Joint downlink power control and channel allocation based on a partial view of future channel conditions

Thi Thuy Nga Nguyen, Olivier Brun, Balakrishna Prabhu

HAL Id: hal-03006880

https://hal.laas.fr/hal-03006880

Submitted on 16 Nov 2020

Joint downlink power control and channel allocation based on a partial view of future channel conditions

T. T. Nga Nguyen[1,2], Olivier Brun[2] and Balakrishna J. Prabhu[2]
[1]Continental Digital Services France
[2]LAAS-CNRS, Université de Toulouse, CNRS, UPS, Toulouse, France

Abstract

We propose two downlink scheduling algorithms that take advantage of partial information on future channel conditions for improving the sum utility. The scheduling model allows for both power control and channel allocation. The objective of the scheduler is the long-term utility under an average power constraint. The two algorithms incorporate the channel predictions in their decisions. The STO1 algorithm computes the decision in each slot based on the means of future channel gains. Depending on the horizon considered, this can require solving a large-dimensional problem in each slot. The STO2 algorithm reduces the dimensionality by operating on two time-scales. On the slower scale it computes an estimation over a larger horizon, and in the faster scale of a slot, it computes the decision based on a shorter horizon. Numerical experiments with both fixed number of users as well as a dynamic number of users show that the two algorithms provide gains in utility compared to agnostic ones.

Keywords — Scheduling, utility maximization, average power constraint

1 Introduction

Downlink scheduling and power control has been widely investigated for wireless networks [1–4]. In each time-slot, one or multiple base-stations have to decide the users to serve and the transmit powers with the objective of maximizing the sum of the utilities of the users. Two different types of utilities can be defined: (i) opportunistic [1, 2]; and *(ii)* long-term [3, 5]. In the opportunistic model, the utility function operates on the rate obtained in a time-slot whereas in the long-term model, the utility operates on the average rate obtained over an horizon.

The focus of this paper is on the long-term utility model. When transmit powers cannot be varied, the celebrated Proportional-Fair algorithm (see [6] and references therein) is known to work well when the utilities are logarithmic functions of the total rate. This algorithm belongs to the class of gradient-based algorithms that choose the user that maximizes the marginal utility or the gradient of the utility function. In [3] the gradient-based solution was then extended to the setting in which joint power control and user scheduling is possible.

We revisit this problem in the context of vehicles which share their itineraries with the decision maker. With the availability of SINR maps in urban zones, the decisions can now be based also upon the future channel conditions of the users (or vehicles) [7]. The future channel conditions are, however, not known exactly as the they vary randomly in time. The SINR maps are assumed to give the expected values of the channel gains on the routes takes by the users. With this additional information on the future expected channel gains, performance improvements can be expected compared to the setting when this information is not available.

We consider a downlink scheduling and power control problem for one base station with an average power constraint. The objective is the sum utility of the users, and the dependence of the utility function on the channel allocation and transmit power is similar to that in [3] with only one sub-channel and we do not have queues. In that respect, our model is a special case of that in [3]. However, the scheduling

algorithms that we propose can be extended to the multiple sub-channel case though for this first paper we restrict ourselves to the simpler case of one sub-channel. There are two difference with the models studied previously. The first, and the main difference is that the base station is also aware of the mean channel gains in the future slots, and the second one is that we include an average power constraint over time in the optimization problem.

1.1 Contributions

We propose two heuristics for that use information on future channel conditions in order to improve the total utility. Both the heuristics were first proposed in [8] in a setting without power control. The first one optimizes over future time-slots in every current slot whereas the second one reduces the complexity by optimizing over the future time-slots only once in a certain number of slots. The performance of these heuristics will be evaluated on three types of stochastic models for the channel gain processes. In the first, the channel gains are a stationary process with fixed mean channel gains (which can be different for different users) whereas in the second model, the mean channel gains are themselves varying on a slower time-scale to that of the channel gains themselves. Finally in the third model, the mean channel gains vary in every slot. For three models, these heuristics are shown to perform better than the other setting in which future information is not available.

1.2 Related work

Scheduling on the downlink but with fixed power examined in numerous papers [9–11]. For a logarithmic utility function, the Proportional-Fair (PF) scheduler has been known to be optimal for stationary channel conditions [6] and its performance has been analyzed in both static and dynamic user scenarios [12]. A restless bandit framework for network utility maximization for channel states modeled as partially observable Markov chains is proposed in [4].

In [13], future rates are assumed to be known accurately over the optimization horizon and improvement in data rates and fairness is compared with algorithms such as greedy (or max-rate) and equal share. Proportional-Fair algorithms with partial information on future channel conditions have been proposed in [14] (on short-time scales) and [7] (on longer-time scales). The algorithm in [7] is based on SINR maps that can be obtained from vehicles or users that were present earlier. They proposed a PF-like index algorithm that takes into account the future rate allocations using round-robin or a few other heuristics. In [8], we proposed two heuristics that solve in each time-slot the utility maximization problem over a short-term horizon assuming that the means of the future rates are known over this short horizon.

The joint optimization of channel allocation and transmit power control has been investigated for different multiplexing schemes such as like CDMA [5] and OFDM [3]. The proposed algorithms based their decisions on the current channel conditions and previous decisions. In the context of high-speed trains, [15] solves the opportunistic utility maximization problem assuming all the future rates are known and with average power constraints.

