Agreement protocols in the presence of a mobile adversary

Chris Dowden (Royal Holloway)

Suppose various processors in a network wish to reach agreement on a particular decision. Unfortunately, some unknown subset of these may be under the control of a malicious adversary who desires to prevent such an agreement being possible.

To this end, the adversary will instruct his “faulty” processors to provide inaccurate information to the non-faulty processors in an attempt to mislead them. The aim is to construct an “agreement protocol” that will always foil the adversary and enable the non-faulty processors to reach agreement successfully (perhaps after several rounds of communication).

In traditional agreement problems, it is usually assumed that the set of faulty processors is “static”, in the sense that it is chosen by the adversary at the start of the process and then remains fixed throughout all communication rounds. In this talk, we shall instead focus on a “mobile” version of the problem, providing results both for the case when the communications network forms a complete graph and also for the general case when the network is not complete.