Large games and population protocols

R. Ramanujam

The Institute of Mathematical Sciences, Chennai, India

email: `jam@imsc.res.in`

Formal Methods Update Meeting, BITS-Pilani, Goa Campus

July 20, 2018

First words . . .

- Thanks to Baskar for the excellent hospitality.

First words . . .

- Thanks to Baskar for the excellent hospitality.
- Much of the work on games I talk about here is joint with Soumya Paul, currently at Univ. Luxembourg.

First words . . .

- Thanks to Baskar for the excellent hospitality.
- Much of the work on games I talk about here is joint with Soumya Paul, currently at Univ. Luxembourg.
- I do not know much about population protocols but am hoping to learn.

Summary

This talk is about *large* games and (large) *population protocols*.

- Games with a large number of players.

Summary

This talk is about *large* games and (large) *population protocols*.

- Games with a large number of players.
  - Payoffs determined by choice distributions and not profiles.

Summary

This talk is about *large* games and (large) *population protocols*.

- Games with a large number of players.
  - Payoffs determined by choice distributions and not profiles.
  - Players are anonymous, interaction is simple. Pure strategy Nash equilibria exist for a large class of games.
- Population protocols:
  - Systems with a large number of identical finite state automata.

Summary

This talk is about *large* games and (large) *population protocols*.

- Games with a large number of players.
  - Payoffs determined by choice distributions and not profiles.
  - Players are anonymous, interaction is simple. Pure strategy Nash equilibria exist for a large class of games.
- Population protocols:
  - Systems with a large number of identical finite state automata.
  - Interaction is simple, outcome based on states of interacting automata.

Summary

This talk is about *large* games and (large) *population protocols*.

- Games with a large number of players.
  - Payoffs determined by choice distributions and not profiles.
  - Players are anonymous, interaction is simple. Pure strategy Nash equilibria exist for a large class of games.
- Population protocols:
  - Systems with a large number of identical finite state automata.
  - Interaction is simple, outcome based on states of interacting automata.
  - Compute exactly the semi-linear predicates.
- Are there interesting connections between the two ? I do not know, but suspect so.

Joining clubs

I would never join a club that would take members like me.

Groucho Marx

- We like to go to restaurants that are not crowded, but not deserted either.

Joining clubs

> *I would never join a club that would take members like me.*
>
> Groucho Marx

- We like to go to restaurants that are not crowded, but not deserted either.
- *The Santa Fe bar problem*: The payoff depends on how many others act as I do.

Joining clubs

I would never join a club that would take members like me.

Groucho Marx

- We like to go to restaurants that are not crowded, but not deserted either.
- *The Santa Fe bar problem*: The payoff depends on how many others act as I do.
- Network congestion problems.

Expectations of a population

How does intersubjectivity work in large games ?

- Each person is to choose a real number $x \in [0, 100]$.

Expectations of a population

How does intersubjectivity work in large games ?

- Each person is to choose a real number $x \in [0, 100]$.
- The one who gets closest to two-thirds of the average wins the game.

Expectations of a population

How does intersubjectivity work in large games ?

- Each person is to choose a real number $x \in [0, 100]$.
- The one who gets closest to two-thirds of the average wins the game.
- In almost all experiments, the winning bid is close to 20, far from Nash equilibrium.

Expectations of a population

How does intersubjectivity work in large games ?

- Each person is to choose a real number $x \in [0, 100]$.
- The one who gets closest to two-thirds of the average wins the game.
- In almost all experiments, the winning bid is close to 20, far from Nash equilibrium.
- What would be a logical basis for expecting others to act in a particular way ?

Expectations of a population

How does intersubjectivity work in large games ?

- Each person is to choose a real number $x \in [0, 100]$.
- The one who gets closest to two-thirds of the average wins the game.
- In almost all experiments, the winning bid is close to 20, far from Nash equilibrium.
- What would be a logical basis for expecting others to act in a particular way ?
- This is hard for a one-shot game, but in repeated play, or in games of long duration, rationale based on observation can significantly affect game dynamics.

