OFDMA Downlink Resource Allocation via ARQ Feedback

Rohit Aggarwal, Mohamad Assaad, C. Emre Koksal, and Phil Schniter

Asilomar Conference

Nov. 3, 2009

Setup:

- Single-antenna downlink with K users
- OFDMA with N subchannels
- Channels are Markov time-varying with L taps
- ACK/NAK feedback from previously scheduled users

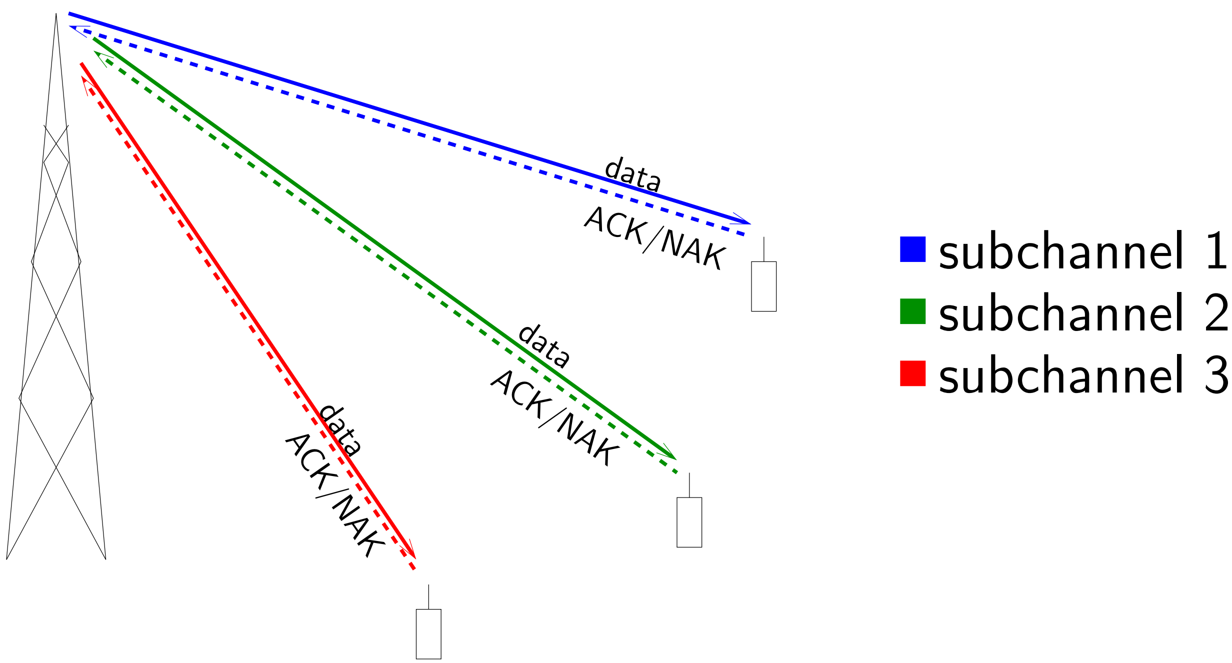

The Basic Resource Allocation Problem:

- At each time t, we want to schedule the "best" users (*multiuser diversity*) to their "best" subchannels (*frequency diversity*).
- We also want to optimize the powers and data-rates of assigned users.
- To make informed choices, we need channel state information (CSI).
- Feedback of each user's CSI about each subchannel is very costly!

Is it possible to do near-optimal resource allocation using only ACK/NAK feedback from previously scheduled users?

Can we learn **enough** *about the CSI from such limited feedback?*

Detailed Objective:

At each time t and subchannel n, choose each user k's next...

- rate $r_{n,k,t+1} \in \{1, \ldots, M\}$,
- power $p_{n,k,t+1} \geq 0$,

based on ACK/NAK feedback $\boldsymbol{F}_1^t$, to maximize the total future utility

$$G_{t+1}^T = \sum_{\tau=t+1}^{T} \sum_{k=1}^{K} \mathrm{E}\left\{ \sum_{n=1}^{N} U\Big(\underbrace{\big(1 - \epsilon_{r_{n,k,\tau}}(\gamma_{n,k,\tau}, p_{n,k,\tau})\big) r_{n,k,\tau}}_{\text{goodput from } k \text{ on } n \text{ at } \tau} \Big) \Bigg| \boldsymbol{F}_1^t \right\}$$

subject to the power constraint $\sum_{n,k} p_{n,k,\tau} \leq P_{\max}, \quad \forall \tau,$

and subject to a one-user-per-subchannel constraint.

Here, $\epsilon_r(\gamma, p)$ is packet error rate and $U(\cdot)$ is a concave utility function.

Optimal ACK/NAK-based Resource Allocation:

- Notice that the current resource allocation affects not only the immediate utility, but also the subsequent ACK/NAK feedback, and hence the future utilities.
- Intuitions:
  - if we assign transmission params that are *very likely to yield ACKs*, we will learn very little about the changing CSI! ($\rightsquigarrow$ "*exploitation*")
  - if we assign transmission params to best inform us of CSI, the expected utility will be low. ($\rightsquigarrow$ "*exploration*")

  Classic tradeoff: *exploration vs exploitation*.
- The optimal allocator is a *partially observable Markov decision process* (POMDP), at least in the simpler case of a finite set of powers. POMDP complexity is impractically high, however, forcing us to consider a suboptimal approach.

Greedy Resource Allocation:

- For ACK/NAK-based rate adaptation in the single-user single-channel case, we previously found that *greedy adaptation* is nearly optimal (at practical fading rates):

  R. Aggarwal, P. Schniter, and C. E. Koksal, "Rate Adaptation via Link-Layer Feedback for Goodput Maximization over a Time-Varying Channel," *IEEE Transactions on Wireless Communications*, Aug. 2009.

- Thus, we propose to use *greedy resource allocation* for our multi-user multi-channel problem.