This paper generalizes the algorithms in [8] for joint channel allocation and transmit power control in order to maximize the total long-term utility. In addition to this, we also include an average power constraint on the base-station. These algorithms can be seen as being similar in spirit to model predictive control [16].

1.3 Organization

The system model and the optimization problem is presented in Section 2 where we also explain the different channel-gain models. In Section 3, we describe the two algorithms that we propose as well as a baseline locally optimal algorithm which can be seen as gradient-based scheduling. These performance of these algorithms is then evaluated numerically in Section 4. Finally, we state the conclusions and future research direction in Section 5.

2 Problem formulation

Consider a base station with K mobile users in its coverage range. In time-slot t, the base station transmits to user i with power $p_i(t)$. Assume that user i has a channel gain of $\gamma_i(t)$ in slot t. These gains will be assumed to stochastic and independent between vehicles but not necessarily stationary for each user.

The received data rate for user i is computed according to the Shannon formula:

$$r_i(t) = x_i(t) \log \left(1 + \frac{\gamma_i(t) p_i(t)}{x_i(t)} \right), \tag{1}$$

where $x_i(t)$ is the fraction of the channel assigned to user i is slot t. User i gets a utility of $U_i(z)$ when it obtains an average rate of z. In order to keep the notation light, we will write $r_i(t)$ as a shorthand for $r_i(x, p, \gamma)$. Let $\mathcal{S}$ be the appropriate dimensional simplex. We shall the use the notation

$$[x_i(t)] \in \mathcal{S} \tag{2}$$

to mean that the vector $[x_1(t), \dots, x_K(t)] \in \mathcal{S}$, where $\mathcal{S}$ is the K-dimensional simplex.

The utility function will be assumed to concave and differentiable. A widely-used class of utility functions is that of the α-fair functions [17] that are parametrized by $\alpha \geq 0$:

$$U_i(z) = \begin{cases} \frac{z^{1-\alpha}}{1-\alpha}, & \alpha \neq 1; \\ \log(z), & \alpha = 1. \end{cases}$$

The special case of $\alpha = 1$ is also known as the proportional-fair utility function.

The objective of the base station is to choose the power and the channel allocation so as to maximize the total utility of these K users over a horizon of T time slots. That is, the base station solves the optimization problem:

$$\text{maximize} \sum_{i=1}^{K} \left[U_i \left(\frac{1}{T} \sum_{t=1}^{T} r_i(t) \right) \right] \tag{OPT}$$

$$\text{subject to } [x_i(t)] \in \mathcal{S}, \forall t; \tag{3}$$

$$\frac{1}{T} \sum_t \sum_i p_i(t) \leq \bar{P}; \tag{4}$$

$$\sum_i p_i(t) \leq P_{max} \ \forall t. \tag{5}$$

Here $\bar{P}$ is the average transmit power budget available to the base station, and (4) is the average power constraint. This constraint also makes the problem different from that in [3] where there was no constraint on the average power.

Remark 1 (Short-term fairness). *A drawback of the utility function defined on the average rate is that if the scheduler knows that, for a particular user, the channel gain may be very high some time in the future then it might wait until this time to serve this user. This user may be starved of allocations in the short-term and the solution may be unfair to this user on short-time scales. One way to resolve the short-term unfairness is to introduce additional quality of service constraints such as requiring each user be scheduled at least once every given number of slots. This constraint was imposed in, for example, [7]. We can also include this constraint in the optimization problem. For simplicity, we do not impose it. We believe the results will be similar as long as this constraint is not very restrictive.*

Remark 2 (Fractional channel allocation). *In (OPT) we have allowed for fractional channel allocations. If the system imposes a binary constraint, that is only one user on one channel in any given slot, then these constraints can be imposed in OPT as well as in the algorithms we propose. In the experiments with a logarithmic utility function, we observed that allocations were mostly binary. So, we expect the qualitative conclusions will be valid whether allocations are binary or not.*
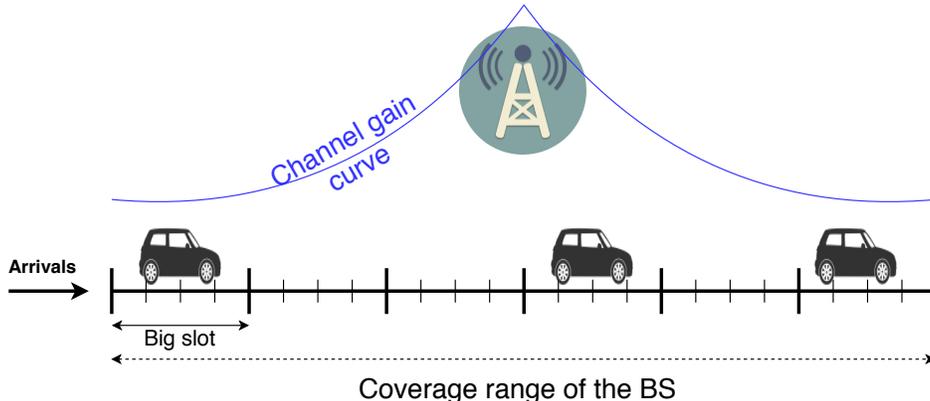
Figure 1: Mobility model.

Remark 3 (Maximum power constraint). *For conciseness, we shall not write the maximum power constraint explicitly in the optimization problems that we will define from now on. This constraint will be implicit and assumed to be applicable in all slots.*

The current literature mostly solves (OPT) when the base station is aware of only the channel gains in the current slot[1]. For vehicles sharing their itineraries, partial information on the future channel conditions could also be available to the decision maker in the current time-slot. We assume that this partial information is in the form of the mean of the channel gains in the future slots. One would expect to improve the value of the objective when this information is incorporated in the decision-making.