Main issues

Framework: outcomes determined by choice distributions.

- Players act individually, though the effect is collective.

Main issues

Framework: outcomes determined by choice distributions.

- Players act individually, though the effect is collective.
- Though the number of players is large, the number of player types is relatively small.

Main issues

Framework: outcomes determined by choice distributions.

- Players act individually, though the effect is collective.
- Though the number of players is large, the number of player types is relatively small.
- Players observe type distributions to determine their own.

Main issues

Framework: outcomes determined by choice distributions.

- Players act individually, though the effect is collective.
- Though the number of players is large, the number of player types is relatively small.
- Players observe type distributions to determine their own.
- This can lead to interesting stability issues.

Main issues

Framework: outcomes determined by choice distributions.

- Players act individually, though the effect is collective.
- Though the number of players is large, the number of player types is relatively small.
- Players observe type distributions to determine their own.
- This can lead to interesting stability issues.
- In turn, this can affect players' strategizing.

Illustration: small game

The mismatch game.

- Each of two players, A and B, choose between u or d. If their choices match A is paid 1 and B gets 0, and if they mismatch A is paid 0 and B gets 1.

Illustration: small game

The mismatch game.

- Each of two players, A and B, choose between u or d. If their choices match A is paid 1 and B gets 0, and if they mismatch A is paid 0 and B gets 1.
- The game has no pure strategy Nash equilibrium, but a mixed strategy Nash equilibrium (NE).

Illustration: small game

The mismatch game.

- Each of two players, A and B, choose between u or d. If their choices match A is paid 1 and B gets 0, and if they mismatch A is paid 0 and B gets 1.
- The game has no pure strategy Nash equilibrium, but a mixed strategy Nash equilibrium (NE).
- One property of this NE is that it is not *information proof*: once you are informed of the other player's move, you have an incentive to switch (from the mixed strategy to the pure mismatch strategy).

Illustration: large game

n-player mismatch game.

- Simultaneously, each of n players of type A and n players of type B have to choose between u or d.
- The payoff to every player of type A equals the proportion of players of type B that her choice matches.
- The payoff to every player of type B equals one minus the proportion of players of type A that his choice matches.

Illustration: large game

n-player mismatch game.

- Simultaneously, each of n players of type A and n players of type B have to choose between u or d.
- The payoff to every player of type A equals the proportion of players of type B that her choice matches.
- The payoff to every player of type B equals one minus the proportion of players of type A that his choice matches.
- When $n = 2$, this is the earlier game. So clearly, it inherits some of the trouble.

When n is large

The n-player mismatch game, as $n \to \infty$.

- Suppose that every player, either type, chooses one of the two randomly with equal probability.

When n is large

The n-player mismatch game, as $n \to \infty$.

- Suppose that every player, either type, chooses one of the two randomly with equal probability.
- Then within each group the proportions of the two selected choices are likely to be close to one half, and no player would be able to gain much by switching.

When n is large

The n-player mismatch game, as $n \to \infty$.

- Suppose that every player, either type, chooses one of the two randomly with equal probability.
- Then within each group the proportions of the two selected choices are likely to be close to one half, and no player would be able to gain much by switching.
- There is a high probability for the events of no possible improvement greater than some given epsilon holding *simultaneously* for all players.

Majority mismatch game

n-player mismatch game with a twist.

- The payoff to every player of type A is 1 if her choice matches that of at least one half of the choices of type B, and 0 otherwise.
- The payoff to every player of type B is 0 if his choice matches that of at least one half of the choices of type A, and 1 otherwise.

Majority mismatch game

n-player mismatch game with a twist.

- The payoff to every player of type A is 1 if her choice matches that of at least one half of the choices of type B, and 0 otherwise.
- The payoff to every player of type B is 0 if his choice matches that of at least one half of the choices of type A, and 1 otherwise.
- If n is odd, no matter what strategies are played, at every known outcome at least one half of the players will have a strong incentive to unilaterally revise their choices.

