Mechanism Design
Powered by Social Interactions

Dengji Zhao

ShanghaiTech University, Shanghai, China

A tutorial @ AAMAS, IJCAI 2019

2009 DARPA Red Balloon Challenge

- The $40,000 challenge award would be granted to the first team to submit the locations of 10 moored, 8-foot, red weather balloons at 10 previously undisclosed fixed locations in the continental United States.

2009 DARPA Red Balloon Challenge

MIT Crowdsourced Solution (The Winner):

- "*We're giving $2000 per balloon to the first person to send us the correct coordinates, but that's not all – we're also giving $1000 to the person who invited them. Then we're giving $500 whoever invited the inviter, and $250 to whoever invited them, and so on ...*"
- got over 5,000 of participants, won the competition in under 9 hours.

2009 DARPA Red Balloon Challenge

MIT Crowdsourced Solution (The Winner):

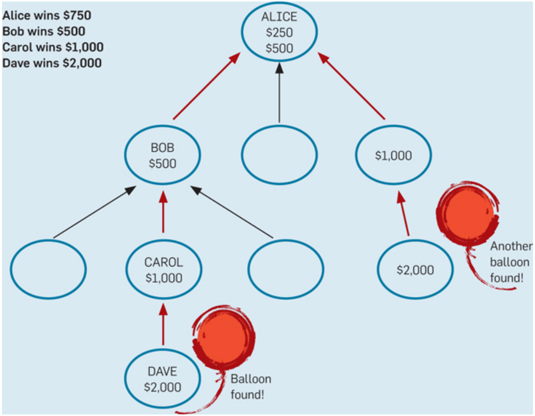

- Pickard, G., et al., Time-Critical Social Mobilization. Science, 2011. 334(6055): p. 509-12.

PinDuoDuo (like Groupon)

PinDuoDuo (like Groupon)

Achievements:

- went online in Sep 2015
- got over 2 million users in two weeks
- by Feb 2016, got over 20 million users
- IPO in Jul 2018

PinDuoDuo (like Groupon)

Their group buying model:

1. choose one product
2. join a group buying deal or initiate a new group buying deal
3. wait or invite friends to join the deal
4. when the required number of buyers is reached, they all buy the product with a cheaper price

What are the incentives?

More participants, higher chance to win!!!

- 2009 DARPA Red Balloon Challenge
  - Inviting more friends has higher chance to win (higher utility)
- PinDuoDuo
  - Inviting more friends has higher chance to get cheap items (higher utility)

What if it is a competition?

- resource allocation such as auctions
- task allocation

More participants means lower chance to win!!!

Diffusion Mechanism Design

Mechanism Design on Social Networks

Design mechanisms/markets under competitive environment such that participants are incentivized to invite more participants/competitors to join the mechanisms.

Starter: Promote a Sale via Social Networks

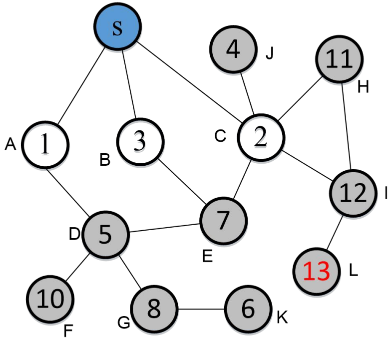

- The seller (blue node) sells one item and has only three connections/neighbours in the network (A,B,C).
- Each node is a potential buyer and the value is her highest willing payment to buy the item (valuation).
- The seller's revenue of applying second price auction without promotion is 2.
- but the highest willing payment of the network is 13.

Starter: Promote a Sale via Social Networks

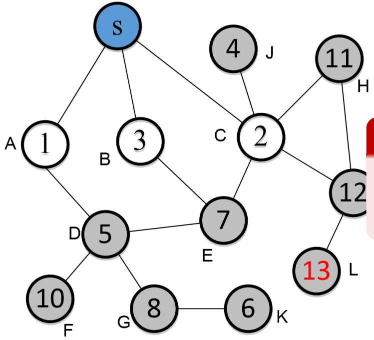

Question

How the seller could do to increase her profit?

Traditional Sale Promotions

Traditional sale promotions:

- Promotions in shopping centres
- Keywords based ads via search engines such as Google
- Ads via social media such as WeChat, Facebook, Twitter

Traditional Sale Promotions

Traditional sale promotions:

- Promotions in shopping centres
- Keywords based ads via search engines such as Google
- Ads via social media such as WeChat, Facebook, Twitter

Challenge

- The return of these promotions are unpredictable.
- The seller may LOSE from the promotions.

