ICST - DISTTRUST 2006

Computational economics for self-organizing networks

Reputation for autonomic systems

Pietro Michiardi – Institut Eurecom

Computational Economics

- Game Theory (GT) and Mechanism Design (MD) with an algorithmic complexity flavor
- Focus of this presentation:
  - Algorithms and protocols concerned with self-organizing communication networks
  - System designers' perspective
- Analytical framework
  - Informal background information
  - Application to real systems

Motivations

- Fundamental assumption made in the past:
  - Entities correctly execute protocols/algorithms, except for faulty or malicious ones
- However, the system designer has no control on individuals’ (software agents, humans …) behavior
  - Agents can/will deviate from prescribed behavior, if beneficial for them
- Opportunistic/selfish behavior prevent from the realization of system designers’ objectives

Game Theory: traditional approach

- Framework to model systems and participants' behavior
- Definition of a game:
  - Set of players
  - Preference ordering over system outcomes
  - Set of available strategies
- Rationality and Self-interest : individuals are able to predict and evaluate the outcome of a game, and maximize their gains

GT: traditional approach (2)

- Utility Function: represents the (real) system
- Solution concepts: Dominant strategy, Nash/Bayesian
  - Prove of existence is hard, finding the equilibrium point requires many assumptions
- GOAL 1: find the system operating point
- GOAL 2: study strategies, their stability properties, repeated interactions

Mechanism Design: traditional approach

- How to design systems so that individuals' selfish behavior results in the desired system-wide goals
  - System objectives termed Social Choice Function: *e.g.*, efficiency (system welfare), stability, and fairness
- GT models the system and the entities
  - Equilibrium solution concepts borrowed from GT
- Individuals *i* hold private information about their preferences, termed *type* θ_i
  - Selfishness = players lie about their type

Traditional MD at a glance

- Quasi-linear utility function: $u_i(\theta_i, o) = v_i(\theta_i, o^*) + p_i$
- Incentive & Allocation (that implements the SCF)
  - Allocation example: g(θ) = o* = argmax (Sum v_i(θ, o)) = SCF(θ)
  - Payment example: pay as much as a cheater would gain
- Desirable properties:
  - Truthfulness / strategy proofness: truth revelation is dominant
  - Incentive compatibility: system outcome = social function

Traditional MD: a sketch

- *σ, θ* *strategies*
- *i individuals*
- *Mechanism = center*

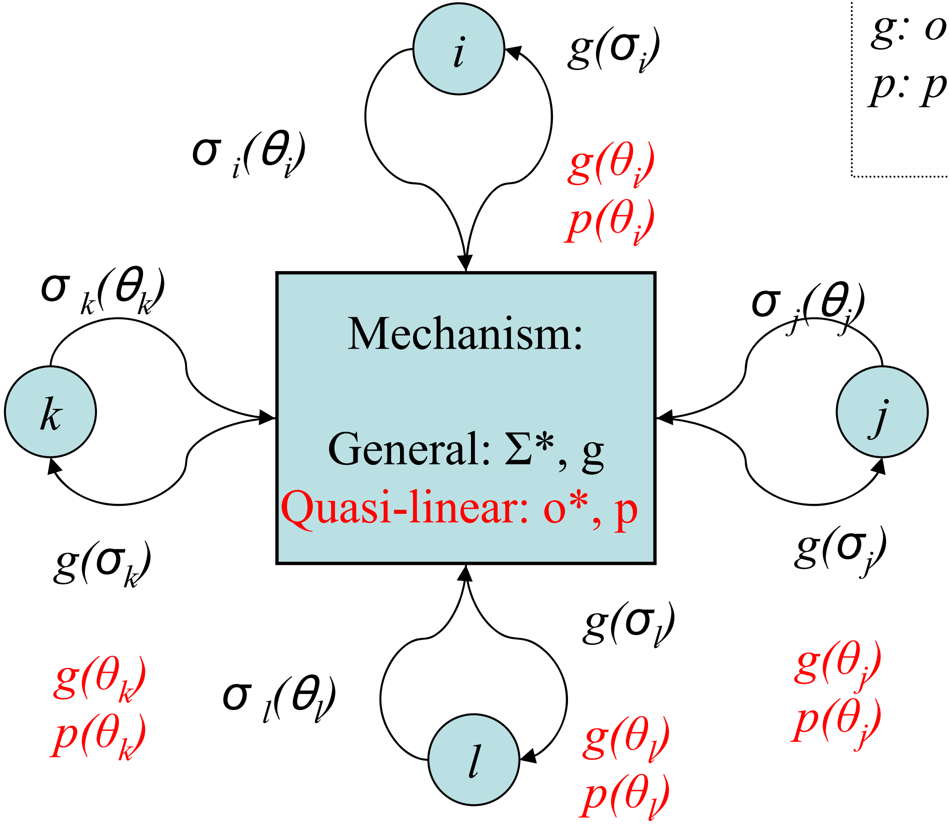

σ, θ: reports by agents

g: outcome rule
p: payment rule

Traditional MD: goals

- Guide the system operating point to an optimal position, from the designer's perspective