Majority mismatch game

n-player mismatch game with a twist.

- The payoff to every player of type A is 1 if her choice matches that of at least one half of the choices of type B, and 0 otherwise.
- The payoff to every player of type B is 0 if his choice matches that of at least one half of the choices of type A, and 1 otherwise.
- If n is odd, no matter what strategies are played, at every known outcome at least one half of the players will have a strong incentive to unilaterally revise their choices.
- There is no information proof equilibrium because of the discontinuity in the payoff function.

Many questions

Large games raise issues of foundational interest.

- Discontinuities in payoff functions can be critical: for instance, consider a game in which the payoff for me depends on matching at least one half of players of the other type.

Many questions

Large games raise issues of foundational interest.

- Discontinuities in payoff functions can be critical: for instance, consider a game in which the payoff for me depends on matching at least one half of players of the other type.
- What about epistemic foundations for reasoning in large games ?

Many questions

Large games raise issues of foundational interest.

- Discontinuities in payoff functions can be critical: for instance, consider a game in which the payoff for me depends on matching at least one half of players of the other type.
- What about epistemic foundations for reasoning in large games ?
- What are good models for large games ?

Many questions

Large games raise issues of foundational interest.

- Discontinuities in payoff functions can be critical: for instance, consider a game in which the payoff for me depends on matching at least one half of players of the other type.
- What about epistemic foundations for reasoning in large games ?
- What are good models for large games ?
- What are the implications for social algorithms ?

Any good news ?

In some respects, large games are easier to reason about than small ones.

- Behaviour for large n can smooth out many individual irregularities.

Any good news ?

In some respects, large games are easier to reason about than small ones.

- Behaviour for large n can smooth out many individual irregularities.
- Many problems related to mutual intersubjectivity and surprise moves disappear.
- When the number of players is large but the number of player types is small, we can sometimes reduce the analysis to small games.

The model

We study infinite play, since we are interested in long run stability issues.

The model

We study infinite play, since we are interested in long run stability issues.

- We work with repeated normal form games.

The model

We study infinite play, since we are interested in long run stability issues.

- We work with repeated normal form games.
- We will assume that the action sets for all players are identical: $A_1 = \ldots = A_n = A$. Let $|A| = k$.

The model

We study infinite play, since we are interested in long run stability issues.

- We work with repeated normal form games.
- We will assume that the action sets for all players are identical: $A_1 = \ldots = A_n = A$. Let $|A| = k$.
- An action distribution is a tuple $\mathbf{y} = (y_1, y_2, \ldots, y_k)$ such that $\forall i$, $y_i \geq 0$ and $\sum_{i=1}^{k} y_i \leq n$.

The model

We study infinite play, since we are interested in long run stability issues.

- We work with repeated normal form games.
- We will assume that the action sets for all players are identical: $A_1 = \ldots = A_n = A$. Let $|A| = k$.
- An action distribution is a tuple $\mathbf{y} = (y_1, y_2, \ldots, y_k)$ such that $\forall i$, $y_i \geq 0$ and $\sum_{i=1}^{k} y_i \leq n$.
- A function $f_i : \mathbf{Y} \to \mathbb{Q}$ for every player i.

Limit average payoffs

The analysis is quantitative.

- Given an initial vertex v_0 consider an infinite play $\rho = v_0 \xrightarrow{y_1} v_1 \xrightarrow{y_2} \ldots$ in the arena.

Limit average payoffs

The analysis is quantitative.

- Given an initial vertex v_0 consider an infinite play $\rho = v_0 \xrightarrow{y_1} v_1 \xrightarrow{y_2} \dots$ in the arena.
- Player i gets a limit average payoff:

$$p_i(\rho) = \liminf_{m \to \infty} \frac{1}{m} \sum_{j=1}^{m} f_i(y_j).$$

Main questions

We study player types specified by formulas that code up beliefs of players about others. We will discuss the logic and its formulas later.

- Main question: Given an initial type distribution of players, which types are eventualy stable ?