2.1 Channel gains

The heuristics will be evaluated on three different stochastic models for channel gains: *(i)* the stationary model in which, for each vehicle i, $\gamma_i(t)$ are independent and identically distributed across t with mean $\bar{\gamma}_i$; *(ii)* non-stationary and slowly varying model in which, for each vehicle i, the mean of the channel gains varies on a slow time-scale. For example, for a time-slot of 2 ms, the means may vary every 100 slot or 200 ms; and *(iii)* a mobility model on a stretch of road, where cars can come and leave as shown in Fig. 2.1. For this model $\gamma_i(t)$ are non-stationary, and also the mean of $\gamma_i(t)$ changes every slot.

Partial information on future channel conditions will be shown to be helpful for the second model and the third model. Nevertheless, the first (or the stationary) model will still be useful to illustrate the benefits of being able to vary the allocated power under an average power constraint.

3 Algorithms

In this section, we present the three algorithms that will be evaluated in the numerical experiments. The first one is the standard gradient-based scheduling which will be taken as the baseline. The other two are the ones we propose in this paper.

3.1 Locally optimal algorithm

The locally optimal algorithm is a special case of the one in [3] for the scenario in which there is only one sub-channel and the power budget is $\bar{P}$ in each slot. We remark that the algorithm can be handle the multiple sub-channel case but that in this paper our focus is on the single sub-channel case. When there

---

[1]We shall use slot and time-slot interchangeably to mean the decision making instants.

is no average-power constraint as in [3], this algorithm chooses maximizes the objective function computed without the knowledge of the future scheduling decisions. That is, in slot t, the scheduler solves

$$\text{maximize } \sum_{i=1}^{K}\left[U_i\left(\frac{1}{t}\sum_{t=1}^{t}r_i(t)\right)\right] \tag{LA}$$

$$\text{subject to } [x_i(t)]\in\mathcal{S},\forall t; \tag{6}$$

$$\sum_i p_i(t)\le\bar{P},\forall t. \tag{7}$$

Here, the past rate $r_i(s)$ for $s=1,2,...,t-1$ are known to the scheduler, so the decision variables are the channel allocations $x_i(t)$ and the transmit powers $p_i(t)$ in slot t. The maximum power constraint is set to $\bar{P}$ in order to make the comparison fair with algorithms which are subject to the average power constraint of $\bar{P}$. For the logarithmic utility function, we shall call this the Proportion Fair (PF) solution.

3.2 Short-term Objective 1 (STO1)

We begin by defining a big-slot to be B consecutive time-slots. The first heuristic we propose uses future mean channel gains to improve the objective. It includes in its decisions the future channel allocations and transmit powers but only on the scale of big-slots. In each time-slot, the allocations are recomputed for the current time-slot as well as the future big-slots. The intuition behind this is to maximize in each slot the long-term objective based on the best information available. We explain its workings in the context of the non-stationary slowly varying channel model and the mobility model.

For the slowly varying model, B will be the number of time-slots during which mean channel gain remains constant whereas in for the mobility model B can be set by the system designer depending on how fast the mean channel gains vary. Let $\hat{T}=T/B$ be the number of big-slots in the horizon, and let $\hat{\tau}_t\in\{1,\ldots,\hat{T}\}$ be the big-slot to which time-slot t belongs to, and let θ_t be the number of slots remaining in big-slot τ_t not including t, that is

$$\theta_t=(\tau_t+1)B-t.$$

We shall use the notation $\hat{x}$ for a quantity that is computed over a big-slot. For example $\hat{p}_i(\tau)$ will denote the power used in all the slots inside big-slot τ. Similarly,

$$\hat{r}_i(\tau)=B\hat{x}_i(\tau)\log\left(1+\frac{\bar{\gamma}_i(\tau)\hat{p}_i(\tau)}{\hat{x}_i(\tau)}\right) \tag{8}$$

is the total rate obtained by vehicle i in big-slot τ when it is served $\hat{x}_i(\tau)$ fraction of time at a transmit power of $\hat{p}_i(\tau)$. Note that here the rate is computed assuming that the channel gain is its mean value in big-slot τ. With slight abuse of notation,

$$\hat{r}_i(\tau_t)=\theta_t\hat{x}_i(\tau_t)\log\left(1+\frac{\bar{\gamma}_i(\tau_t)\hat{p}_i(\tau_t)}{\hat{x}_i(\tau_t)}\right), \tag{9}$$

shall denote the total rate in the remaining slots in current big-slot τ_t. Also, define

$$P_t=T\bar{P}-\sum_{i=1}^{K}\sum_{s=1}^{t-1}p_i(s)$$

to be the total remaining power available to the scheduler in slot t.

5

In each slot, STO1 maximizes

$$\sum_{i=1}^{K} U_i \left(\frac{1}{T} \left(\sum_{s=1}^{t-1} r_i(s) + r_i(t) + \sum_{\tau=\tau_t}^{\hat{T}} \hat{r}_i(\tau) \right) \right) \tag{STO1}$$

subject to

$$[x_i(t)] \in \mathcal{S}; \; [\hat{x}_i(\tau)] \in \mathcal{S}, \; \tau = \hat{\tau}_t, \dots, \hat{T}; \tag{10}$$

$$\sum_{i=1}^{K} \left(p_i(t) + \theta_t \hat{p}_i(\tau_t) + \sum_{\tau=\tau_t+1}^{\hat{T}} B \hat{p}_i(\tau) \right) \leq P_t. \tag{11}$$

The variables in this problem are $[p_i(t)]$ and $[\hat{p}_i(\tau)], \tau = \tau_t \dots \hat{T}$, and the corresponding channel allocations. In (11), the LHS is the total transmit power starting from the current slot which has to be less than the remaining power P_t.