- Formulate and study incentives or payment schemes to stimulate participants to behave as prescribed by the original protocol/algorithm, assuming selfish behavior

From traditional MD to Algorithmic MD to Distributed MD (1)

- Limitations of traditional MD
  - Computationally unbounded agents
  - Centralized approach (the mechanism, the center)
  - Cost-free, error-free communication with the center

- Limitations that derive from GT
  - Solution concepts often too "demanding"
  - No repeated interactions

From traditional MD to Algorithmic MD to Distributed MD (2)

Algorithmic MD (AMD)

- Considers computational complexity of MD
  - Agent level: valuation and strategic complexity
  - Infrastructure level: outcome computation and communication complexity
- Approaches:
  - Polynomial time solutions (Nisan *et. al.*)
  - Protocol compatibility (Feigenbaum, Shenker)

From traditional MD to Algorithmic MD to Distributed MD (3)

Distributed (Algorithmic) MD

- Eliminates the need for a computational center
  - Agents compute the outcome rule and payments
- Theory is still in its infancy, many open problems:
  - Communication security
  - Privacy issues
  - Self-sustaining systems issues: budget balance is not always guaranteed

Recent advances

- Automatic Mechanism Design
  - Consider complexity of mechanism definition
  - Deals with the limited (a couple!!) number of available mechanism/incentive schemes and their limitations
    - Quasi-linear utilities
    - No payments
    - Different social choice functions (SCF, defined by the system designer)
- Interaction with Security and Cryptography
  - Cryptography has emerged as a fundamental tool, especially in distributed settings
  - Reputation/trust has emerged as a useful metric to prioritize interaction
- Towards autonomic systems: the designers' role is still preponderant in all analytical models, because of the SCF
  - Let the system participants dynamically define the SCF

Application: Mobile ad hoc networks

- Prescribed behavior: forward packets for others
- Problem: forwarding cost energy, which is scarce
- Nodes deviate from prescribed behavior

- Game theoretical models: equilibrium = no cooperation
  - Prisoner’s Dilemma
  - Cyclic dependent networks

- Solution: incentive schemes or (?) reputation
- Cooperation strategies: TFT, G-TFT, CORE, …
- *Evolutionary game theory*: blends game theory and genetic programming

Application: Mobile ad hoc networks (Continued)

- Application of mechanism design: Ad hoc-VCG
  - Node declare its cost for forwarding a packet
  - Node can lie to be on the shortest path or to be excluded from it

- Payment for an intermediate node on the shortest path from a source *S* to a destination *D* = declared cost for forwarding a packet plus a premium

- Premium = the difference of the total cost of the shortest path from *S* to *D* without *vi* and the shortest path from *S* to *D* with *vi* (and its declared cost)

- The payment is defined in such a way that node *vi* will get the same amount independent of what it declares as its forwarding cost.

- Budget balance problem
- Centralized solution`

Application: Content Distribution

- Prescribed behavior: share upload capacity
- Problem: limited resource, free-riding

- Game theoretical models
  - Marriage theorem
  - Bipartite matching
  - Network formation games

- Practical instance: BitTorrent, peer selection strategy

Conclusion

- Powerful analytical frameworks
- From theory to practice is often an open problem
  - Many impossibility results!
- Question: what is the difference between reputation in MANET and peer selection in P2P?
  - Hint: “shadow of the future”

MD Example: VCG example

- Second-best sealed bid auction
  - Quasi-linear utility
  - Valuation function over (single) item allocation
- Players reveal their type to the center (*i.e.* mechanism)
  - But, don't need to know other players' values/types
- Center computes system allocation to maximize SCF
  - Center allocates item to highest bid
  - Winner pays the second highest price
- Truthful revelation is dominant
- Not budget-balanced

Generalized VCG, formal definition

- Center collects $\theta=(\theta_1,\ldots,\theta_n)$ from agents
- $g(\theta)$: select o* to maximize $\mathrm{sum}(v_i(\theta,o))$
- $p(\theta)$: agent i pays $\mathrm{sum}_{j\neq i}(v_j(\theta_j,o^{-i})) - \mathrm{sum}_{j\neq i}(v_j(\theta_j,o^*))$ where
  o^{-i} solves $\max_a \mathrm{sum}_{j\neq i}(v_j(\theta_j,o))$

- Theorem: the VCG mechanism is truthful and implements the SCF