Type based analysis

When the number of types is $t << n$ where n is the number of players, can one carry out all the analysis using only the t types and then lift the results to the entire game ?

Type based analysis

When the number of types is $t << n$ where n is the number of players, can one carry out all the analysis using only the t types and then lift the results to the entire game ?

- Why should such an analysis be possible ? This is because outcomes are determined by player choice distributions rather than strategy profiles.

Type based analysis

When the number of types is $t << n$ where n is the number of players, can one carry out all the analysis using only the t types and then lift the results to the entire game ?

- Why should such an analysis be possible ? This is because outcomes are determined by player choice distributions rather than strategy profiles.
- When types describe finite memory strategies (as in the case of first order logic specifications) we can consider them to be *finite state transducers* that observe play, make boundedly many observations and output the moves to be played.

Type based analysis

When the number of types is $t << n$ where n is the number of players, can one carry out all the analysis using only the t types and then lift the results to the entire game ?

- Why should such an analysis be possible ? This is because outcomes are determined by player choice distributions rather than strategy profiles.
- When types describe finite memory strategies (as in the case of first order logic specifications) we can consider them to be *finite state transducers* that observe play, make boundedly many observations and output the moves to be played.
- Can we use the structure of these transducers to do this reduction ?

Products of transducers

When we have n finite memory players, the analysis space is the n-fold product of these automata. We wish to to map this space into a t-fold product.

Products of transducers

When we have n finite memory players, the analysis space is the n-fold product of these automata. We wish to to map this space into a t-fold product.

- We show in the case of deterministic transducers, that the the product of a type with itself is isomorphic to the type.

Products of transducers

When we have n finite memory players, the analysis space is the n-fold product of these automata. We wish to to map this space into a t-fold product.

- We show in the case of deterministic transducers, that the the product of a type with itself is isomorphic to the type.
- Thus a population of 1000 players with only two types needs to be represented only by pairs of states and not 1000-tuples.

Products of transducers

When we have n finite memory players, the analysis space is the n-fold product of these automata. We wish to to map this space into a t-fold product.

- We show in the case of deterministic transducers, that the the product of a type with itself is isomorphic to the type.
- Thus a population of 1000 players with only two types needs to be represented only by pairs of states and not 1000-tuples.
- However, there is no free lunch: an exponential price has to be paid for determinization. But we characterize when this can be worthwhile.

Transducer reduction

We use output preserving homomorphisms to equate transducers.

Transducer reduction

We use output preserving homomorphisms to equate transducers.

- Theorem: Suppose that we have n players, k choices and t types. Let $p = \max_i |R_i|$, where the ith type formula induces a nondeterministic FST with state space R_i. Then the type based analysis is more efficient when

$$\frac{n}{t} > 0.693 \cdot k \cdot \pi(p)$$

where $\pi(p)$ is the number of primes less than or equal to p.

Existence of Nash equilibria

In general, games possess only mixed-strategy Nash equilibria.

Existence of Nash equilibria

In general, games possess only mixed-strategy Nash equilibria.

- Theorem: In large games, pure strategy Nash equilibria exist and are information proof for a class of games whose best-response function is direction preserving.

Existence of Nash equilibria

In general, games possess only mixed-strategy Nash equilibria.

- Theorem: In large games, pure strategy Nash equilibria exist and are information proof for a class of games whose best-response function is direction preserving.
- For $p, q \in Z^d$, let $A_{p,q} = \{r \in Z^d \mid p \leq r \leq q\}$.

Existence of Nash equilibria

In general, games possess only mixed-strategy Nash equilibria.

- Theorem: In large games, pure strategy Nash equilibria exist and are information proof for a class of games whose best-response function is direction preserving.
- For $p, q \in Z^d$, let $A_{p,q} = \{r \in Z^d \mid p \leq r \leq q\}$.
- A map $F : A_{p,q} \to R^d$ is said to be direction-preserving if for any $r_1, r_2 \in A_{p,q}$ with $|r_1 - r_2|_\infty \leq 1$, we have, for all $i, 1 \leq i \leq d$: $(F_i(r_1) - r_1^i)(F_i(r_2) - r_2^i) \geq 0$.