For the mobile model, instead of solving (STO1) over the whole horizon T, we solve it on a shorter horizon, which is equal to the maximum staying time of the users currently inside the system. This shorter time horizon can vary from one slot to another, and it explains the words 'short-term objective' in the name of the algorithm. The advantages of this is to reduce the computation time which can be helpful when the algorithm has to be executed every 1ms.

3.3   Short-term Objective 2 (STO2)

The STO1 algorithm recomputes in each time slot the optimal solution of the future big-slots. In STO2, we recompute the solution of the future big-slots only at the beginning of each big-slot. Inside a big-slot, we compute only the solution for the current slot assuming the solution for the future big-slots to be the same as that computed at the the start of the current big-slot. That is if $t \equiv 1 \pmod{B}$, STO2 first maximizes

$$\sum_{i=1}^{K} U_i \left(\frac{1}{T} \left(\sum_{s=1}^{t-1} r_i(s) + \sum_{\tau=\tau_t}^{\hat{T}} \hat{r}_i(\tau) \right) \right) \tag{STO2-Big}$$

subject to

$$[\hat{x}_i(\tau)] \in \mathcal{S}, \; \tau = \hat{\tau}_t, \dots, \hat{T} \tag{12}$$

$$B \sum_{i=1}^{K} \sum_{\tau=\tau_t}^{\hat{T}} \hat{p}_i(\tau) \leq P_t. \tag{13}$$

The variables in this problem are $[\hat{x}_i(\tau)]$ and $[\hat{p}_i(\tau)], \tau = \tau_t \dots \hat{T}$.

Next, in each slot t, we compute the optimal allocation and transmit power assuming that the allocations and transmit powers in the future big-slots are those computed from solving (STO2-Big). In slot t, STO2 maximizes

$$\sum_{i=1}^{K} U_i \left(\frac{1}{T} \left(\sum_{s=1}^{t-1} r_i(s) + r_i(t) + \sum_{\tau=\tau_t}^{\hat{T}} \hat{r}_i(\tau) \right) \right) \tag{STO2-Small}$$

subject to

$$[x_i(t)] \in \mathcal{S}, \; [\hat{x}_i(\tau_t)] \in \mathcal{S}; \tag{14}$$

$$\sum_{i=1}^{K} (p_i(t) + \theta_t \hat{p}_i(\tau_t)) \leq P_t - \sum_{i=1}^{K} \sum_{\tau=\tau_t+1}^{\hat{T}} B \hat{p}_i(\tau). \tag{15}$$

The variables in this problem are $[x_i(t)]$, $[p_i(t)]$, $[\hat{x}_i(\tau_t)]$ and $[\hat{p}_i(\tau_t)]$. As can be seen, in a slot the dimension of the problem is no bigger than the one for STO1. As in STO1, for the mobile model, STO2 solves (STO2-Big) and (STO2-Small) over a shorter horizion which is the maximum staying time of the users currently in the coverage range. The pseudo code for STO2 is shown in Algorithm 1.

Algorithm 1 The STO2 algorithm

$t \leftarrow 1$
while $t \leq T$ **do**
    if $t \equiv 0 \mod B$ **then**
        Solve (STO2-Big) and obtain $\hat{\mathbf{x}}(\tau)$ and $\hat{\mathbf{p}}(\tau)$.
        Solve (STO2-Small)
    else
        Solve (STO2-Small)
    end if
end while

In terms of computational effort, compared to STO1, STO2 solves a lower dimensional problem in each slot except at the starting of every big-slot where the dimension is same as for STO1. Thus, one can except it to be faster but further from the optimal solution.

4 Numerical experiments

The numerical experiments were run in Python, and all optimization problems were solved using the python package CVXPY [18] and the solver MOSEK. The results will be presented according to the channel-gain models. For the stationary and slowly varying channel models, we assume that the number of users is fixed. The third model will be evaluated in a dynamic setting in which users will arrive and leave the network.

In all the experiments, $\bar{P}$ is set to 15 and P_{max} is set to 30. Whenever we show any performance measure of the optimal solution, it will be assumed to mean that the optimal is computed assuming all the future channel gains are known exactly.

4.1 Stationary channel

For the first experiment, we take $K = 4$, that is four users, and a logarithmic utility function for every user. The vector of means $[\bar{\gamma}_i] = [6.76, 5.45, 4.35, 1.31]$. The channel gain in slot t for user i is generated as follows:

$$\gamma_i(t) = \bar{\gamma}_i A_i(t; \eta) \tag{16}$$

Here $A_i(t; \eta)$ is a sequence of i.i.d. uniform random variable in the range $[1-\eta, 1+\eta]$. We shall refer to η as the *noise level*. It is assumed that $A_i(t)$ and $A_j(t)$ are assumed to be independent for $i \neq j$. Varying η from 0 to 1 changes the variance of $A_i(t)$ from low to high. For $\eta = 0$, the channel gains become deterministic and known to the decision maker.

