Balancing Spreads of Influence in a Social Network

Ruben Becker, Federico Corò, Gianlorenzo D'Angelo, Hugo Gilbert

Gran Sasso Science Institute, L'Aquila, 67100, Italy

{firstname.lastname}@gssi.it

Keywords : influence maximization, diversity maximization, approximation algorithms.

1 Introduction

One of the promises of a highly connected world is that of an impartial spread of opinions driven by free and unbiased sources of information leading to an equitable exposure of opinions to the wide public. On the contrary, the social network platforms that are currently governing news diffusion, while offering many seemingly-desirable features like searching, personalization, and recommendation, are reinforcing the centralization of information spreading and the creation of what is often termed *echo chambers* and *filter bubbles* [7].

Consequently, instead of giving users a diverse perspective and balancing users opinions by exposing them to challenging diverse ideas, social media platforms are likely to make users more extreme by only exposing them to views that reinforce their pre-existing beliefs [4, 5].

To address this issue from an algorithmic perspective, Garimella et al. [6] introduced the problem of *balancing information exposure* in a social network. Following the *influence maximization paradigm* going back to the seminal work of Kempe, Kleinberg and Tardos [8, 9], their problem involves two opposing viewpoints or campaigns that propagate in a social network following the *independent cascade model*. Given initial seed sets for both campaigns, they consider the optimization problem of selecting at most k additional seed nodes for both campaigns in order to maximize the expected number of nodes that are reached by either both or none of the campaigns. The authors studied two different settings, namely the *heterogeneous* and *correlated* settings. The heterogeneous setting corresponds to the general case in which there is no restriction on the probabilities with which the campaigns spread. Contrarily, in the correlated setting, the probability distributions for different campaigns are identical and completely correlated. After proving the *NP*-hardness of balancing information exposure, the authors designed efficient approximation algorithms with an approximation ratio of $(1-1/e-\epsilon)/2$ for any $\epsilon > 0$ for both the correlated and heterogeneous settings.

2 Our Contribution

This communication aims to present a natural generalization of this optimization problem in which μ different campaigns propagate in the network and we aim to maximize the expected number of nodes that are reached by at least ν or none of the campaigns, where $\mu \geq \nu \geq 2$. We term this problem the μ-ν-Balance problem. This generalization is motivated by the fact that for most problems, not only two, but a multitude of viewpoints are perceivable. Hence, given this possibly large number μ of viewpoints, the ν threshold parameter aims to guarantee that influenced users are exposed to a sufficiently large subset of viewpoints, hopefully providing them with a more representative picture of the problem at hand.

Interestingly, we will present results that surprisingly differ from the ones Garimella et al. [6] obtained for the special case where $\mu = \nu = 2$. Indeed, in the heterogeneous setting, one can obtain a reduction which shows that the μ-ν-Balance problem is as hard to approximate as the Densest-k-Sub-d-hypergraph problem [3]. This finding easily leads to strong approximation hardness results. In particular, when $\nu \geq 3$, then under the gap exponential time hypothesis [10], there is no $n^{-g(n)}$-approximation algorithm with $g(n) = o(1)$ for the μ-ν-Balance problem, where n is the number of nodes. Moreover, when $\nu \geq 4$, then, if a certain class of one-way functions exists [1], there is no $n^{-\epsilon}$-approximation algorithm for the μ-ν-Balance problem, where $\epsilon > 0$ is a constant which depends on ν. On a more positive side, we will detail an algorithm with an approximation factor of $\Omega(n^{-1/2})$ for the case where $\nu = 3$ and μ is an arbitrary constant and we will see that in the correlated setting, there exists a constant approximation algorithm for the μ-ν-Balance problem. The building blocks of these algorithms are all of a greedy flavor. They exploit clever decompositions of the objective function (in the μ-ν-Balance problem) as well as the submodularity of the influence diffusion function.

Detailed proofs for all results discussed in this communication are available online [2].

Références

[1] B. Applebaum. Pseudorandom generators with long stretch and low locality from random local one-way functions. *SIAM Journal on Computing*, 42(5) :2008–2037, 2013.

[2] Ruben Becker, Federico Corò, Gianlorenzo D'Angelo, and Hugo Gilbert. Balancing spreads of influence in a social network. *CoRR*, abs/1906.00074, 2019.

[3] E. Chlamtác, M. Dinitz, C. Konrad, G. Kortsarz, and G. Rabanca. The densest k-subhypergraph problem. *SIAM Journal on Discrete Mathematics*, 32(2) :1458–1477, 2018.

[4] Michael D. Conover, Jacob Ratkiewicz, Matthew R. Francisco, Bruno Gonçalves, Filippo Menczer, and Alessandro Flammini. Political polarization on Twitter. In *Proceedings of the 5th International Conference on Weblogs and Social Media, ICWSM 2011, Barcelona, Catalonia, Spain, July 17-21, 2011*, 2011.

[5] M. Del Vicario, A. Bessi, F. Zollo, F. Petroni, A. Scala, G. Caldarelli, H. E. Stanley, and W. Quattrociocchi. The spreading of misinformation online. *Proceedings of the National Academy of Sciences*, 113(3) :554–559, 2016.

[6] K. Garimella, A. Gionis, N. Parotsidis, and N. Tatti. Balancing information exposure in social networks. In *Advances in Neural Information Processing Systems, 30th Annual Conference on Neural Information Processing Systems, NIPS 2017, 4-9 December 2017, Long Beach, CA, USA*, pages 4666–4674, 2017.

[7] Kiran Garimella, Gianmarco De Francisci Morales, Aristides Gionis, and Michael Mathioudakis. Political discourse on social media : Echo chambers, gatekeepers, and the price of bipartisanship. In *Proceedings of the 2018 World Wide Web Conference, WWW 2018, Lyon, France, April 23-27, 2018*, pages 913–922, 2018.

[8] David Kempe, Jon M. Kleinberg, and Éva Tardos. Maximizing the spread of influence through a social network. In *Proceedings of the 9th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD 2003, Washington, DC, USA, August 24 - 27, 2003*, pages 137–146, 2003.

[9] David Kempe, Jon M. Kleinberg, and Éva Tardos. Influential nodes in a diffusion model for social networks. In *Proceedings of the 32nd International Colloquium on Automata, Languages and Programming, ICALP 2005, Lisbon, Portugal, July 11-15, 2005*, pages 1127–1138, 2005.

[10] Pasin Manurangsi. Almost-polynomial ratio ETH-hardness of approximating densest k-subgraph. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2017, Montreal, QC, Canada, June 19-23, 2017*, pages 954–961, 2017.