

We formulate and study a multi-user multi-armed bandit problem that exploits the temporal-spatial opportunistic spectrum access (OSA) of primary user (PU) channels, so that secondary users (SUs) who do not interfere with each other can make use of the same PU channel. We first propose a centralized channel allocation policy that has logarithmic regret, but requires a central processor to solve an NP-complete optimization problem at exponentially increasing time intervals. To overcome the high computation complexity at the central processor, we also propose heuristic distributed policies that, however, have linear regrets. Our first distributed policy utilizes a distributed graph coloring and consensus algorithm to determine SUs' channel access ranks, while our second distributed policy incorporates channel access rank learning in a local procedure at each SU at the cost of a higher regret. We compare the performance of our proposed policies with other distributed policies recently proposed for temporal (but not spatial) OSA. We show that all these policies have linear regrets in our temporal-spatial OSA framework. Simulations suggest that our proposed policies have significantly smaller regrets than the other policies when spectrum temporal-spatial reuse is allowed.