Remark 4. *The method for generating $\gamma_i(t)$ need not be necessarily multiplicative as in* (16). *Our heuristics can be used as long as the means of the future channel gains are known. In this paper, we limit the numerical evaluation to the form in* (16).

Figure 4.1 shows the total utility as a function of the noise level η. The time horizon T was taken to be 500 with 5 big-slots, that is, one big-slot has $B = 100$. If the scheduling slot is 1 ms as in 4G [19], then the scheduling horizon is of 500 ms. Five sample paths for channel gains were generated, and the plot shows the average of these 5 samples. The label PF is for the local optimization algorithm. The suffix FP attached to STO1 means that STO1 was run wilth a fixed power budget of $\bar{P}$ in each slot.
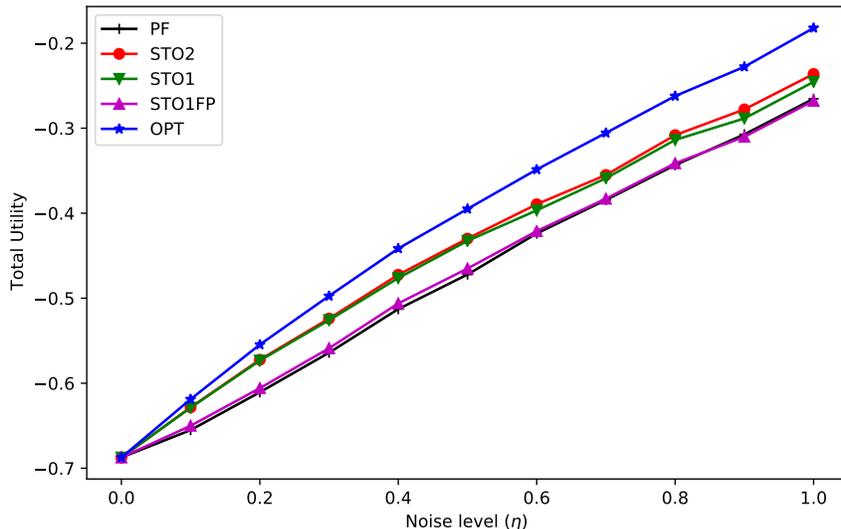
Figure 2: Total utility as a function of noise level. Channel is stationary. Log utility function.

As expected, allowing for an average transmit power constraint and using future information (even if it is just the mean channel gains) improves the utility. When $\eta = 0$, all the algorithms are equivalent and give the same utility since the channel gain is the same is every time slot and is known. Further, performance improvement is more when the noise variance is higher which is again to be expected.

4.2 Slowly varying channel

Next, we conduct experiments with the slowly varying channel model. It is assumed that the channel means are constant during $B = 100$ slots. The optimization horizon is $T = 2000$, that is there are 20 big-slots in each run. The mean channel gains were first determined for each big-slot. Within a big-slot, the channel gains were then generated using the same method as for the stationary case and given in (16). The number of users was again set to 4.

In the first experiment, the mean channel gains are relatively homogeneous with an empirical average of the mean channel gains being $[5.98, 5..55, 4.69, 5.30]$. Figure 3a, plots the total utility as a function of the noise level while Fig. 3b shows the total transmit power in a slot as a function of time-slot. The data in the latter plot was obtained on a separate run with $B = 50$, $\eta = 1$, and only STO2 and OPT are shown so as to have a more readable figure.

We observe that all algorithms except PF are close to optimal almost throughout. Since the total transmit power in a slot is not far from $\bar{P}$, STO1FP is almost as good as STO1. However, in this scenario prediction is still useful as PF is away from OPT even for $\eta = 0$, that is when there is no noise.

The mean channel gains in the second experiment are widely varying with the empirical average of the mean channel gains being $[2.63, 10.3, 3.76, 0.009]$. One user is in a very bad channel state, whereas another one has much better mean channel gains than the others. The plots for the total utility and total transmit power in a slot are shown in Fig. 4.2. Again, the total power was computed from a separate run.