The Greedy Resource-Allocation Problem:

Using the indicator $I_{n,m,k,t} \in \{0,1\}$ to denote time-t assignment of subchannel n to user k at MCS index m, the time-t problem becomes

$$\max_{\substack{I_{n,k,m,t+1}\in\{0,1\} \\ p_{n,k,m,t+1}\geq 0}} \sum_k \mathrm{E}\left\{ \sum_{n,m} U\Big(I_{n,k,m,t+1}(1-a_m e^{-b_m p_{n,k,m,t+1}\gamma_{n,k,t+1}})r_m \Big) \Big| \boldsymbol{F}_1^t \right\}$$

$$\text{subject to } \sum_{n,k,m} I_{n,k,m,t}\, p_{n,k,m,t} \leq P_{\max}, \quad \forall t,$$

$$\text{and } \sum_{k,m} I_{n,k,m,t} \leq 1, \quad \forall n,\ \forall t,$$

where

- $\gamma_{n,k,t}$ is SNR of user k at subchannel n at time t,
- (a_m, b_m, r_m) determine data rate and error rate for MCS index m
- $\boldsymbol{F}_1^t$ collects all ACK/NAK feedbacks collected from times 1 to t.

Greedy Allocation — Practical Approximation:

Say that we relax the binary indicators to $\tilde{I}_{n,m,k,t} \in [0,1]$.

Then the KKT conditions become (suppressing the time-t notation):

$$\forall n,k,m, \quad \mu = a_m b_m r_m \,\mathrm{E}\{\gamma_{n,k} e^{-b_m p_{n,k,m}\gamma_{n,k}} \mid \boldsymbol{F}\} \tag{1}$$

$$\forall n,k,m, \quad \lambda_n = r_m \,\mathrm{E}\{1 - a_m e^{-b_m p_{n,k,m}\gamma_{n,k}} \mid \boldsymbol{F}\} - \mu p_{n,k,m} \tag{2}$$

where $\{\lambda_n\}_{n=1}^N$ and μ are Lagrange multipliers. A practical alg is then:

1. Initialize μ at a small value.
2. For each subchannel n,
   - For each $(k,m)\ldots$
     - calculate $p_{n,k,m}$ from (1) with $\tilde{I}_{n,k,m} = 1$, forcing $p_{n,k,m} \geq 0$.
     - plug $p_{n,k,m}$ into (2) and calculate the corresponding $\lambda_n(k,m)$.
   - Find $(k^*,m^*) = \arg\max_{(k,m)} \lambda_n(k,m)$.
   - Set $I_{n,k^*,m^*} = 1$ and $I_{n,k,m}|_{(k,m)\neq(k^*,m^*)} = 0$.
3. If $\sum_n p_{n,k^*,m^*} > P_{\max}$, increase μ and repeat, else stop.

Example Performance of Greedy Approximation:

N	K	M	greedy goodput	approximation
1	3	9	5.9884	5.988
1	5	9	6.3501	6.3499
2	3	9	10.3251	10.3249
2	5	9	10.9778	10.9774
3	3	9	14.0573	14.0571
3	5	9	14.9653	14.9651

The practical approximation yields 99.99% of the goodput attained by the true greedy scheme.

Tracking the SNR distribution:

The greedy allocator tracks the SNR by updating the SNR distributions

$$p(\gamma_{n,k,t+1} \mid \boldsymbol{F}_1^t), \quad \forall \text{ users } k \text{ and subchannels } n.$$

The SNR evolves as follows:

- Markov evolution of time-domain channel taps:

$$h_{l,k,t+1} = (1-\alpha)h_{l,k,t} + \alpha w_{l,k,t}, \quad w_{l,k,t} \sim \mathcal{CN}(0,1),$$

- subchannel gains as a function of time-domain channel taps:

$$H_{n,k,t} = \sum_{l=0}^{L-1} h_{l,k,t} e^{-j\frac{2\pi}{N}nk},$$

- subchannel SNRs as a function of subchannel gains:

$$\gamma_{n,k,t} = K|H_{n,k,t}|^2.$$

Tracking the SNR distribution (cont.):

SNR tracking can be done as follows:

$$p(\gamma_{n,k,t+1} \mid \boldsymbol{F}_1^t) = \int_{\boldsymbol{h}_{k,t+1}} \underbrace{p(\gamma_{n,k,t+1} \mid \boldsymbol{h}_{k,t+1})}_{\text{(approx of) Dirac delta}} p(\boldsymbol{h}_{k,t+1} \mid \boldsymbol{F}_1^t) \tag{3}$$

$$p(\boldsymbol{h}_{k,t+1} \mid \boldsymbol{F}_1^t) = \int_{\boldsymbol{h}_{k,t}} \underbrace{p(\boldsymbol{h}_{k,t+1} \mid \boldsymbol{h}_{k,t})}_{\text{Markov prediction}} p(\boldsymbol{h}_{k,t} \mid \boldsymbol{F}_1^t) \tag{4}$$

$$p(\boldsymbol{h}_{k,t} \mid \boldsymbol{F}_1^t) = \frac{p(\boldsymbol{f}_{k,t} \mid \boldsymbol{h}_{k,t}) p(\boldsymbol{h}_{k,t} \mid \boldsymbol{F}_1^{t-1})}{\int_{\boldsymbol{h}'_{k,t}} p(\boldsymbol{f}_{k,t} \mid \boldsymbol{h}'_{k,t}) p(\boldsymbol{h}'_{k,t} \mid \boldsymbol{F}_1^{t-1})} \quad \text{(Bayes rule)} \tag{5}$$

$$p(\boldsymbol{f}_{k,t} \mid \boldsymbol{h}_{k,t}) = \prod_{n=1}^{N} p(f_{n,k,t} \mid \gamma_{n,k,t}(\boldsymbol{h}_{k,t})) \tag{6}$$

$$p(f_{n,k,t} = f \mid \gamma_{n,k,t}) = \begin{cases} \sum_m I_{n,k,m,t} a_m e^{-b_m p_{n,k,m,t} \gamma_{n,k,t}} & f = 0 \\ \sum_m I_{n,k,m,t} (1 - a_m e^{-b_m p_{n,k,m,t} \gamma_{n,k,t}}) & f = 1 \\ 1 - \sum_m I_{n,k,m,t} & f = \emptyset \end{cases} \tag{7}$$