Existence of Nash equilibria

In general, games possess only mixed-strategy Nash equilibria.

- Theorem: In large games, pure strategy Nash equilibria exist and are information proof for a class of games whose best-response function is direction preserving.
- For $p, q \in Z^d$, let $A_{p,q} = \{r \in Z^d \mid p \leq r \leq q\}$.
- A map $F : A_{p,q} \to R^d$ is said to be direction-preserving if for any $r_1, r_2 \in A_{p,q}$ with $|r_1 - r_2|_\infty \leq 1$, we have, for all $i, 1 \leq i \leq d$: $(F_i(r_1) - r_1^i)(F_i(r_2) - r_2^i) \geq 0$.
- Note that the fixed point computation happens in a discrete space (where we do not have Brouwer - Kakutani fixed point theorems). So we use a different technique due to Chen.

Local equilibrium

Nash equilibrium is not the best notion in large games. We should ask when a strategy profile in a large game constitute an equilibrium.

Local equilibrium

Nash equilibrium is not the best notion in large games. We should ask when a strategy profile in a large game constitute an equilibrium.

- Let π be a profile. For every player i, $\mathcal{T}(\pi^{-i})$ is a type projection, giving the set of types visible to i.

Local equilibrium

Nash equilibrium is not the best notion in large games. We should ask when a strategy profile in a large game constitute an equilibrium.

- Let π be a profile. For every player i, $\mathcal{T}(\pi^{-i})$ is a type projection, giving the set of types visible to i.
- In general, though there are n players in the game, player i sees k_i types, and hence is involved in a $k_i + 1$-player game.

Local equilibrium

Nash equilibrium is not the best notion in large games. We should ask when a strategy profile in a large game constitute an equilibrium.

- Let π be a profile. For every player i, $\mathcal{T}(\pi^{-i})$ is a type projection, giving the set of types visible to i.
- In general, though there are n players in the game, player i sees k_i types, and hence is involved in a $k_i + 1$-player game.
- Let $\sigma \in \Sigma_i$. We say σ is a best response to a set T of player types, if for every profile π such that $\mathcal{T}(\pi^{-i}) = T$, $u_i(\sigma; \pi^{-i}) \geq u_i(\pi)$.

Local equilibrium

Nash equilibrium is not the best notion in large games. We should ask when a strategy profile in a large game constitute an equilibrium.

- Let π be a profile. For every player i, $\mathcal{T}(\pi^{-i})$ is a type projection, giving the set of types visible to i.
- In general, though there are n players in the game, player i sees k_i types, and hence is involved in a $k_i + 1$-player game.
- Let $\sigma \in \Sigma_i$. We say σ is a best response to a set T of player types, if for every profile π such that $\mathcal{T}(\pi^{-i}) = T$, $u_i(\sigma; \pi^{-i}) \geq u_i(\pi)$.
- A profile π is in local equilibrium if for all i, $\pi(i)$ is the best response to $\mathcal{T}(\pi^{-i})$.

A new notion

When every strategy defines a unique type this is Nash equilibrium.

A new notion

When every strategy defines a unique type this is Nash equilibrium.

- As a rule, local equilibria are conservative; they constitute response to potential strategies based on observations rather than strategies.

A new notion

When every strategy defines a unique type this is Nash equilibrium.

- As a rule, local equilibria are conservative; they constitute response to potential strategies based on observations rather than strategies.
- Stability in this notion is sensitive to the way projections of strategies to types is defined.

A new notion

When every strategy defines a unique type this is Nash equilibrium.

- As a rule, local equilibria are conservative; they constitute response to potential strategies based on observations rather than strategies.
- Stability in this notion is sensitive to the way projections of strategies to types is defined.
- The projection function is uniform in the definition above. In general, it would be indexed by players, or better, by types again !

