Asynchronous Recommendation Systems

Awerbuch, Patt-Shamir & Nisgav

JHU TAU

The Basic Motivation

- There are a zillion objects out there… Which would I like?
  - If I probe one, I'll know whether I like it or not
  - But I can't read them all!
- Simple solution: share the load with others
  - System provides means to disseminate actual experience reports
  - "collaborative filtering," "recommendation system," …
- Problem now: different users may have wildly different preferences. Whose view should I adopt?

Formal Model

- N **objects**, N **users**
- Each user has a binary **grade** for each object (like/dislike)
- Initially, all grades are unknown. Goal: find grades!
- A user can **probe** an object: pay cost and find grade.
- Results of probes are available to everyone (public billboard)
- Execution proceeds in atomic **steps** (synchronous?):
  - Read **billboard**
  - Probe an object (pay cost!)
  - Post result

What's Known

- DKR [STOC'2]: noise-free, restricted preferences
- AALPT [SPAA'5]: noise-free, synchronous
- AAAP [SPAA'6]: synchronous

Our Contribution

- Theorem: If $\tilde{\Omega}(1)$ of the first probes of an asynchronous schedule are by players with similar preferences, then with high probability, after $\tilde{O}(n)$ work by these players, their output stabilizes on their true preferences.
- Idea: a general randomized simulation technique, applied to the synchronous algorithms. Cost increases by a polylog factor.

Thanks!