Tackle the Challenge

Build promotion inside the market mechanism such that

1. the promotion will never bring negative utility/revenue to the seller.
2. all buyers who are aware of the sale are incentivized to diffuse the sale information to all her neighbours.

Tackle the Challenge

Build promotion inside the market mechanism such that

1. the promotion will never bring negative utility/revenue to the seller.
2. all buyers who are aware of the sale are incentivized to diffuse the sale information to all her neighbours.

"Diffusion Mechanism Design"

New Challenges

Why a buyer would bring more buyers to compete with her?

- only if their efforts are rewarded, *but the seller doesn't want to lose*!
- we cannot just pay each node a fixed amount to incentivise them to diffuse the information.

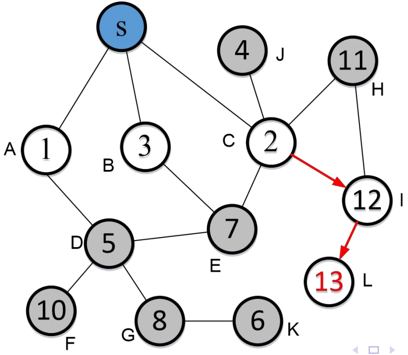

What is Mechanism Design

What is Mechanism/Market Design?

- it is known as Reverse Game Theory

What is Game Theory

- **Game theory** is the study of mathematical models of conflict and cooperation between intelligent rational decision-makers (wiki) [von Neumann and Morgenstern 1944].

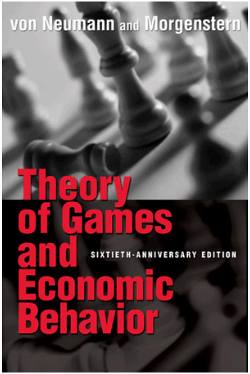

- Non-cooperative games: Go, poker, rock-paper-scissors
- Cooperative games: coordination games

Mechanism Design (Reverse Game Theory)

Mechanism Design is to answer...

Question

How to design a mechanism/game, toward desired objectives, in strategic settings?

	Participants	Game	Outcome
Game Theory			
Mechanism Design		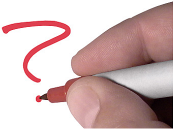	

Mechanism Design (Reverse Game Theory)

Mechanism Design is to answer...

Question

How to design a mechanism/game, toward desired objectives, in strategic settings?

Roger B. Myerson (born March 29, 1951, University of Chicago, US)

- Nobel Prize for economics (2007), for "having laid the foundations of mechanism design theory."
- *Eleven game-theorists have won the economics Nobel Prize*.

Algorithmic Game Theory (AGT)

- **Algorithmic game theory** is an area in the intersection of game theory and algorithm design, whose objective is to design algorithms in strategic environments (wiki) [Nisan et al. 2007].

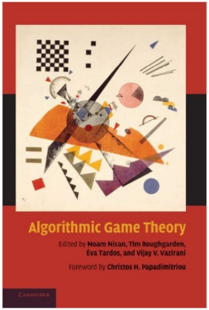

- *Computing in Games*: algorithms for computing equilibria
- *Algorithmic Mechanism Design*: design games that have both good game-theoretical and algorithmic properties
- ...

Algorithmic Game Theory in Artificial Intelligence

- Algorithmic game theory research in AI:
  - Game Playing: computation challenges, AlphaGo, poker
  - Social Choice: preferences aggregation, voting, prediction
  - Mechanism Design: the allocation of scarce resources, ad auctions
- Many IJCAI Computers and Thought Award (outstanding young scientists in artificial intelligence) winners had worked on AGT:
  - Sarit Kraus (1995), Nicholas Jennings (1999), Tuomas Sandholm (2003), Peter Stone (2007), Vincent Conitzer (2011), and Ariel Procaccia (2015)