A new notion

When every strategy defines a unique type this is Nash equilibrium.

- As a rule, local equilibria are conservative; they constitute response to potential strategies based on observations rather than strategies.
- Stability in this notion is sensitive to the way projections of strategies to types is defined.
- The projection function is uniform in the definition above. In general, it would be indexed by players, or better, by types again !
- We can show that local equilibrium is a new notion, in the sense that we can define games that have local but no global equilibria, or the other way.

Two kinds of stability

Local equilibria predict stable play in the dynamics of strategy improvement. But this assumes visibility to be static.

Two kinds of stability

Local equilibria predict stable play in the dynamics of strategy improvement. But this assumes visibility to be static.

- In large games, visibility is dynamic as well.

Two kinds of stability

Local equilibria predict stable play in the dynamics of strategy improvement. But this assumes visibility to be static.

- In large games, visibility is dynamic as well.
- This results in a dynamic game form.

Two kinds of stability

Local equilibria predict stable play in the dynamics of strategy improvement. But this assumes visibility to be static.

- In large games, visibility is dynamic as well.
- This results in a dynamic game form.
- Note that the two dynamics are recursive in each other.

Two kinds of stability

Local equilibria predict stable play in the dynamics of strategy improvement. But this assumes visibility to be static.

- In large games, visibility is dynamic as well.
- This results in a dynamic game form.
- Note that the two dynamics are recursive in each other.
- We describe the game form dynamics by neighbourhoods.

Last words on large games

We have not presented theorems on large games.

- Game theorists have mainly studied utility functions and learning, interaction / communication models are very simplistic.

Last words on large games

We have not presented theorems on large games.

- Game theorists have mainly studied utility functions and learning, interaction / communication models are very simplistic.
- Our results:
  - Soumya Paul and R. Ramanujam, "Dynamics of choice restriction in large games", *Journal of Game Theory Review*, vol 15, no. 4, 156-184, 2013.
  - Soumya Paul and R. Ramanujam, "Subgames within large games and the heuristic of imitation", *Studia Logica*, vol 102, no. 2, 361-388, 2014.

Population protocols

Population protocols were introduced by Angluin, D., Aspnes, J., Diamadi, Z., Fischer, M.J., Peralta, R.: Computation in networks of passively mobile finite-state sensors. Distributed Computing 18(4), 2006.
The defining features of the basic model are:

- Finite-state agents and uniformity.
- Computation by direct interaction, and unpredictable interaction patterns.
- Distributed inputs and outputs.
- Convergence rather than termination.

The basic model

An n-agent PP is a tuple $P = (Q, \delta, \iota, \omega)$ over (Σ, Γ) where Σ is the input alphabet, Γ is the output alphabet, $\iota : \Sigma \to Q$, $\omega : Q \to \Gamma$ and $\delta \subseteq Q^4$.

- The initial configuration is determined by the inputs via ι.
- δ describes pairwise interaction and thus configuraion change.
- Via ω all automata constantly produce output.
- Fairness assumption: if C appears infinitely often in a computation and $C \to C'$ then C' appears infinitely often in it.
- A protocol computes a function f that maps multisets of elements of Σ to Γ if, for every such multiset I and every fair execution that starts from the initial configuration corresponding to I, the output value of every agent eventually stabilizes to $f(I)$.

The OR protocol

The aim of the protocol is to output the 'or' of all input bits.

- $\Sigma = \Gamma = Q = \{0, 1\}$ and the input and output maps are the identity functions.
- The only interaction in δ is $(0, 1) \rightarrow (1, 1)$.
- If all agents have input 0, no agent will ever be in state 1.
- If some agent has input 1 the number of agents with state 1 cannot decrease and fairness ensures that it will eventually increase to n.

The dancers protocol

The agents are dancers, and each dancer is (exclusively) a leader or a follower. The problem is to determine whether there are more leaders than followers.

