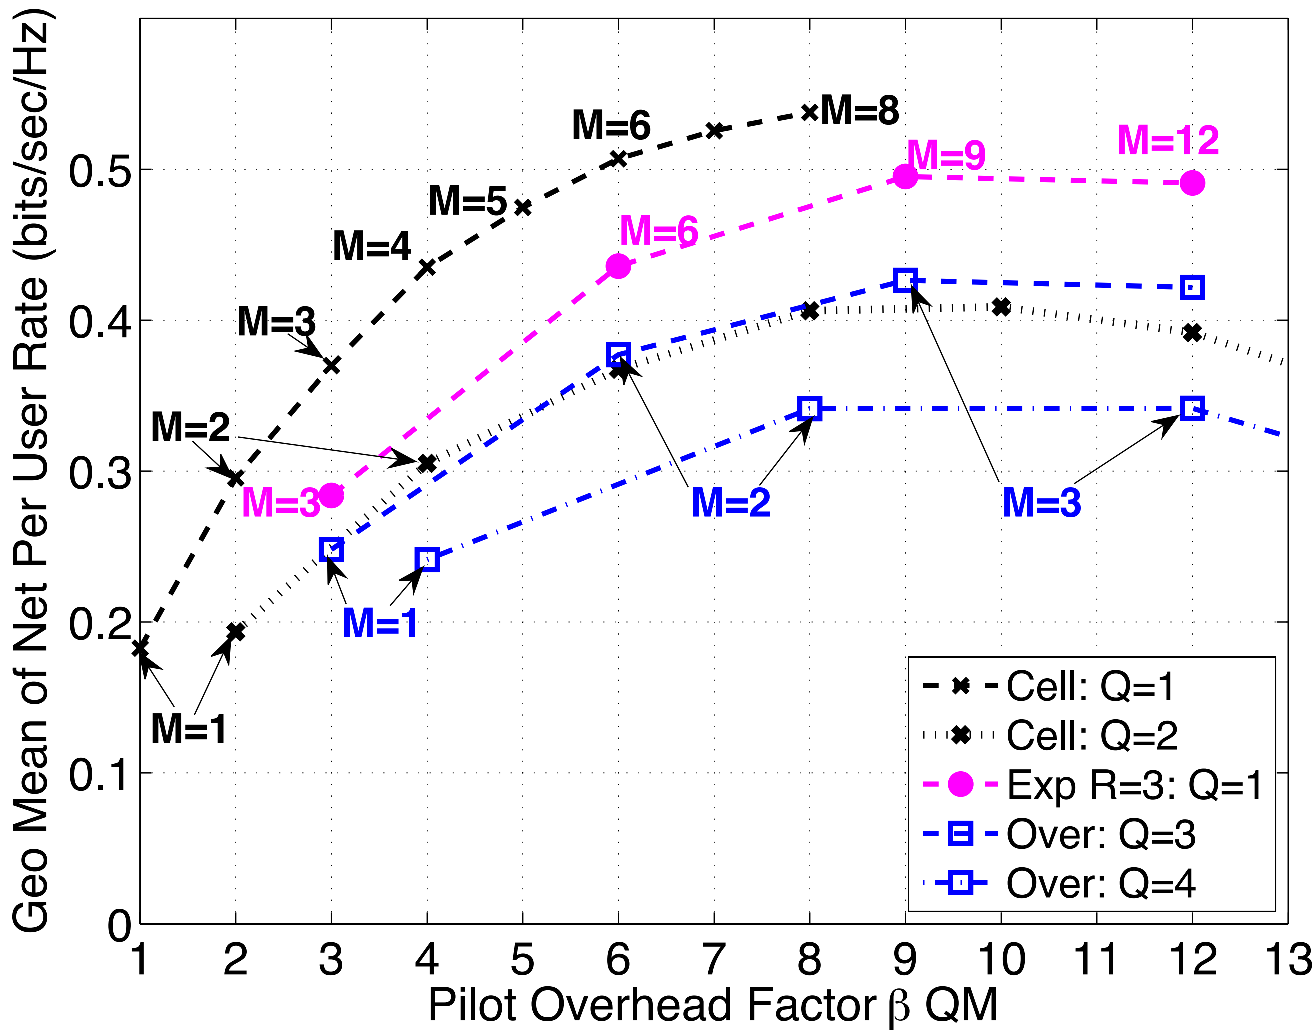

Conclusions

- Ergodic rates (as opposed of outage rates) are meaningful in data-oriented communications.
- The gains of implementing incremental redundancy at the PHY layer jointly with scheduling can be significant, especially for edge users and high-mobility users.
- At least in FDD, the intrinsic cost of estimating CSIT must be taken into account in order to reach any meaningful system design guideline.
- Grouping large number of antennas and avoinding base-station joint processing appears the best strategy, when overhead and fairness are properly taken into account.