Ex Ante Coordination and Collusion in Zero-Sum Multi-Player Extensive-Form Games

Gabriele Farina[1*], Andrea Celli[2*], Nicola Gatti[2], Tuomas Sandholm[1]

[1] Computer Science Department, Carnegie Mellon University [2] DEIB, Politecnico di Milano [*] Equal contribution

{gfarina,sandholm}@cs.cmu.edu {andrea.celli,nicola.gatti}@polimi.it

EXTENSIVE-FORM GAMES

Extensive-Form Game:

- Game tree
- Branches denote actions
- Information sets: set of nodes where the player does not know which node they are at
- Payoffs at leaves

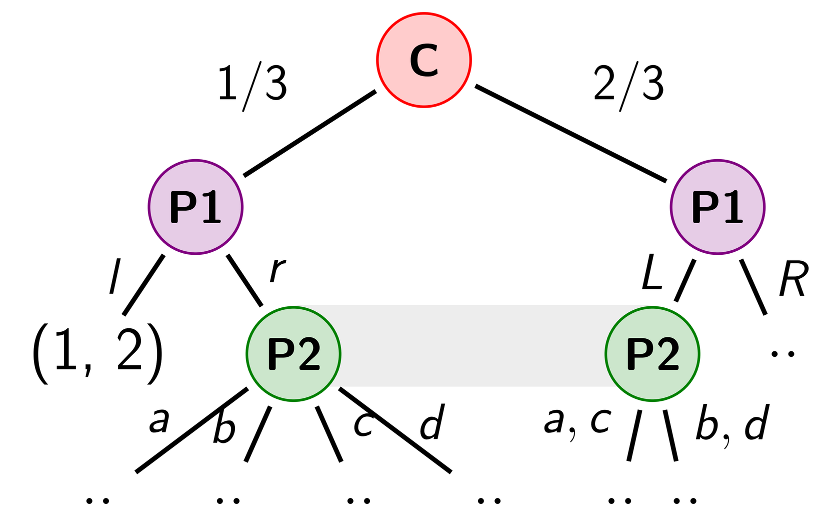

Figure 1: Sample game tree

TEAM GAMES & *EX ANTE* COORDINATION

- **Team**: set of players willing to collaborate (e.g., sharing the same objectives)
- **Key Idea**: team members can exploit coordination to reach higher rewards
- Focus on ***ex ante* coordination**: team members discuss and agree on tactics before the game starts, but are unable to communicate during the game, except through publicly-observed actions. Examples: Bridge, multi-player Poker with collusion, collusion in bidding
- *Ex ante* coordination can be modeled through an external *coordination device*
- Finding an equilibrium with *ex ante* coordination is NP-hard and inapproximable

BEHAVIORAL vs MIXED STRATEGIES

- Behavioral strategies are a compact representation of players' strategies but they cannot directly be applied to *ex ante* coordination without incurring in a loss of expressiveness

- Normal-form strategy assigning probability 1/2 to (ℓ, L) and 1/2 to (r, R) does not have an equivalent behavioral strategy

