Byzantine Preferential Voting – Abstract HALG

Darya Melnyk dmelnyk@ethz.ch

Yuyi Wang yuwang@ethz.ch

Roger Wattenhofer wattenhofer@ethz.ch

ETH Zürich, Switzerland

September 25, 2018

Abstract

A fundamental question in distributed computing is how to establish consensus in a system where some of the communicating parties show arbitrary behavior. This problem is called Byzantine agreement: Given a set of n nodes, each of whom has a certain input value, find a protocol that establishes agreement on one of these input values under the condition that t of the nodes are Byzantine. The Byzantine nodes can act arbitrarily, and in the worst case collaborate with the aim of disturbing the protocol and preventing the correct, i.e. non-Byzantine, nodes from reaching agreement. It is assumed that the Byzantine nodes have full knowledge about the protocol, about all input values and all decisions of the correct nodes, and that the correct nodes cannot detect the Byzantine behavior as long as the Byzantine nodes follow the protocol.

We present a generalization of Byzantine agreement that lets the input values of the nodes be preference rankings of three or more candidates. Consensus on preferences, which is an important question in social choice theory, complements already known results from Byzantine agreement. For example, social choice theory shows that the plurality rule is the best voting rule when there is only two alternatives. This rule is also widely used in binary Byzantine agreement in order to select a consensus value. In addition to requiring some basic properties preferential voting also raises new questions about how to approximate consensus vectors. For the synchronous communication model we propose a deterministic algorithm to solve Byzantine agreement on rankings under a generalized validity condition, which we call Pareto-Validity. This validity condition makes sure that the agreement algorithm satisfies the weak Pareto condition, i.e., any pair of candidates that are ranked in the same way by all correct voters will be ranked in the same way in the consensus ranking.

Next to the general properties we consider a special voting rule, the Kemeny median, which minimizes the Kendall's τ distance to all correct rankings. For this rule, we show that no deterministic algorithm can approximate the Kemeny median better than by a factor of $\frac{n}{n-2t}$, which is a 3-approximation of the Kemeny median in the worst case. We also provide an algorithm that matches this lower bound. Both algorithms are based on the idea of the King algorithm (Berman et al., FOCS'89) which selects $t+1$ leaders who dictate the consensus value. This way our algorithms terminate within $t+1$ rounds and they further tolerate up to $n/3$ Byzantine nodes, which is optimal for the synchronous communication model.

Previous generalizations of binary Bazyntine agreement either struggled to be practical for multiple dimensions (Doerr et al., SPAA'11 and Stolz et al., OPODIS'15) or the number of tolerated Byzantine nodes decreased with the number of dimensions (Vaidya et al., PODC'13 and Mendes et al., STOC'13). To our knowledge, this is the first non-trivial multi-dimensional approach which can tolerate a constant fraction of Byzantine nodes.