

Report on "Coordination Capacity" by P. Cuff, H. H. Permuter, and T. M. Cover

This paper introduces and develops an information theoretic framework for problems involving coordinating action among distributed but interconnected nodes. While there has been a somewhat lengthy history of work addressing different aspects of specific instances of such problems, this paper appears to be the first to frame the broader problem in some generality.

Perhaps the main contribution of the paper is the proposed model for such problems. Specifically, action at a node takes the form of a distribution, and joint action across nodes as a joint distribution. Given (joint) actions exogenously imposed on a subset of nodes, and given a set of rate constraints between various nodes, one asks what is the set of joint actions possible over the full set of nodes. There are a couple of possible performance measures introduced, the weaker of which is based on empirical distributions ("empirical coordination"). There are no delay constraints in the formulation.

A few basic results on the general problem are established (convexity of coordination, common randomness doesn't help, etc), then the paper narrows its focus to a handful of 2- and 3-node problems with exogenously imposed actions that are i.i.d. These problems are: basic 2-node, 3-node with one isolated node, 3-node cascade, and 3-node "relay". These example special cases effectively illustrate both the richness of the framework and the kinds of tradeoffs involved. Finally, the paper makes the connection between coordination capacity and rate-distortion problems by demonstrating regions of the latter as projections of regions of the former.

What is great about this paper is that it highlights an important area of opportunity for research by the community, and develops a novel and meaningful formulation of the associated fundamental questions. As such it broadens the scope of applications of information theoretic analysis. From this perspective, this is precisely the kind of work that the society would want to highlight to the community both within and beyond the IT society. However, at the same time, the specific results presented never quite transcend the realm of "toy problems", so the paper stops short of providing the kind of compelling example of the potential of the framework that would really get the attention of the community. As such, the paper certainly represents a fine starting point, but the question is whether it goes far enough to warrant a paper award. This would appear to be the question we should focus on.