AGT on Networks

- Algorithmic Game Theory started with Routing Networks

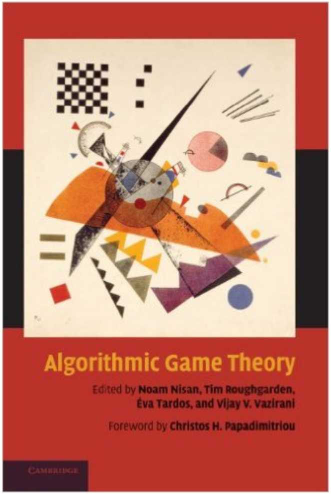

AGT on Networks

- Algorithmic Game Theory started with Routing Networks

AGT on Networks

- Algorithmic Game Theory started with Routing Networks

AGT on Networks

- Another book regarding Game Theory and Networks

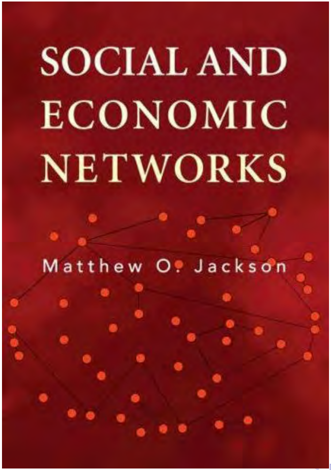

A Simple Mechanism Design Example

Design Goal

How can a house-seller sell her house with the "highest" profit?

A Mechanism Design Example

Design Goal

How can a house-seller sell her house with the "highest" profit?

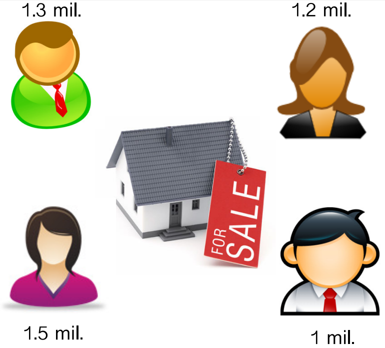

- **Challenge**: the seller doesn't know how much the buyers are willing to pay (their valuations).

A Mechanism Design Example

Design Goal

How can a house-seller sell her house with the "highest" profit?

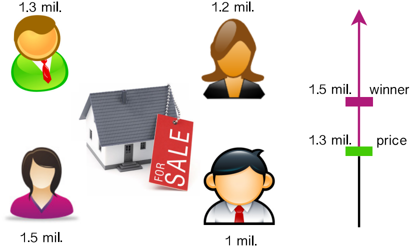

Solution: Second Price Auction (Vickrey Auction/VCG)

- Input: each buyer reports a price/bid to the seller
- Output: the seller decides
  - *allocation*: the agent with the highest price wins.
  - *payment*: the winner pays the second highest price.

A Mechanism Design Example

Design Goal

How can a house-seller sell her house with the "highest" profit?

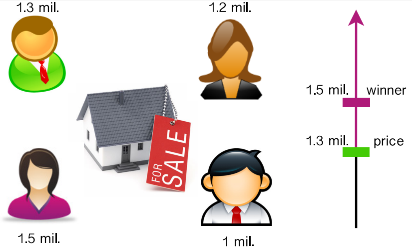

Solution: Second Price Auction (Vickrey Auction/VCG)

Properties:

- Efficient: maximising social welfare
- Truthful: buyers report their valuations truthfully

Is this the BEST the seller can do?

Question

What can the seller do to FURTHER increase her profit?

Is this the BEST the seller can do?

Question

What can the seller do to FURTHER increase her profit?

- estimate a good reserve price [Myerson 1981]
  - *requires a good estimation of buyers' valuations*
- promotions: let more people know/participate in the auction

Is this the BEST the seller can do?

Question

What can the seller do to FURTHER increase her profit?

- estimate a good reserve price [Myerson 1981]
    - *requires a good estimation of buyers' valuations*
- promotions: let more people know/participate in the auction

Recap: Promote a Sale via Social Networks

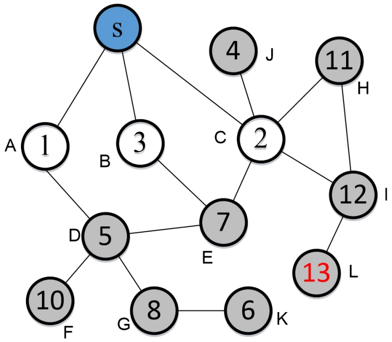

- The seller (blue node) sells one item and has only three connections in the network (A,B,C).
- Each node is a potential buyer and the value is her highest willing payment to buy the item (valuation).
- Profit of applying second price auction without promotion is 2.
- but the highest willing payment of the network is 13.

Traditional Sale Promotions

Traditional sale promotions:

- Promotions in shopping centres
- Keywords based ads via search engines such as Google
- Ads via social media such as WeChat, Facebook, Twitter

Challenge

- The return of these promotions are unpredictable.
- The seller may LOSE from the promotions.

Tackle the Challenge

Build promotion inside the market mechanism such that

1. the promotion will never bring negative utility/revenue to the seller.
2. all buyers who are aware of the sale are incentivized to diffuse the sale information to all her neighbours.

