

Online Progressive Recovery of Interdependent Networks

Siddartha Devic, Genya Ishigaki, Prof. Jason Jue

Department of Computer Science, The University of Texas at Dallas



Abstract

The progressive recovery problem is an important problem in physical and communication network recovery where the network service provider wishes to maximize the total interim recovery utility after a catastrophic simultaneous failure of multiple network nodes. Ishigaki et al. study the static progressive recovery problem in [1]. We extend the problem to the online setting, with multiple problem instances arriving in a sequence. We attempt to derive a convex relaxation of the problem in order to apply regret bounds from the rich online convex optimization literature. Lastly, we discuss the difficulty in providing convex relaxations for complex time-dependent combinatorial optimization problems.

The Progressive Recovery Problem

The progressive recovery problem concerns the incremental recovery of physical network nodes—represented by a graph—under resource constraints. The goal of the problem is to maximize the sum of the utilities of functional nodes over all time steps (Fig. 1). Utility (u) can take the form of computation power or number of people serviced by that node. The resource constraint r represents manpower or virtual resources available at each time step. Finally, the demand of a node d represents the amount of resources r which must be allocated to it for full recovery.

The progressive recovery problem as defined in [1] also has an additional constraint. A node can only be brought online if there exists a path from an independent node (a in Fig. 1). Due to this additional constraint, not all recovery orders are optimal. For example, Plan II in Fig. 1 is a worse recovery order than Plan I since we waste resources recovering node d before node b . Since there is no path from $a \rightarrow d$ without first recovering d , we gain no utility in the first time step $t = 1$.

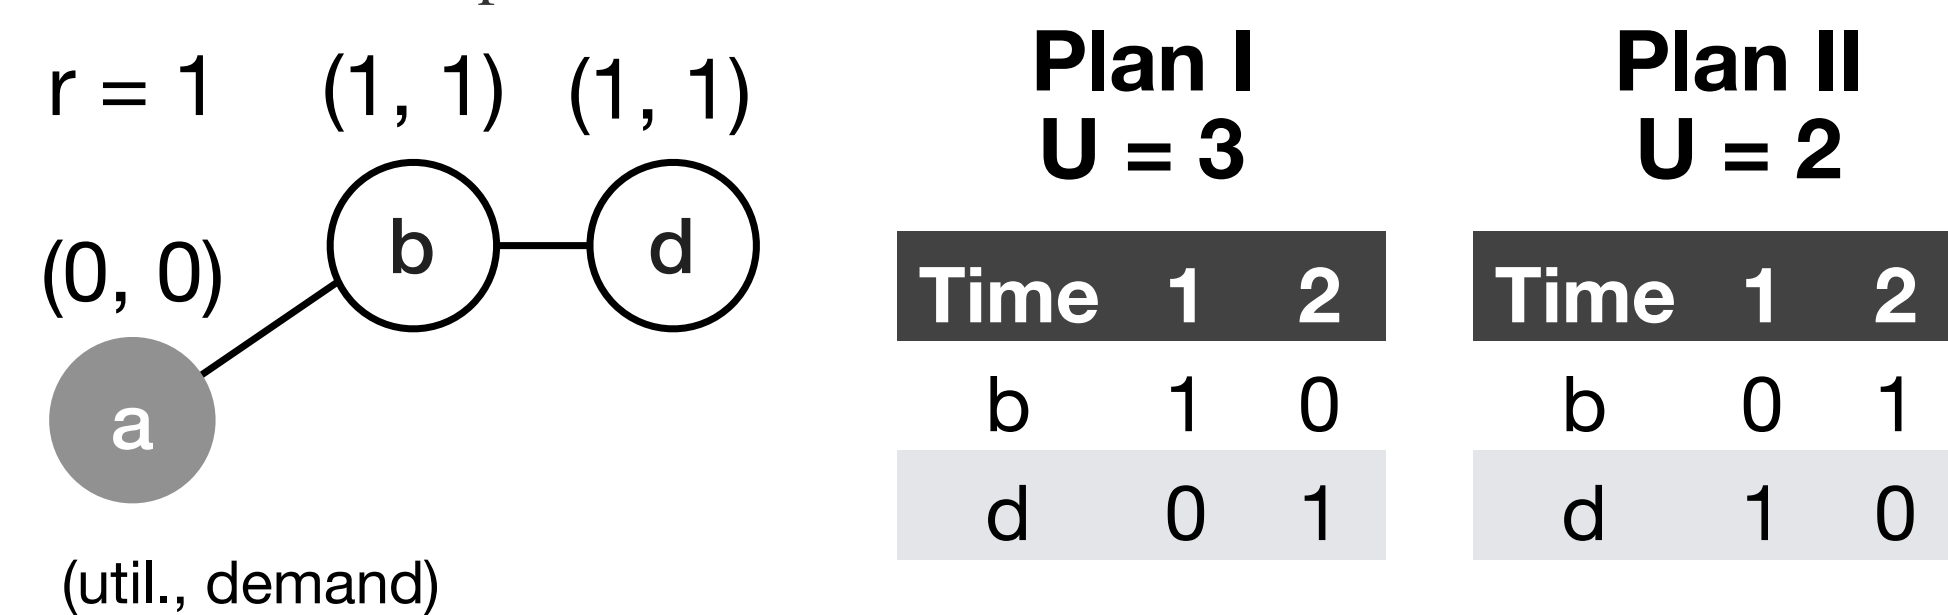


Figure 1: A simple progressive recovery problem instance. (Left) Each node in the recovery graph has (u,d) attributes, representing utility that node provides and the demand required to recover it. (Right) After a given node’s demand is satisfied, that node provides (util) resources at each subsequent time step.