- $\Gamma = \{0, 1\}$. We set $\Sigma = \{L, F\}$ and $Q = \{L, F, 0, 1\}$.
- The input map is the identity; the output maps L and 1 to 1, F and 0 to 0.
- δ has: $(L, F) \to (0, 0)$, $(L, 0) \to (L, 1)$, $(F, 1) \to (F, 0)$ and $(0, 1) \to (0, 0)$.
- In case of a tie, the last rule ensures that the output stabilizes to 0.

Convergence

It is not obvious that this protocol converges.

- Consider the sequence of configurations:

$$(L, L, F), (0, L, 0), (1, L, 0), (0, L, 0), (0, L, 1), (0, L, 0)$$

- Repeating the last four transitions yields a non-converging execution, but it is not fair.

Some exercises

The notion of fairness is subtle: it is distinct from the condition that each pair of agents must interact infinitely often. E.g. consider $(L, L, L)^\omega$ where all interactions take place between the first two agents: it is fair.

- Show the dancers protocol converges in every fair execution.
- Design a protocol to determine whether more than $2/3^{rds}$ of the dancers are leaders.
- Design a protocol to determine whether more than $2/3^{rds}$ of the dancers play the same role.

Modular arithmetic

Suppose each agent is given an input from $\Sigma = \{0, 1, 2, 3\}$. Consider the problem of computing the sum of the inputs, modulo 4.

- The protocol gathers the sum (modulo 4) into a single agent. Once an agent has given its value to another agent, its value becomes null, and it obtains its output value from the eventually unique agent with a non-null value.

Modular arithmetic

Suppose each agent is given an input from $\Sigma = \{0, 1, 2, 3\}$. Consider the problem of computing the sum of the inputs, modulo 4.

- The protocol gathers the sum (modulo 4) into a single agent. Once an agent has given its value to another agent, its value becomes null, and it obtains its output value from the eventually unique agent with a non-null value.
- $Q = \{0, 1, 2, 3, n_0, n_1, n_2, n_3\}$, where n_v stands for null value with output v.
- δ has $(v_1, v_2) \rightarrow (v_1 + v_2, n_{v_1+v_2})$ (addition modulo 4) and $(v_1, n_{v_2}) \rightarrow (v_1, n_{v_1})$.

Computability

We can represent multisets over Σ by vectors. For instance (a, b, a, b, b) over $\Sigma = \{a, b, c\}$ by $(2, 3, 0)$. Thus we can speak of input vectors $(x_1, \ldots, x_d)$ in $\mathcal{N}^d$ where $d = |\Sigma|$.

- Threshold predicates are of the form $\Sigma_{i=1}^{d} c_i x_i < a$, and remainder predicates are: $\Sigma_{i=1}^{d} c_i x_i = a(mod b)$.
- Angluin et al (easily) show that population protocols can compute these and their boolean combinations.
- Surprisingly, the converse also holds: these are the only predicates that a population protocol can compute.

The theorem

Theorem (Angluin et al): A predicate is computable in the basic population protocol model if and only if it is semilinear.

- The proof is quite involved, the main tool is Higman's Lemma. There are three main steps to the proof.

The steps

The three main steps:

- Show that any predicate stably computed by a population protocol is a finite union of monoids: sets of the form $\{(b + k_1 a_1 + k_2 a_2 + \ldots) \mid k_i \in \mathcal{N} \text{ for all } i\}$, where the number of terms may be infinite.

The steps

The three main steps:

- Show that any predicate stably computed by a population protocol is a finite union of monoids: sets of the form $\{(b + k_1 a_1 + k_2 a_2 + \dots) \mid k_i \in \mathcal{N} \text{ for all } i\}$, where the number of terms may be infinite.
- Show that when detecting if a configuration x is output-stable, it suffices to consider its truncated version:

$$\tau_k(x_1, \dots, x_d) = (min(x_1, k), \dots, (x_d, k))$$

provided k is large enough to encompass all of the minimal non-output-stable configurations. (There are only finitely many such configurations.)