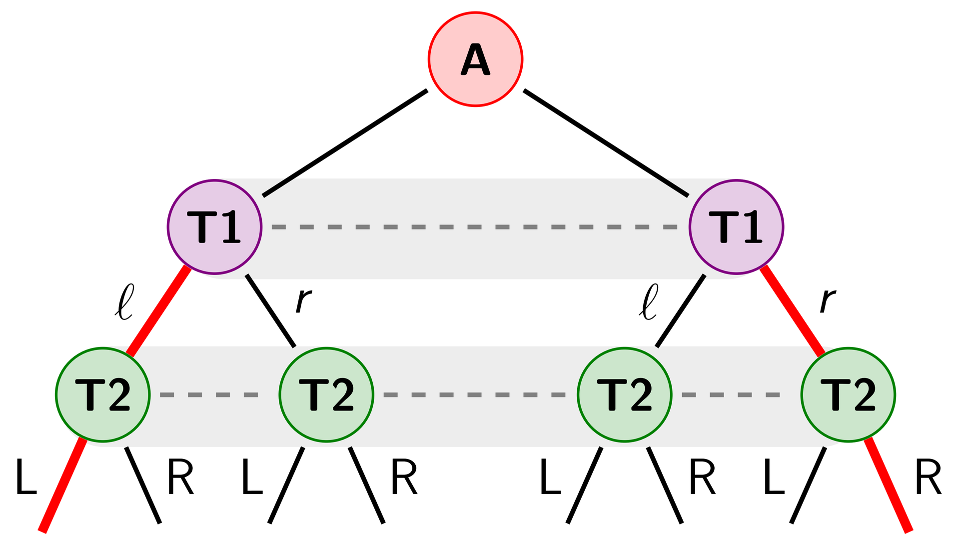

Figure 2: Extensive-form game with a team.

- The signaling power of *ex ante* coordinated strategies may not be enough to propagate the information observed

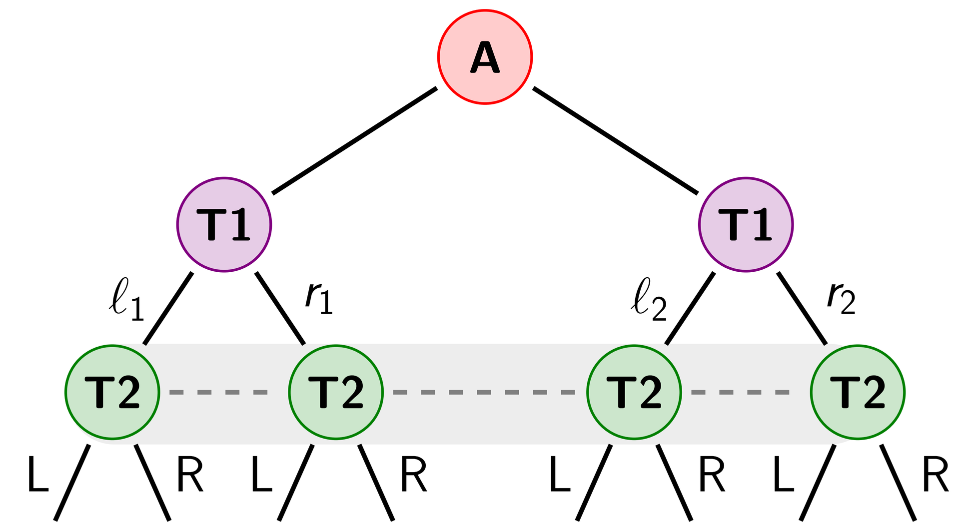

Figure 3: A game where coordinated strategies have a weak signaling power.

- Team members can be modeled as a single *meta* player with imperfect-recall (e.g., Figure 3)

REALIZATION FORM

- How can we compress the normal-form action space (exponential in the size of the tree)?
- $\rho_i^x(z)$: probability with which player i plays to reach terminal nodes z
- **Definition**: The *realization function* of player $i \in \mathcal{P}$ is the function $f_i^\Gamma : \mathcal{X}_i \to [0,1]^{|Z|}$ that maps every normal-form strategy for player i to the corresponding vector of realizations for each terminal node: $f_i^\Gamma : \mathcal{X}_i \ni x \mapsto (\rho_i^x(z_1), \ldots, \rho_i^x(z_{|Z|}))$. Player i's *realization polytope* Ω_i^Γ in game Γ is the range of f_i^Γ
- **Properties**: f_i^Γ is a linear function and Ω_i^Γ is a convex polytope. Applicable to imperfect-recall games (Ω_i^Γ has an exponential number of constraints)
- **Theorem**: the realization polytope of a non-absent-minded player is the convex hull of the set of realizations that are reachable starting from behavioral strategies

AUXILIARY GAME

Γ with team members $\{1, 2\}$ and adversary $\mathcal{A}$. The new game Γ^* is s.t. ϕ is a decision node of Player $\mathcal{T}$ with an action a_σ for each normal-form plan σ of Player 1. Each a_σ is followed by a subtree Γ_σ, obtained by fixing σ in Γ. $\mathcal{A}$ does not observe the action chosen by $\mathcal{T}$ at ϕ.