"Diffusion Mechanism Design"

Our Solutions

Information Diffusion Mechanisms

- Bin Li, Dong Hao, Dengji Zhao, Tao Zhou: Mechanism Design in Social Networks. AAAI'17.
- Dengji Zhao, Bin Li, Junping Xu, Dong Hao, Nick Jennings: Selling Multiple Items via Social Networks. AAMAS'18.
- Bin Li, Dong Hao, Dengji Zhao, Tao Zhou: *Customer Sharing in Economic Networks with Costs*. IJCAI-ECAI'18.
- Bin Li, Dong Hao, Dengji Zhao, Makoto Yokoo: *Diffusion and Auction on Graphs*. IJCAI'19.
- Tianyi Zhang, Dengji Zhao, Wen Zhang, Xuming He: *Fixed-price Diffusion Mechanism Design*. CoRR abs/1905.05450 (2019)
- Wen Zhang, Yao Zhang, Dengji Zhao: *Crowdsourcing Data Acquisition via Social Networks*. CoRR abs/1905.05481 (2019)

Information Diffusion Paths

An information diffusion path from the seller to node L:
$s \rightarrow C \rightarrow I \rightarrow L$

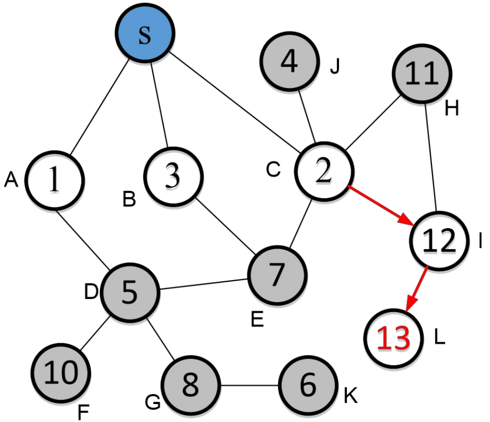

Information Diffusion Paths

An information diffusion path from the seller to node L:
$s \rightarrow C \rightarrow I \rightarrow L$

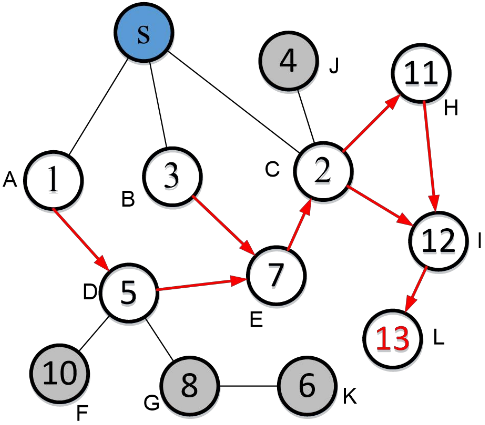

Diffusion Critical Nodes

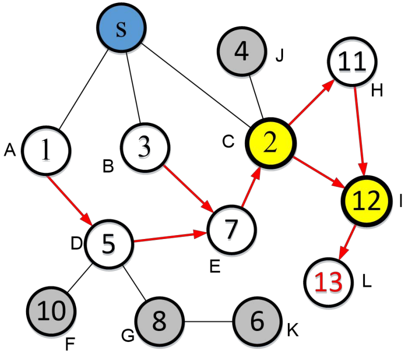

Definition

i is *j*'s diffusion critical node if all the information diffusion paths started from the seller *s* to *j* have to pass *i*.

- nodes C and I are L's only diffusion critical nodes.

Information Diffusion Mechanism [Li et al. AAAI'17]

The payment definition (second-price-like):

- If a buyer or one of her "*diffusion critical children*" gets the item, then the buyer pays the highest bid of the others (without the buyer's participation);
- otherwise, her payment is zero.

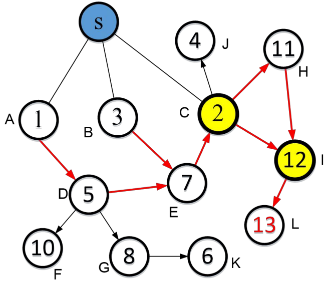

Information Diffusion Mechanism [Li et al. AAAI'17]

The payment definition (second-price-like):

- If a buyer or one of her "*diffusion critical children*" gets the item, then the buyer pays the highest bid of the others (without the buyer's participation);
- otherwise, her payment is zero.

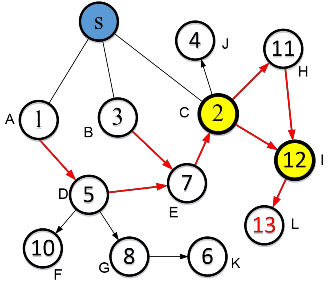

If the item is allocated to *L*, the payments of C, I and L are 10, 11, 12 respectively .

