Proportional Participatory Budgeting

Grzegorz Pierczyński

Nowadays, a growing number of cities allocate part of their budget funds to projects indicated directly by citizens within *participatory budgeting*. The voting process is simple: usually, voters can choose any number of projects they support. Sometimes they additionally assign them e.g. point weights or the order of their preferences in the ranking. However, the maximum amount of funds that we can spend on projects is limited, and the implementation of each of them is associated with a specific cost. How should we decide which projects to choose?

Imagine a situation in which the costs of all projects are equal, and the budget limit allows us to implement 10 of them. 40% of voters support 10 urban greening projects, another 30% of voters support 10 new playground projects, and the remaining 30%—10 bicycle infrastructure projects. In such a case, the solution that would maximize total voter satisfaction would be to choose 10 urban greening projects—however, such a solution would be unfair, satisfying only 40% of voters and completely ignoring 60% of them. Instead, we very often aim to *proportionality*—we expect each x% group of voters with consistent preferences to have the right to decide how to spend x% of the budget. In this sense, we can say that e.g. in the above situation, the only fair solution would be to choose 4 projects from the first group and 3 from each of the other two groups.

However, such clear and unambiguous cases as described above rarely occur in practice—most often, it is not possible to divide voters and projects into disjoint groups. Moreover, the problem is further complicated by unequal project costs and unequal voter support for them, as well as logical dependencies between the projects—e.g. the implementation of some of them may be mutually exclusive (because e.g. several projects are planned in the same location), or on the contrary—the selection of some project requires funding another one.

In such situations, defining what proportionality is, developing proportional algorithms for selecting projects and analyzing their computational complexity becomes an interesting and challenging problem from the perspective of computer science and mathematics. In addition, it turns out that our considerations find practical applications also in areas completely unrelated to the organization of society—e.g. in determining what results the web browser should return in response to a query, in blockchain consensus protocols, or in improving selection procedures for genetic algorithms.

Unfortunately, our current knowledge on the subject is quite limited—in fact, the most important theoretical results have so far only been achieved for specific cases of this problem (e.g. equal project costs, no logical dependencies between projects, the simple binary form of voters' preferences). The goal of our project is to formulate and analyse new mathematical definitions (axioms) and algorithms for its more general variants. For instance, we will propose two completely new approaches to understanding fairness, based on the concepts of (1) market equilibria known from theoretical economics (i.e. the Lindahl equilibrium for public goods), (2) Condorcet consistency, an important criterion for achieving a group consensus in a situation of when we only have one decision to make. We also plan to design proportional algorithms for a model with general logical dependencies between projects—our research in this area will be a pioneering contribution to the literature—as well as to develop theoretical tools for the statistical analysis of the proportionality of algorithms, both on real and randomly generated data. Our preliminary results are promising in all of these topics.