This time we observe that both power control and channel prediction result in improvements. Due to one user being much worse than the others, the optimal transmit power varies much more than in the previous experiment and hits the maximum constraint quite often. Since the fixed-power version is inflexible in this respect, it performs worse for all noise levels.
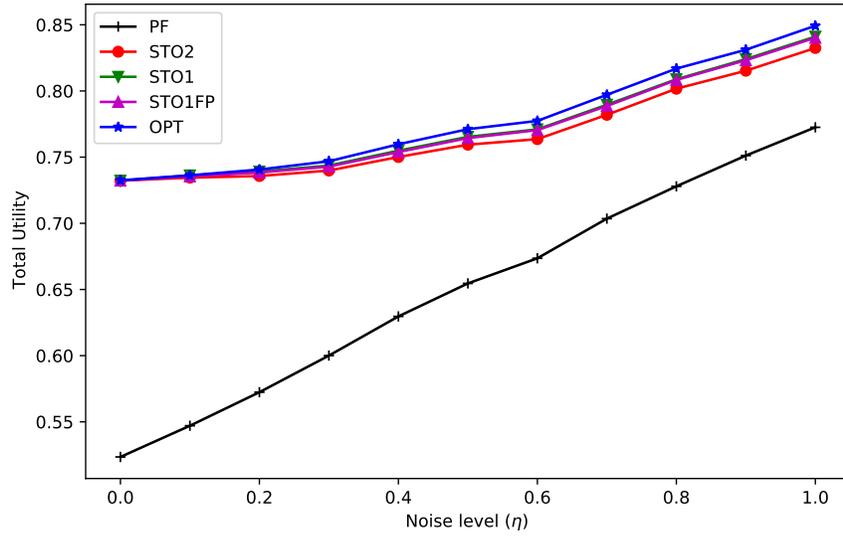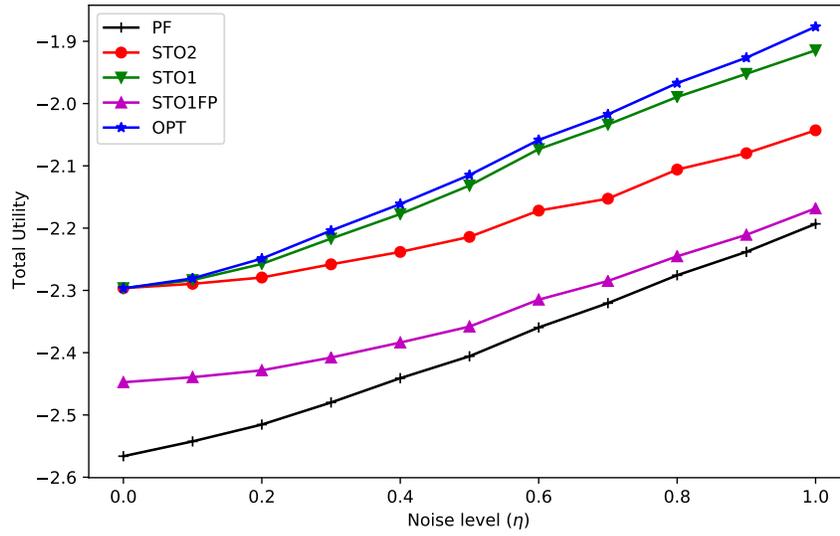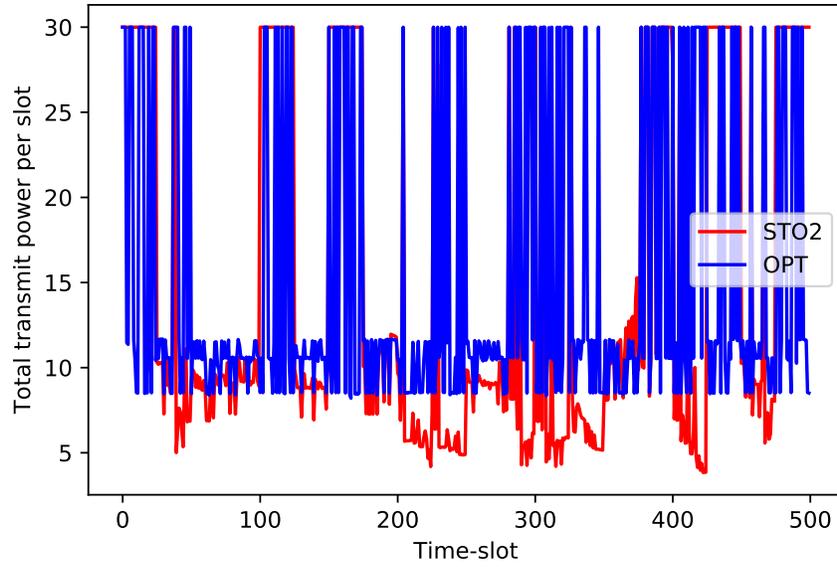
(a) Total utility vs. noise level
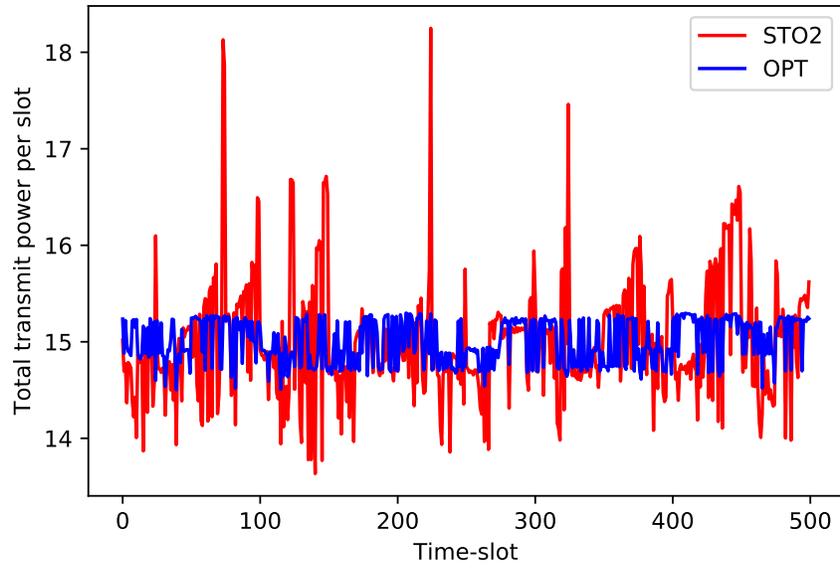
(b) Total allocated power per slot vs. time-slot

Figure 3: Slowly varying channel. Log utility function. Experiment 1. $\eta = 1$.