Information Diffusion Mechanism

The allocation definition:

- Identify the node i with the highest bid and the node's diffusion critical node path $P_{c_i} = (c_i^1, c_i^2, ..., i)$.
- Give the item to the first node of P_{c_i}, the node pays to the seller and then decides to whether keep the item or pass it to the next node in P_{c_i}:
  - If the payment of the next node is greater than the bid of the current node, passes it to the next node and receives the payment from the next node; the next node makes a similar decision;
  - otherwise, keep the item.

The Information Diffusion Mechanism

The outcome of the Information Diffusion Mechanism:

- the item is allocated to node I.
- node I pays 11 to C, C pays 10 to the seller.
- the utilities of I, C, the seller are 1, 1, 10.

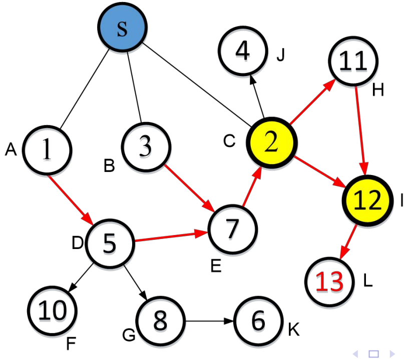

Why Buyers are Happy to Diffuse the Information?

- buyers receive the information earlier have higher priority to win the item (*C* chooses before *I* and *I* chooses before *L*).
- diffuse the information to more buyers will potentially increase their reward (if *C* does not invite *H*, her utility is 0).

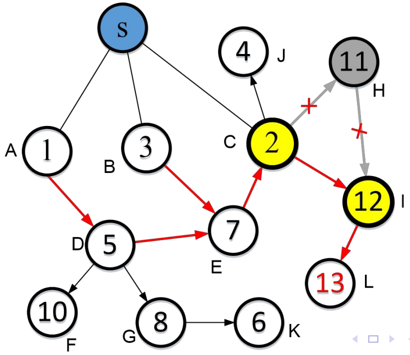

Properties of the Information Diffusion Mechanism

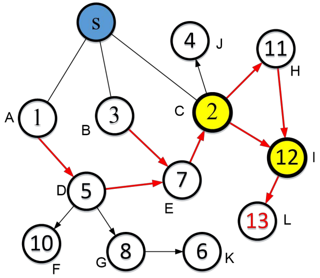

- Truthful: report true valuation and diffuse the sale information to all her neighbours is a dominate strategy.
- Individually Rational: no buyer will receive a negative utility to join the mechanism.
- Seller's Revenue Improved: the seller's revenue is non-negative and is $\geq$ that of the VCG without diffusion.

What Next?

- Diffusion mechanisms for combinatorial exchanges
- Diffusion with costs and delays
- Network structure based revenue analysis
- Applications/implementations in the existing social networks
- Other mechanisms to further improve the revenue and/or the efficiency

Diffusion Mechanisms for Combinatorial Exchanges

Challenge

How to generalise the mechanism to combinatorial settings?

Diffusion Mechanisms for Combinatorial Exchanges

Consider the following simple setting:

- A seller sells three units of one commodity, e.g. MacBook computers.
- Each buyer has a diminishing marginal utility for consuming the goods.

Diffusion Mechanisms for Combinatorial Exchanges

Consider the following simple setting:

- A seller sells three units of one commodity, e.g. MacBook computers.
- Each buyer has a diminishing marginal utility for consuming the goods.

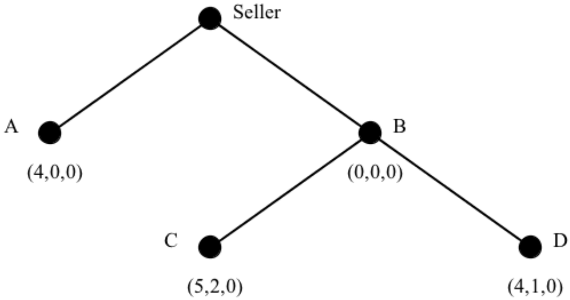

Diffusion Mechanisms for Combinatorial Exchanges

If we simply apply our information diffusion mechanism:

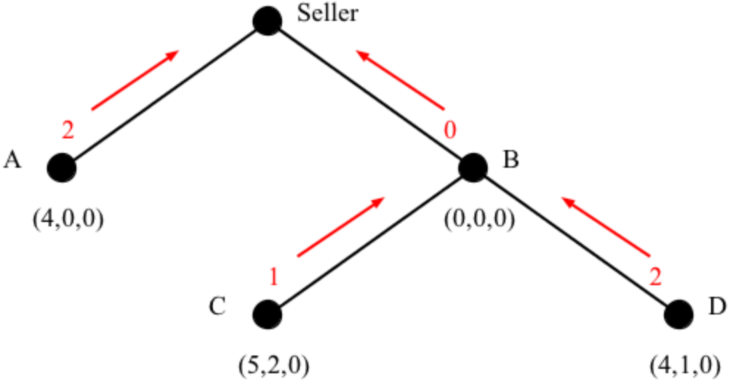

Diffusion Mechanisms for Combinatorial Exchanges

What if buyer B's valuation is $(3, 0, 0)$?

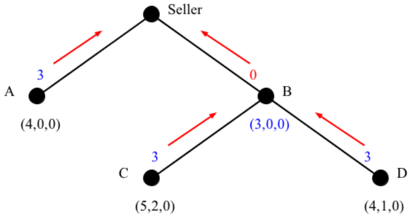

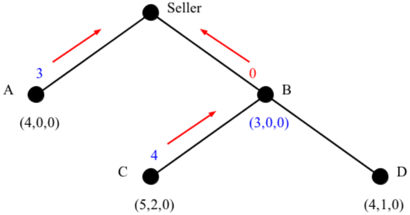

Diffusion Mechanisms for Combinatorial Exchanges

What if buyer D's valuation is $(4, 2.5, 0)$?

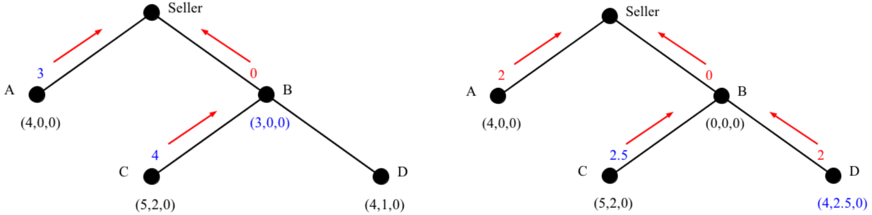

Diffusion Mechanisms for Combinatorial Exchanges

Challenge

There is a very complex **Decision Making** at each node!!!

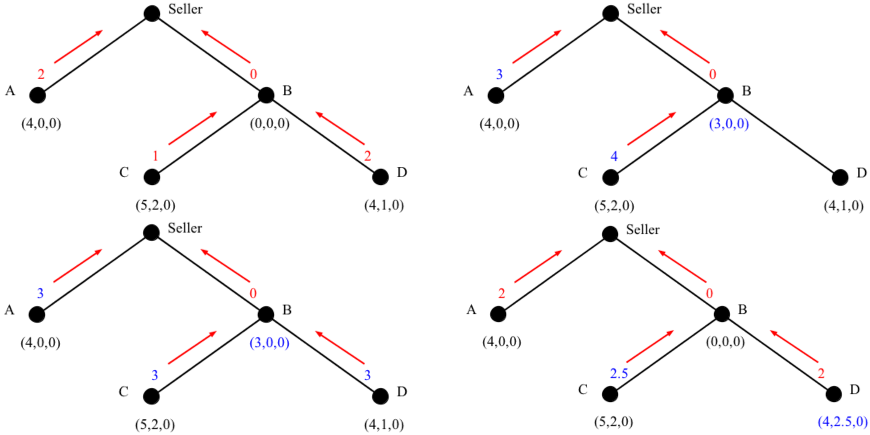

Why is it so complex when there are multiple items?

To achieve truthfulness:

- The mechanism has to maximise each node's utility under truthful reporting.
- Each node's payment should not depend on her valuation.

The complexity we had:

- A node can influence her received payments by controlling the items passed to her children.
- A node can influence the payments of her peers, without changing her own allocation and payments.
- This leads to a decision loop (very complex optimization) and may not able to maximise everyone's utility.

Reduce the Complexity

The Main Idea

A node CANNOT influence the payments she receives by controlling the items passed to her children.

Simplify the decision complexity we had:

- ~~A node can influence her received payments by controlling the items passed to her children.~~
- A node can influence the payments of her peers, without changing her own allocation and payments.
- ~~This leads to a decision loop and may not able to maximise everyone's utility.~~

Solution Example: Sells Multiple Homogeneous Items

Selling Multiple Items via Social Networks [Zhao et al. AAMAS'18]

- generalises the result from [Li et al. 2017];
- agent *i*'s reward/payment doesn't depends on how many of *i*'s children received items;
- agent pays to the seller directly rather than to their parent;

The Generalised Setting

- A seller sells $\mathcal{K} \geq 1$ homogeneous items;
- each buyer requires at most one item (single-unit demand);
- the rest is the same as [Li et al. 2017].

