中国科学院数学与系统科学研究院
Academy of Mathematics and Systems Science, CAS

运筹学与信息科学研究室

Department of Operations Research and Information Science

学术报告

题　　目： **Robust Deterministic Policies for Markov Decision Processes under Budgeted Uncertainty**

报告人： 吴菲，比利时鲁汶大学

时　　间： 10 月 31 日（星期四） 14:00 – 15:00

地　　点： 数学院思源楼 S615

摘　　要： This paper studies the computation of robust deterministic policies for Markov Decision Processes (MDPs) in the Lightning Does Not Strike Twice (LDST) model of Mannor, Mebel, and Xu (2012). In this model, designed to provide robustness in the face of uncertain input data while not being overly conservative, transition probabilities and rewards are uncertain and the uncertainty set is constrained by a budget that limits the number of states whose parameters can deviate from their nominal values.

Mannor et al. (2012) showed that optimal randomized policies for MDPs in the LDST regime can be efficiently computed when only the rewards are affected by uncertainty. In contrast to these findings, we observe that the computation of optimal deterministic policies is NP-hard even when only a single terminal reward may deviate from its nominal value and the MDP consists of 2 time periods. For this hard special case, we then derive a constant-factor approximation algorithm by combining two relaxations based on the *Knapsack Cover* and *Generalized Assignment* problem, respectively. For the general problem with possibly large number of deviations and a longer time horizon, we derive strong inapproximability results for computing robust deterministic policies as well as Σ_2^p-hardness, indicating that the general problem does not even admit a compact mixed integer programming formulation.

报告人简介： 吴菲是比利时鲁汶大学运筹学研究中心的博士生，研究方向和兴趣兴趣包括组合优化、算法设计与分析、鲁棒优化，以及在调度中的应用等。