(a) Total utility vs. noise level

(b) Total allocated power per slot vs. time-slot

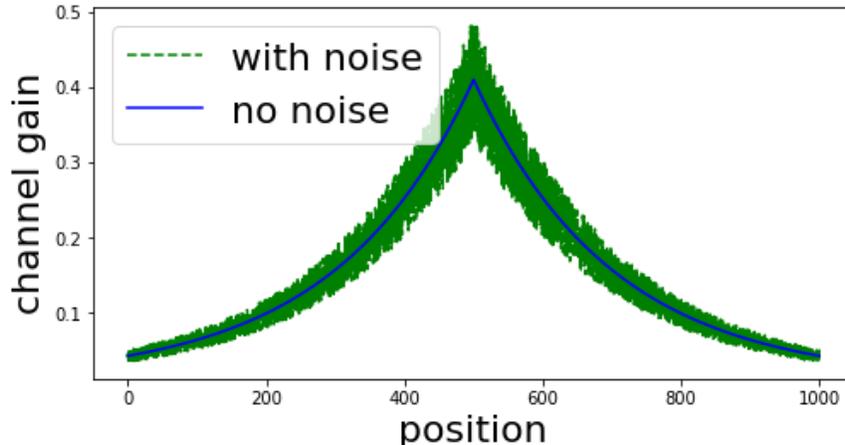Figure 4: Slowly varying channel. Log utility function. Experiment 2. $\eta = 1$.

Figure 5: Channel gain for two cases: with noise ($\eta = 0.2$) and without noise.

4.3 Mobility Model

Consider a stretch of road of length 1 km covered by one base station, in which vehicles enter from the left and leave on the right (Fig. 2.1). To simplify for illustration, we assume they move with same velocity $v = 25$ m/s, but the algorithms presented in Sec. 3 do not depend on this assumption. So they stay in the coverage range of the base station in 40 seconds. New vehicles can enter the network only at the start of the big-slots. The length of a big-slot is set to 1 sec. The probability of a new arrival in a big-slot is set to $p = 0.3$. For each value of noise level (η) we simulate this network for 800 seconds. In this model, the channel gain of mobile user is non-stationary and varies every time slot (1 ms), STO1 and STO2 solves over an horizon of 40 seconds which is the staying time of users in the network. These two aspects make this model different from the two first models. Since the dimension of (OPT) is very large in this case, we do not compute the optimal solution in the experiments. Also, since STO1 solves a high dimensional problem in every slot compared to STO2, it is much slower when the horizon is large. So, we show only the performance of STO2 and PF.

Fig. 4.3 illustrates the channel gain curve in the noise and no noise cases. Here, the rate is a function of users' position which is in fact a function of the distance to the base station. Let us take the left margin of the road be 0, and the right margin be 1000, then for position $x \in [0, 1000]$ inside the coverage range, the channel gain is equal to $f(x) = \beta(1 + \kappa \exp(|500 - x|/\sigma)$. Here, β, κ, σ are adjustable parameters; in our experiment $\beta = 0.01, \kappa = 40, \sigma = 200$.

Fig. 6 illustrates the numerical results for the above mobile system. The total allocated power is shown only for a small interval of time. Again, we observe that channel prediction leads to a better total utility. Further, by observation, STO2 allocates more power when there are many users (than usual) close to the peak where channel gains are good with high probability.

5 Conclusions and future work

We proposed two heuristics for joint power control and channel allocation that exploit partial information future channel conditions to improve the utility. Even little information such as mean channel gains is sufficient to observe improvement compared to when no information is used.

This preliminary work opens several directions for future work. First, these algorithms can be generalized to multiple base-station networks. Second, the robustness of these algorithms with respect to errors in future information will be worth investigating. Third, dual-decomposition methods could be applied to get
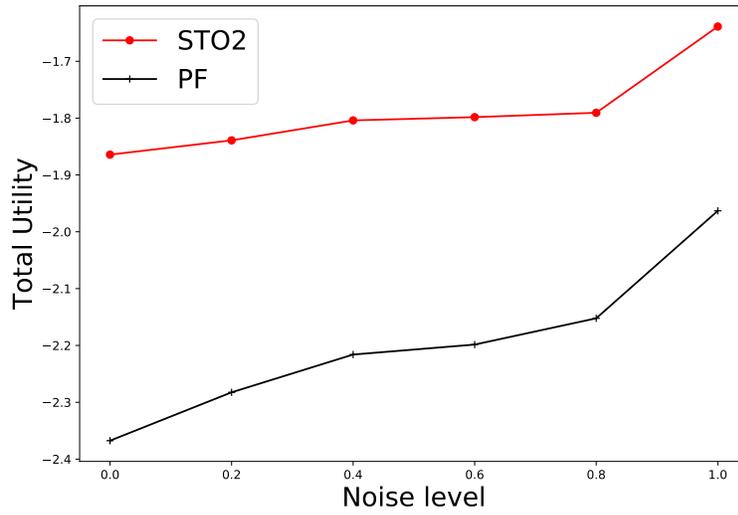
(a) Total utility vs. noise level
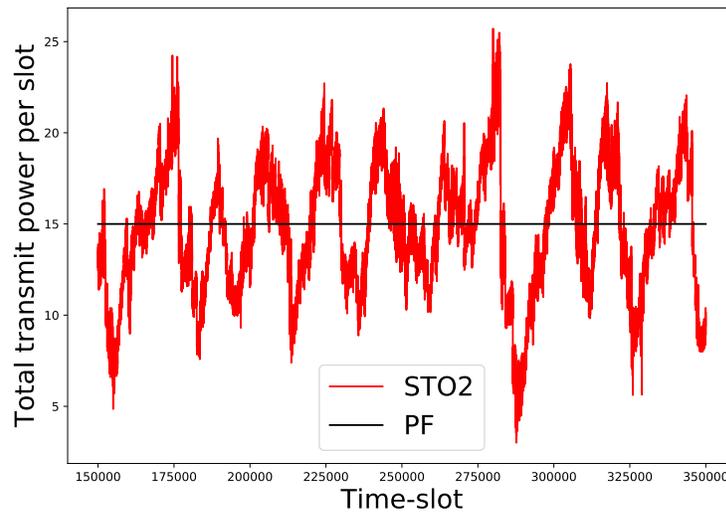
(b) Total allocated power vs. time (in second). $\eta = 0.2$.