The Generalised Diffusion Mechanism

Consider $\mathcal{K} = 5$:

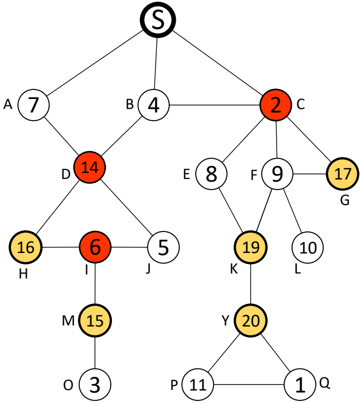

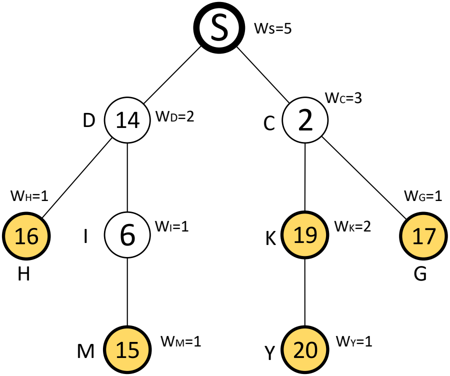

The Generalised Diffusion Mechanism

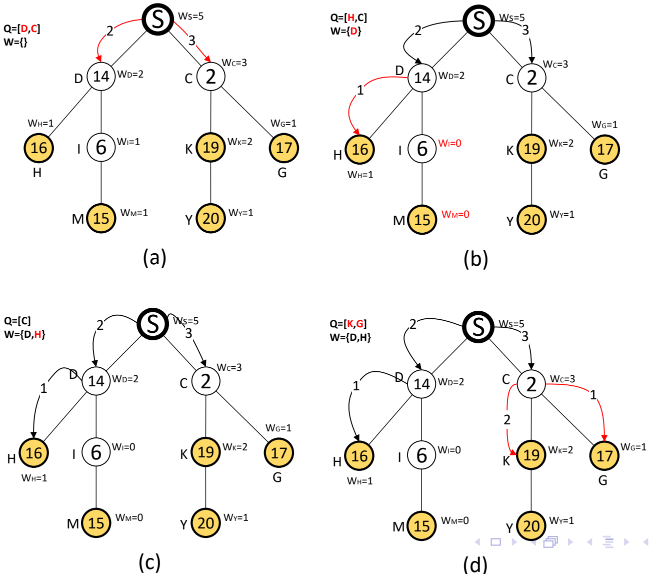

The Generalised Diffusion Mechanism

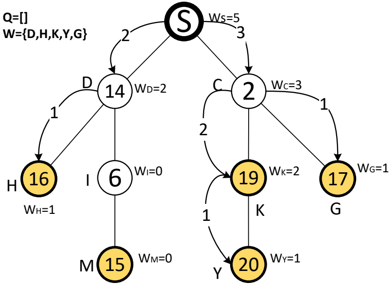

(i)

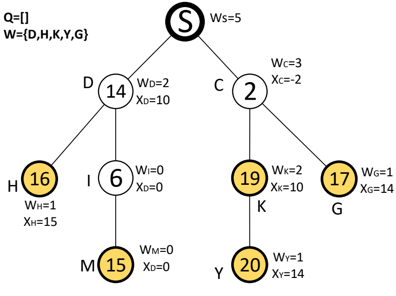

(ii)

The Allocation Policy of the Generalisation

Node/buyer i receives one item if and only if

1. the top $\mathcal{K}$-highest valued children of i (and their parents, who are also i's children) do not participate
2. and i wins under the efficient allocation with their absence

given that all i's (critical) parents' allocation is determined and fixed.

The Payment Policy of the Generalisation

Node i's utility is the social welfare difference of the efficient allocation between

1. the top $\mathcal{K}$-highest valued children of i (and their parents, who are also i's children) do not participate (guarantees that i's payment does not depend on how many items i's children get)
2. and i (and all her children) does not participate

Formally, i's payment is:

$$\begin{cases} \mathcal{SW}_{-D_i} - (\mathcal{SW}_{-\mathcal{C}_i^{\mathcal{K}}} - v_i') & \text{if } i \in W, \\ \mathcal{SW}_{-D_i} - \mathcal{SW}_{-\mathcal{C}_i^{\mathcal{K}}} & \text{if } i \in \bigcup_{j \in W} \mathcal{P}_j(\theta') \setminus W, \\ 0 & \text{otherwise.} \end{cases}$$

where W is the set of nodes each of whom received one item.