The steps

The three main steps:

- Show that any predicate stably computed by a population protocol is a finite union of monoids: sets of the form $\{(b + k_1 a_1 + k_2 a_2 + \dots) \mid k_i \in \mathcal{N} \text{ for all } i\}$, where the number of terms may be infinite.
- Show that when detecting if a configuration x is output-stable, it suffices to consider its truncated version:

$$\tau_k(x_1, \dots, x_d) = (min(x_1, k), \dots, (x_d, k))$$

  provided k is large enough to encompass all of the minimal non-output-stable configurations. (There are only finitely many such configurations.)
- Finally we can reduce the problem to a form of coverability.

Variations on the model

Many variants of the basic model have been studied in the literature.

- One-way interaction: we have sender and receiver agents. This leads to immediate and delayed observation models, and queued transmission models.
- Delayed observation models can detect multiplicity of input symbols (upto a threshold) and essentially only such predicates.
- Immediate observation models can count the number of agents with a particular input symbol (upto a threshold).
- Queued models have the same power as the basic model.
- Interaction graphs, in general, lead to Turing computability.
- Many papers study random interaction models.

Games

We consider two player games in normal form.
The two players are I (initiator) and R (responder).
Let $S(I), S(R)$ denote the (finite) sets of strategies of the players.
$BR_I : S(R) \to S(I)$ is the best response map for player I; BR_R similarly.
Assume that $S(I) = S(R) = S$, for now, and let Δ be a fixed integer constant.

From games to protocols

To each such game we can associate a population protocol as follows:

- $Q = S$.
- $(q_1, q_2, q_1', q_2') \in \delta$ iff:
  - $q_1' = q_1$ if $u_I(q_1, q_2) \geq \Delta$; $q_1' = x \in BR_I(q_2)$, otherwise.
  - $q_2' = q_2$ if $u_R(q_1, q_2) \geq \Delta$; $q_2' = x \in BR_R(q_1)$, otherwise.
- We vary input and output functions and Δ to get a class of protocols associated with the game.

Pavlovian protocol

Call a population protocol Pavlovian if it can be obtained from a game by the rules above.

- **Proposition**: The class of predicates computable by Pavlovian population protocols is closed under negation.
- The proof proceeds by a kind of determinization: by constructing 'equivalent' games that have unique best response, which makes the rules above deterministic.
- It is not clear that predicates computable by Pavlovian population protocols are closed under conjunction or disjunction.

Product protocols

Consider k two-player games, and define (in the natural way) the associated protocol by the k-fold product of the rules above. We call them Multi-Pavlovian protocols.

- **Theorem**: The predicates defined by Multi-Pavlovian protocols are exactly the semi-linear ones. Thus every population protocol corresponds to a finite product of 2-player normal form games.
- There are surprises when we restrict ourselves to symmetric games. (A population protocol is symmetric if whenever $(q_1, q_2, q_3, q_4) \in \delta$ then $(q_2, q_1, q_4, q_3) \in \delta$ as well.

Last words

- Evolutionary game theory is a well-studied subject approach to population dynamics, modelling simple interactions on a large scale.

Last words

- Evolutionary game theory is a well-studied subject approach to population dynamics, modelling simple interactions on a large scale.
- Population protocols provide a very interesting model of computation that is very similar.

Last words

- Evolutionary game theory is a well-studied subject approach to population dynamics, modelling simple interactions on a large scale.
- Population protocols provide a very interesting model of computation that is very similar.
- This may be one way to scale up systems of automata and their interactions the study of which has been too rigid.

Last words

- Evolutionary game theory is a well-studied subject approach to population dynamics, modelling simple interactions on a large scale.
- Population protocols provide a very interesting model of computation that is very similar.
- This may be one way to scale up systems of automata and their interactions the study of which has been too rigid.
- The model of games and automata has been used well in the context of systems with a fixed number of players / components. Moving to large distributed systems, the models of large games and population protocols seem promising and worthy of study.

Discussion time

Thank you.

Questions, comments, suggestions welcome; also, please write to jam@imsc.res.in.