Figure 6: Mobility model. Log utility function.

a STO2-like algorithm but with a penalty term instead of the average power constraint. For this relaxed problem, structural results similar to that for gradient-based scheduling could be hoped for. Finally, these scheduling problem can also be studied on the uplink, and similar algorithms can be investigated for the opportunistic utility model.

Acknowledgment

References

[1] J. Lee, R. R. Mazumdar, and N. B. Shroff, "Opportunistic power scheduling for dynamic multi-server wireless systems," *IEEE Trans. Wireless Communications*, vol. 5, no. 6, pp. 1506–1515, 2006.

[2] M. Chiang, P. Hande, T. Lan, and C. W. Tan, "Power control in wireless cellular networks," *Foundations and Trendső in Networking*, vol. 2, no. 4, pp. 381–533, 2008.

[3] J. Huang, V. Subramanian, R. Agrawal, and R. Berry, "Downlink scheduling and resource allocation for ofdm systems," *IEEE Transactions on Wireless Communications*, vol. 8, pp. 288–296, 01 2009.

[4] C. Li and M. J. Neely, "Network utility maximization over partially observable markovian channels," *Perform. Evaluation*, vol. 70, no. 7-8, pp. 528–548, 2013.

[5] V. G. Subramanian, R. A. Berry, and R. Agrawal, "Joint scheduling and resource allocation in CDMA systems," *IEEE Trans. Inf. Theory*, vol. 56, no. 5, pp. 2416–2432, 2010.

[6] H. J. Kushner and P. A. Whiting, "Convergence of proportional-fair sharing algorithms under general conditions," *IEEE Transactions on Wireless Communications*, vol. 3, no. 4, pp. 1250–1259, July 2004.

[7] R. Margolies, A. Sridharan, V. Aggarwal, R. Jana, N. K. Shankaranarayanan, V. A. Vaishampayan, and G. Zussman, "Exploiting mobility in proportional fair cellular scheduling: Measurements and algorithms," *IEEE/ACM Trans. Netw.*, vol. 24, no. 1, pp. 355–367, Feb. 2016.

[8] N. Nguyen, O. Brun, and B. Prabhu, "An algorithm for improved proportional-fair utility for vehicular users," *The 25th International Conference on Analytical and Stochastic Modelling Techniques and Applications ASMTA-2019*, October 2019.

[9] S. Borst, N. Hegde, and A. Proutiere, "Mobility-driven scheduling in wireless networks," in *IEEE INFOCOM 2009*, April 2009, pp. 1260–1268.

[10] L. Tan, Z. Zhu, F. Ge, and N. Xiong, "Utility maximization resource allocation in wireless networks: Methods and algorithms," *IEEE Transactions on Systems, Man, and Cybernetics: Systems*, vol. 45, no. 7, pp. 1018–1034, July 2015.

[11] P. Chandur, R. M. Karthik, and K. M. Sivalingam, "Performance evaluation of scheduling algorithms for mobile wimax networks," in *2012 IEEE International Conference on Pervasive Computing and Communications Workshops*, March 2012, pp. 764–769.

[12] S. Borst, "User-level performance of channel-aware scheduling algorithms in wireless data networks," *IEEE/ACM Transactions on Networking*, vol. 13, no. 3, pp. 636–647, June 2005.

[13] H. Abou-zeid, H. S. Hassanein, and N. Zorba, "Long-term fairness in multi-cell networks using rate predictions," in *2013 7th IEEE GCC Conference and Exhibition (GCC)*, Nov 2013, pp. 131–135.

[14] H. J. Bang, T. Ekman, and D. Gesbert, "Channel predictive proportional fair scheduling," *IEEE Transactions on Wireless Communications*, vol. 7, no. 2, pp. 482–487, February 2008.

[15] S. Xu, G. Zhu, C. Shen, Y. Lei, and Z. Zhong, "Utility-based resource allocation in high-speed railway wireless networks," *EURASIP Journal on Wireless Communications and Networking*, vol. 2014, no. 1, p. 68, Apr 2014.

[16] A. Mesbah, "Stochastic model predictive control: An overview and perspectives for future research," *IEEE Control Systems Magazine*, vol. 36, no. 6, pp. 30–44, Dec 2016.

[17] J. Mo and J. Walrand, "Fair end-to-end window-based congestion control," *IEEE/ACM Transactions on Networking*, vol. 8, no. 5, pp. 556–567, Oct 2000.

[18] S. Diamond and S. Boyd, "CVXPY: A Python-embedded modeling language for convex optimization," *Journal of Machine Learning Research*, vol. 17, no. 83, pp. 1–5, 2016.

[19] E. Dahlman, S. Parkvall, and J. Sköld, *4G LTE-Advanced Pro and The Road to 5G (Third Edition)*, third edition ed. Academic Press, 2016.