Fundamental properties of Γ^*:

- Two-player, **perfect-recall** game between $\mathcal{A}$ and a *team-player* $\mathcal{T}$
- For every *ex ante* coordinated strategy for the team in Γ, there exists a payoff-equivalent behavioral strategy for $\mathcal{T}$ in Γ^*, and vice versa

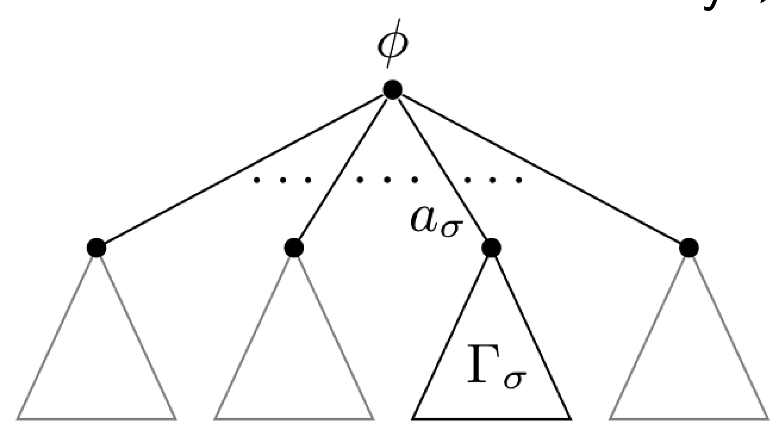

Figure 4: Structure of the auxiliary game Γ^*.

FICTITIOUS TEAM-PLAY

Key Features:

- Convergence to an NE from FP over Γ
- Best-response oracles $\mathrm{BR}_\mathcal{A}$ and $\mathrm{BR}_\mathcal{T}$ work on Γ
- Significantly faster than previous techniques (HCG)

Game	Tree size Inf.	Tree size Seq.	Δ_u	Fictitious team-play 10%	5%	2%	1.5%	1%	0.5%	HCG
K3	25	13	6	0s	0s	0s	1s	1s	1s	0s
K4	33	17	6	1s	1s	4s	4s	30s	1m 12s	9s
K5	41	21	6	1s	2s	44s	1m	4m 15s	8m 57s	1m 58s
K6	49	25	6	1s	12s	43s	5m 15s	8m 30s	23m 32s	25m 26s
K7	57	29	6	4s	17s	2m 15s	5m 46s	6m 31s	23m 49s	2h 50m
L3	457	229	21	15s	1m	14m 05s	30m 40s	1h 34m 30s	> 24h	oom
L4	801	401	21	1s	1m 31s	11m 8s	51m 5s	6h 51m	> 24h	oom

Algorithm 1 Fictitious team-play

```
function FictitiousTeamPlay(Γ)
    Initialize ω̄_A
    λ̄ ← (0, ..., 0), t ← 1
    ω̄_{T,σ} ← (0, ..., 0)   ∀σ ∈ Σ_1
    while within computational budget do
        (σ^t, ω^t_T) ← BR_T(ω̄_A)
        λ̄ ← (1 − 1/t) λ̄ + (1/t) 𝟙_{σ^t}
        ω̄_{T,σ^t} ← (1 − 1/t) ω̄_{T,σ^t} + (1/t) ω^t_T
        ω^t_A ← BR_A(λ̄, {ω̄_{T,σ}}_σ)
        ω̄_A ← (1 − 1/t) ω̄_A + (1/t) ω^t_A
        t ← t + 1
    return (λ̄, (ω̄_{T,σ})_{σ∈Σ_1})
```


Acknowledgments. This material is based on work supported by the National Science Foundation under grants IIS-1718457, IIS-1617590, and CCF-1733556, and the ARO under award W911NF-17-1-0082.