Tracking the SNR distribution (cont.):

Thus, for each user k,

1. measure feedbacks $\boldsymbol{f}_{k,t}$ across all subchannels,
2. compute $p(\boldsymbol{f}_{n,k} \mid \gamma_{n,k,t}(\boldsymbol{h}_{k,t}))$ on $\boldsymbol{h}$-lattice using error-rate rules (6)-(7),
3. compute $p(\boldsymbol{h}_{k,t} \mid \boldsymbol{F}_1^t)$ on $\boldsymbol{h}$-lattice by updating previous posterior via (5),
4. compute $p(\boldsymbol{h}_{k,t+1} \mid \boldsymbol{F}_1^t)$ on $\boldsymbol{h}$-lattice via Markov-prediction step (4),
5. compute $p(\boldsymbol{\gamma}_{k,t+1} \mid \boldsymbol{F}_1^t)$ on $\boldsymbol{\gamma}$-lattice via $\boldsymbol{h}$-to-γ conversion step (3).

This costs $\mathcal{O}\big(KNQ_h^L + KLQ_h^{L+1} + KNQ_\gamma Q_h^L\big)$, where

$Q_h =$ number of grid points used per dimension of $\boldsymbol{h}$-lattice,

$Q_\gamma =$ number of grid points used per dimension of $\boldsymbol{\gamma}$-lattice.

Numerical Experiments:

Setup:

$K = 2$	users
$N = 2$	subchannels
$L = 2$	time-domain channel taps
$\mathrm{E}\{\gamma_{n,k,t}\} = 25\text{dB} = 330$	mean subchannel SNR
$\alpha \in \{0.01, 0.001, 0.0001\}$	channel fading rate
$\rho = 0.33$	subchannel correlation

Plots show (versus packet index t):

- goodput of
  - approximate-greedy with genie-aided CSI
  - approximate-greedy with tracked CSI
  - approximate-greedy with prior CSI (and round robin)
- power/rate/user of approximate-greedy with tracked CSI

Goodput for $\alpha = 0.0001$:

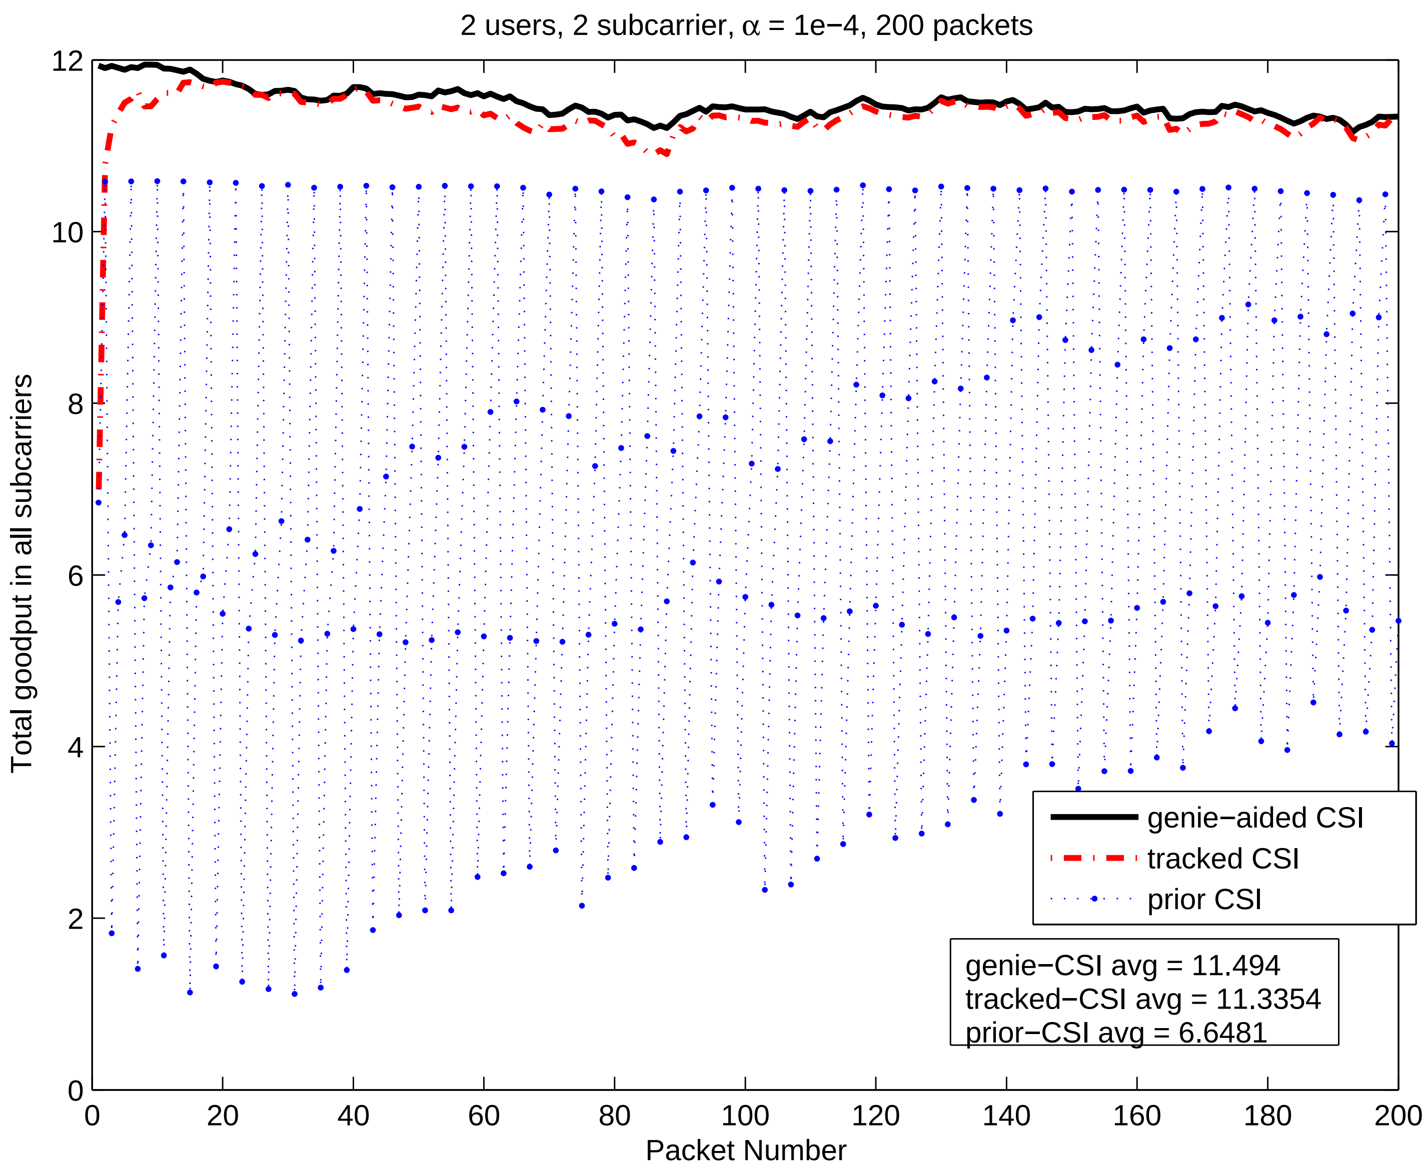

Allocations for $\alpha = 0.0001$:

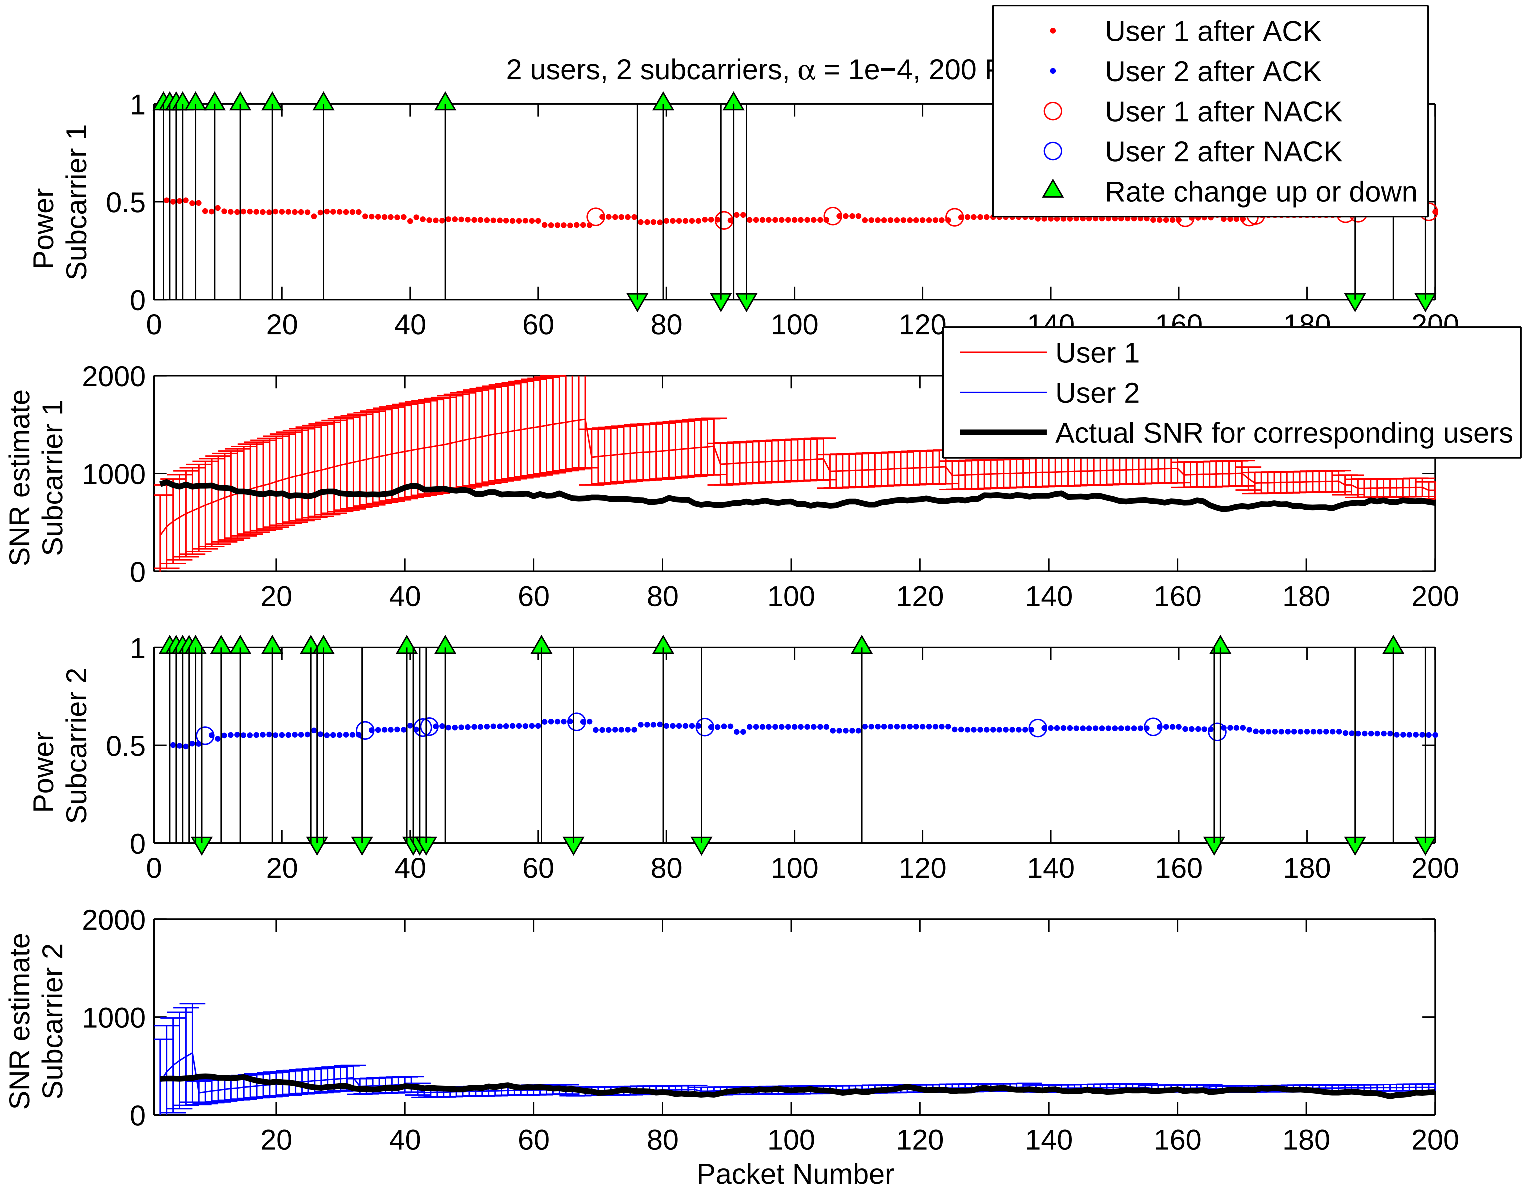

Goodput for $\alpha = 0.001$:

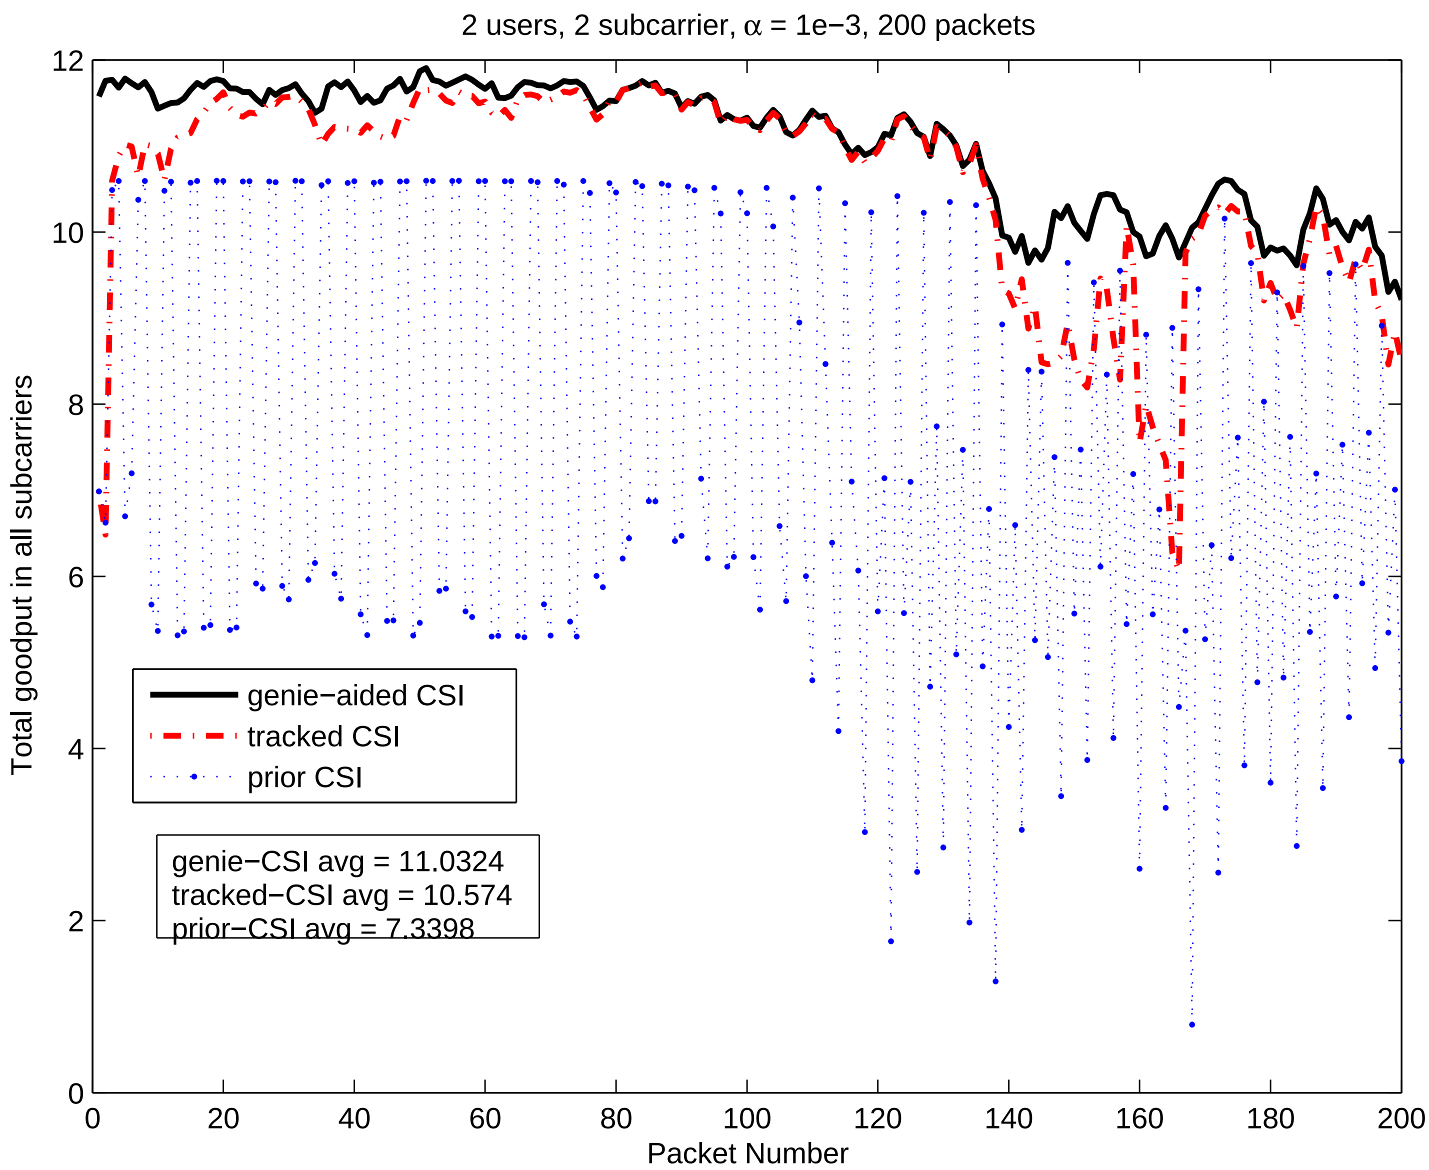

Allocations for $\alpha = 0.001$:

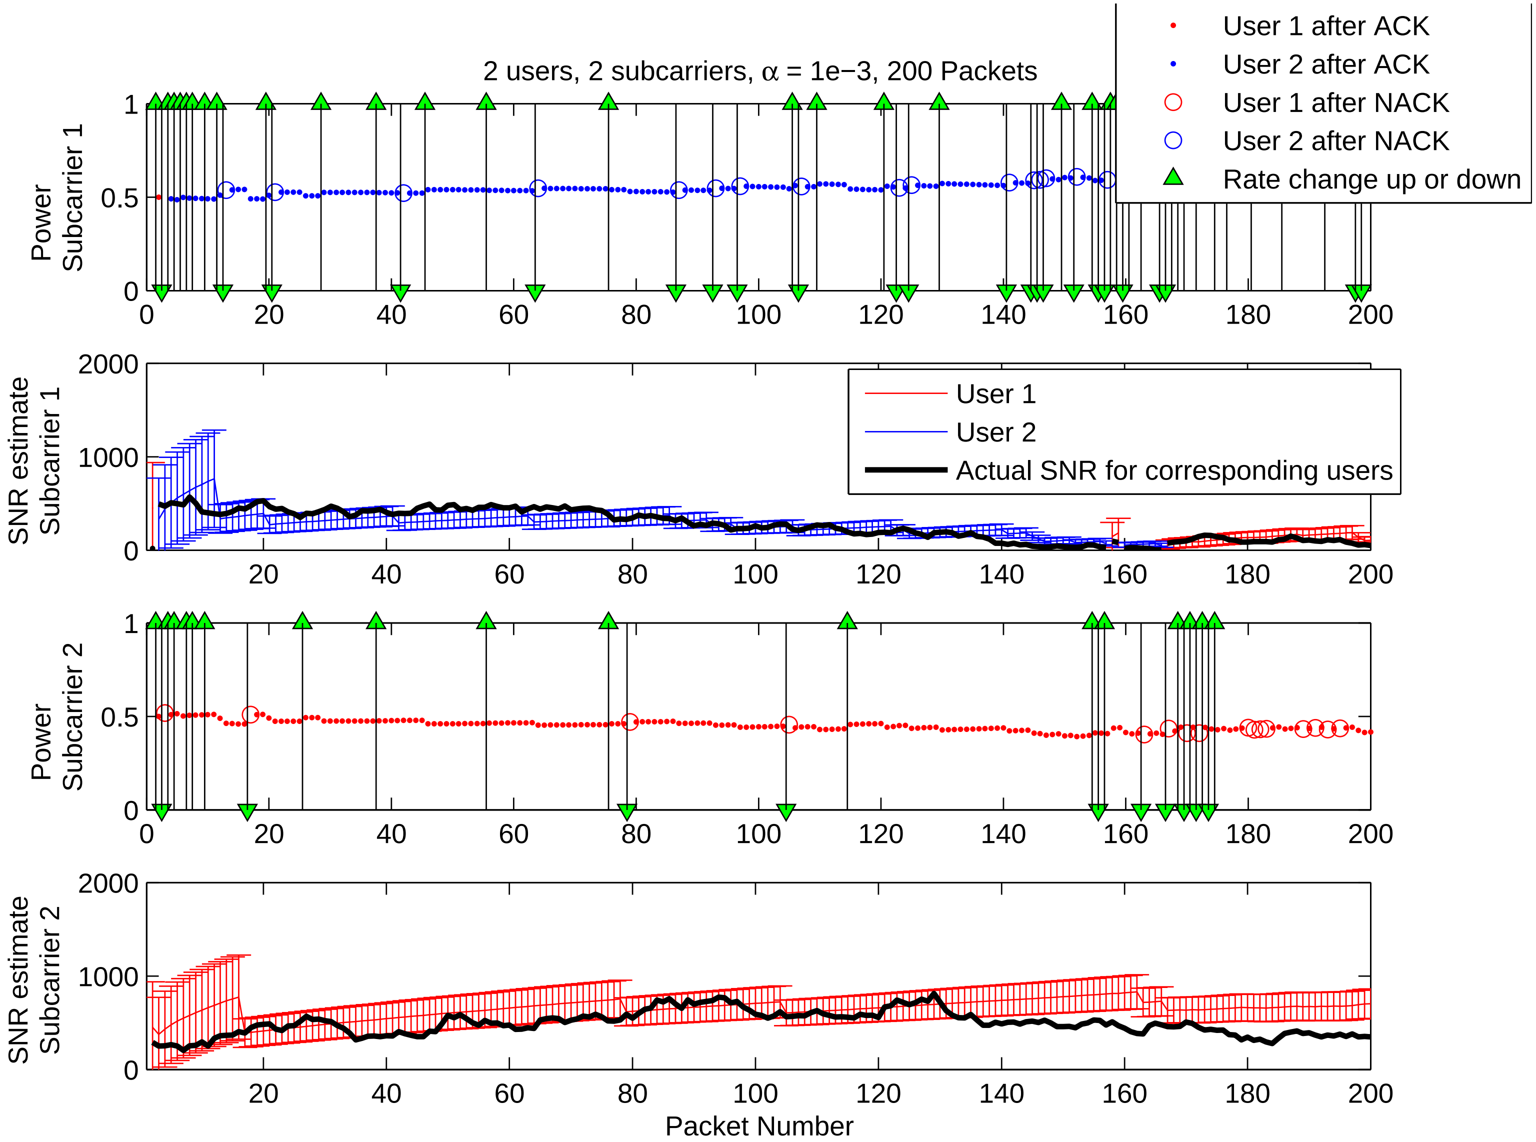

Goodput for $\alpha = 0.01$:

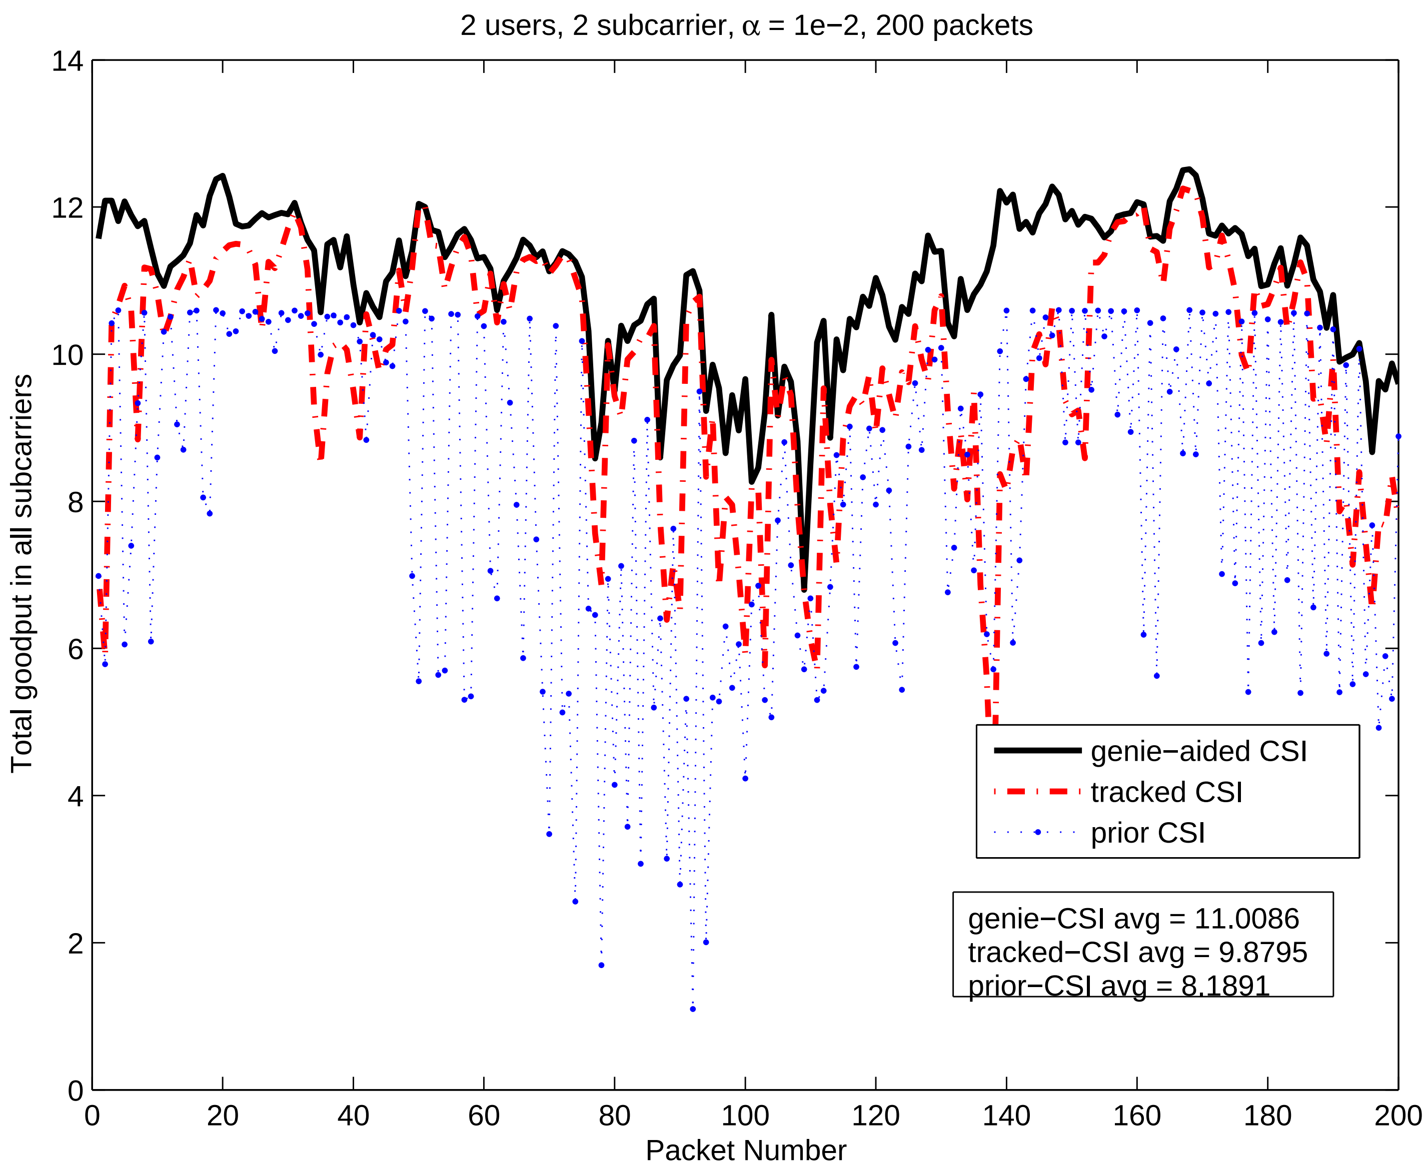

Allocations for $\alpha = 0.01$:

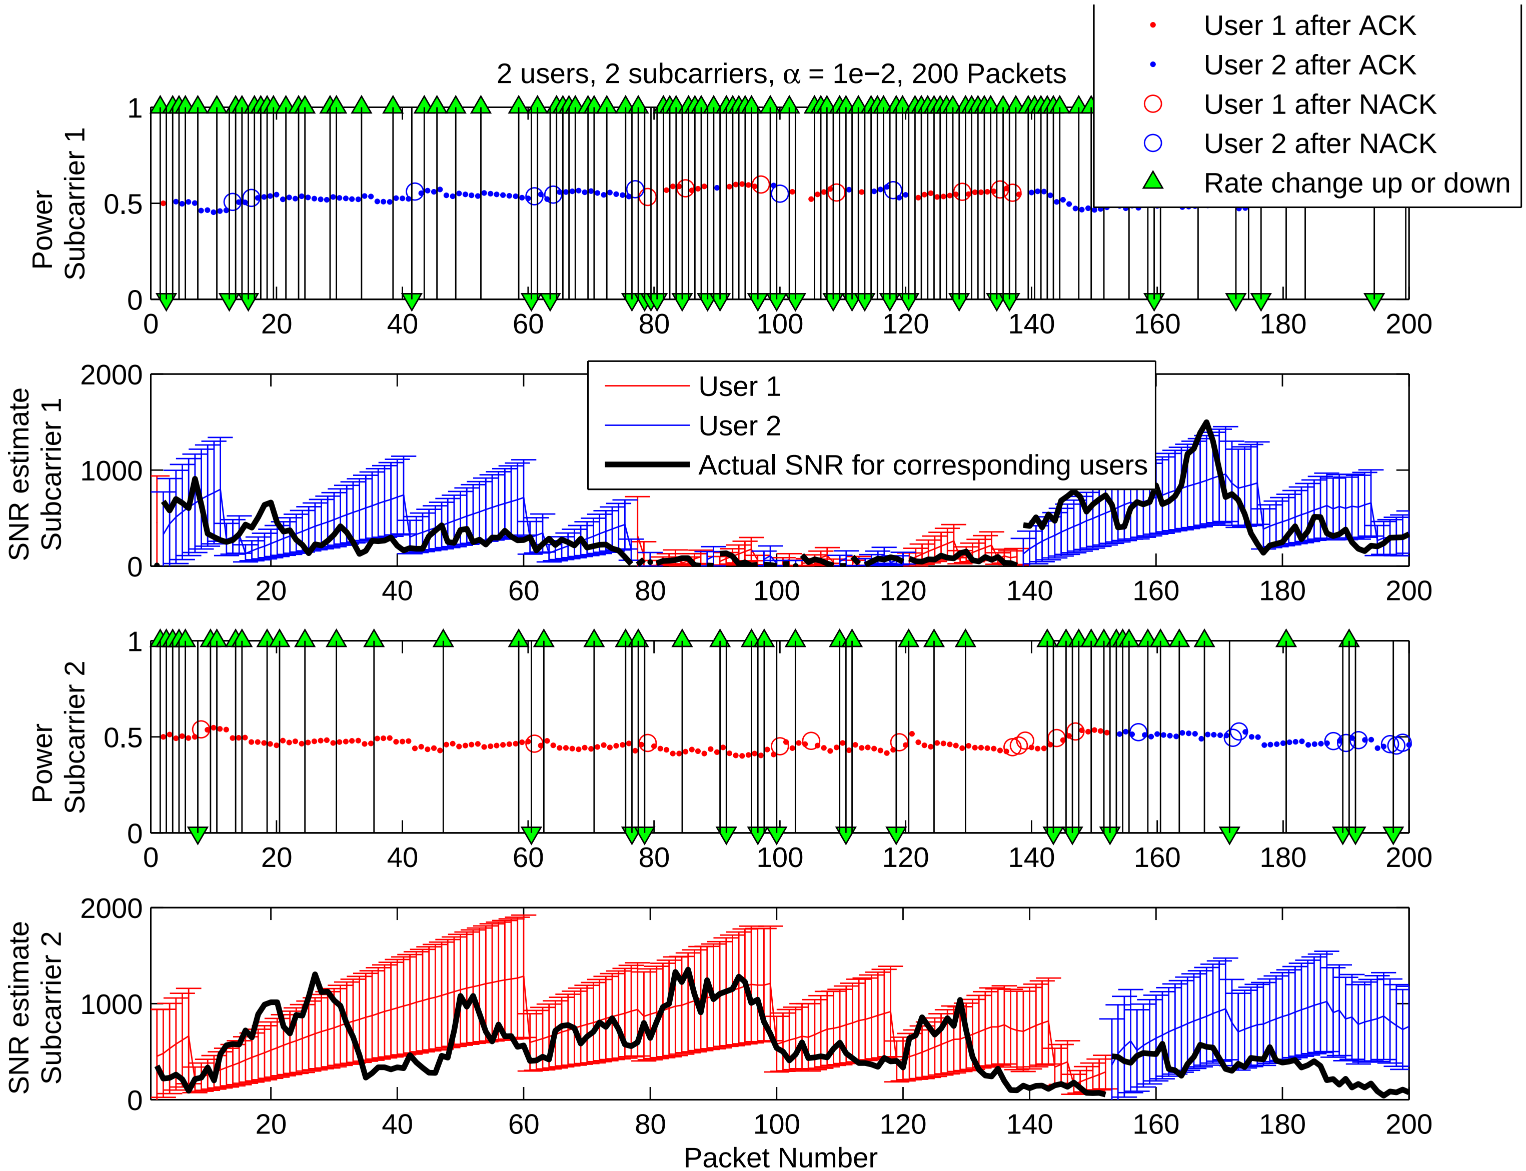

Summary:

- Goal: Allocation of {user schedule, powers, rate} to maximize finite-horizon expected goodput under an instantaneous total-power constraint and a one-user-per-subcarrier constraint.
- The optimal resource allocator is a POMDP, which is computationally impractical.
- We settle for greedy resource allocation, thought to be near-optimal for practical fading rates.
- The greedy allocator itself is computationally impractical, and so we settle for a practical approximation (99.99% exact).
- To maintain CSI, we track the SNR distribution of each user at each subcarrier (conditioned on past ACK/NAK feedback).
- Preliminary experiments for 2 users and 2 subchannels indicates that our practical algorithm does a decent job of SNR tracking and goodput maximization.