Dynamic Facility Location : Minimizing Sum of Radii

Nicolas Blanchard, Nicolas Schabanel

November 4, 2015

Submission to the fourth GT CoA sessions

In the past few years the increased study of social networks of different sorts has given rise to large databases with not just static but also dynamic information. The interactions of individuals in a hospital, for example, are strongly dependent on time, and clustering techniques can help understand the dynamics of those networks. A possible modelization for this is via the Facility Location problem. This model was first studied in the context of a warehouse locating situation, with a heuristic first published in 1963. Since then it has become an important subject of research covered by hundreds of publications. Many different subcases exist depending on whether we have an evolution of the network or not (dynamic or static), on how exactly the cost function behaves, on the structure of the underlying space (euclidean, metric, non-metric), as well as other parameters. Most of those can be proved to be **NP**-complete so the main goal is finding approximation algorithms. Several constant factor approximation algorithms already exist in different settings.

The Facility Location problem consists of assigning a set of n clients $\mathcal{C}$ to a set of m services $\mathcal{F}$ so that every client is served. Although the problem is **NP**-complete and **APX**-hard, some optimal bounds have been found : $O(\log n)$ in the general (non-metric) case and $O(1)$ in the metric case (where distances follow the triangle inequality). As opposed to the static model which has been extensively studied, the dynamic model is still partially unexplored – this model has the added complexity of having distances which vary through time as well as additional costs when clients change facilities. We introduce a variant of the classical setting, where the connection cost is the sum of the radii of facilities instead of the sum of the distances between clients and facilities (the radius is the distance between a facility and the furthest client assigned to it). This second variant is a natural expansion of the problem and had received less attention until now.

We present four results : first, we proved that unless $\mathbf{P} = \mathbf{NP}$, the Dynamic Facility Location with Radius (DFLR) problem in the genera non-metric case cannot be approximated within a factor $(1 - o(1)) \log n$. Second, we designed and proved an algorithm working in expected polynomial time ($O(nm(n+m)^{1/2})$) which gives a $2 \log n$ approximation for DFLR in the non-metric case, and show that this method can be adapted to an algorithm from [Eisenstat, Mathieu, Schabanel 2014] to obtain an $O(\log n)$ approximation for non-metric DFL without radii (down from their original $O(\log nT)$). For the sum of distances between facilities and clients the algorithm present in [An, Norouzi, Svensson 2014] achieves a constant approximation in the metric case. We show that the natural adaptation of this algorithm to the sum of radii has an approximation ratio at least $\Omega(\log \log n)$. Finally, we base our proof on a combinatorial lemma dealing with permutations in trees which is of intrinsic interest.