Formally, the progressive recovery problem requires us to find a recovery plan P represented by a (node \times time step)-matrix that maximizes the sum of

interim utility provided by functional nodes during the recovery.

Problem 1. Progressive Recovery Problem (PR): Given a graph $N = (V = V_0 \cup V_1, A = A_{01} \cup A_{10} \cup E_{00} \cup E_{11})$, a demand vector $\mathbf{d} \in \mathbb{R}^N$, an utility vector $\mathbf{u} \in \mathbb{R}^N$, an initial functionality vector $\mathbf{F}_0 \in \mathbb{R}^N$, and a resource constant $r \in \mathbb{R}$, maximize the network-wide utility $U(P) = \sum_{s=0}^S \mathbf{u}^\top \mathbf{F}_s$ by deciding a resource assignment matrix $P \in \mathcal{P}$ subject to $\sum_{v \in V} P[v][s] = r \ (\forall s)$.

Online Learning

Online optimization problems are a class of problems where a decision maker must learn from problem instances arriving in sequence, minimizing the loss between the online and hindsight decisions. For example, in the online spam email classification problem, there is an agent who receives emails in a sequence and must decide to let an email into the inbox or redirect it to the spam folder. The agent is told whether it made the right decision or not after classifying each email. The problem is considered “online” because the agent must begin learning from scratch, i.e. it has no *a priori* knowledge on what constitutes spam emails before the sequence begins.

Online Learning for Progressive Recovery

The *online* progressive recovery problem concerns multiple instances of the static progressive recovery problem in sequence. For example, if the static progressive recovery problem dealt with a single earthquake, the online problem deals with a sequence of earthquakes and learns how to best recover the network after each subsequent failure.

The online progressive recovery problem is to identify an optimal recovery plan P over different failure scenarios that are characterized by different $(\mathbf{d}, \mathbf{F}_0)$ pairs.

Problem 2. Online Progressive Recovery Problem (Online PR): Assuming a system experiences multiple independent failure events in sequence $(\mathbf{d}_t, \mathbf{F}_0^t)$ ($t = 1, 2, \dots, T$), which could be adversarial scenarios, find a recovery plan P that minimizes the regret:

$$R_T \triangleq \sup_{U_1, \dots, U_T} \left\{ \sum_{t=1}^T U_t(P_t) - \min_{P \in \mathbb{R}^{N \times S}} \sum_{t=1}^T U_t(P) \right\}, \quad (1)$$

where P_t is the recovery plan chosen for a scenario t .

Note that the supremum function over U_t sets the worst-case regrets, insisting that our strategy is robust against adversaries.

Online Convex Optimization

Convex optimization is a subfield of computer science which deals with finding global minima of convex functions: functions whose tangents lie below

the function everywhere. By transforming online progressive recovery to the online convex optimization (OCO) framework, we can derive bounds on the worst case of the solution we find. With OCO, we can account for the worst case scenarios of adversarial attacks and achieve “bounded regret”. Regret is defined as the worse-case difference between the online solution and the optimal solution with hindsight knowledge.

Deriving a Convex Relaxation

Importantly, to apply techniques from OCO, we must have a convex objective U_t . A sketch of our approach is as follows.

- We restrict the action space to an n -dimensional (convex) hypercube
- Instead of directly incorporating the constraints into the optimization problem, we add penalty terms to the objective for violating certain constraints (such as invalid/over-allocation of resources)
- Working in this “dual” space, we need only prove convexity of each penalty function in the objective

We are currently working on providing a full convex relaxation of the progressive recovery problem. This poses a unique challenge since the problem is combinatorial in nature and also requires a recursive evaluation within the objective function. This differs from existing work in the convex optimization literature which typically involves only the evaluation of a single convex objective. Nonetheless, we also derive helpful techniques that may apply more generally within the OCO framework, and hope to extend our results to other problems within optical networking.

Future Work

We plan to continue theoretical work regarding convexity of the online progressive recovery problem. Currently, we are facing an issue with regards to an interdependence of constraints within the problem. We plan to fix this issue and implement practical simulations showcasing the convergence bounds of our method. Furthermore, using the general intuition and techniques discovered while working on this problem, we hope to show that OCO style proofs may be able to provide useful theoretical bounds for other common problems in optical networking, such as routing or defragmentation.

References

- [1] Genya Ishigaki, Siddartha Devic, Riti Gour, and Jason P. Jue. DeepPR: Recovery of interdependent virtualized networks with incremental resources by deep rl. In *IEEE Journal on Selected Areas in Communications - Advances in AI and ML for Networking*, 2020 (Accepted), 2019.