Properties of the Generalisation

- Truthful: report true valuation and diffuse the sale information to all her neighbours is a dominate strategy for each node.
- Individually Rational: no node will receive a negative utility to join the mechanism.
- Seller's Revenue Improved: the seller's revenue is non-negative and is $\geq$ that of the VCG without diffusion.

What happens when $\mathcal{K} = 1$?

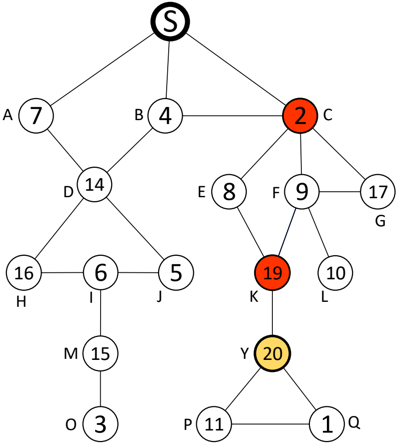

Open Questions

- More general settings
  - characterize truthful diffusion mechanisms, revenue monotonicity is the key?
- When there is a diffusion cost
  - how to guarantee the seller will not lose?
- Privacy concern and the seller's strategies
  - the seller discovery the whole network and she may cheat as well!
- False-name manipulations
  - a node may create multiple ids as her neighbours to gain more payment?
- many more...

Fixed-price Diffusion Mechanism

Fixed-price mechanism is another way to go:

- simple and easy to design/compute
- does not require buyers' valuation reports

Challenges:

- why buyers should invite their neighbours to join the mechanism?
- how to define the fixed-prices?

Exp. Rev. for Selling One Item under Tree

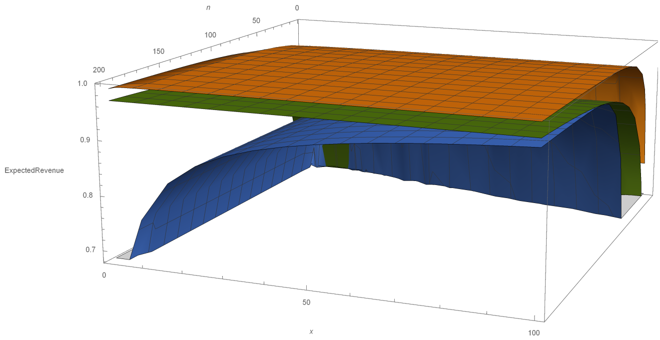

- A tree network with random valuation distribution on [0, 1], x is the number of neighbors of the seller, n is the total number of buyers. Brown is the optimal exp. rev. with fixed-price, Green is the rev. of IDM, and Blue is the exp. rev. with a fixed-price diffusion mechanism.

Diffusion Mechanism Design for Task Allocation

Resource allocation vs task allocation:

1. task requires more participants' contribution (collaboration)
2. but participants' contribution may conflict with each other (competition)

An Example: Crowdsourcing Data Acquisition

- a requester is collecting data from the crowd
- more participants gives richer dataset
- participants' contribution depends on the quality of their provided data
- if two participants offer the same data, how to calculate their contribution?

An Example: Crowdsourcing Data Acquisition

- a requester is collecting data from the crowd
- more participants gives richer dataset
- participants' contribution depends on the quality of their provided data
- if two participants offer the same data, how to calculate their contribution?

Shapley Value?

Shapley Value

The problem of the Shapley value:

- two participants offer the same data will share the same Shapley value
  - the Shapley value is doubled if one of them didn't participate

Solution: Layered Shapley Value

- participants are layered
- the Shapley value is calculated for each lower layer first
- the calculation for higher layer assumes that lower layers' participants are always in the coalition

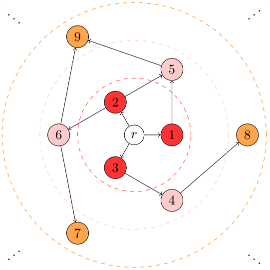

$$\hat{\phi}_i = \sum_{S \subseteq L_{l_i} - \{i\}} \frac{|S|!(|L_{l_i}| - |S| - 1)!}{|L_{l_i}|!} \cdot \left(v\left(D'_{L^*_{l_i - 1} \cup S \cup \{i\}}\right) - v\left(D'_{L^*_{l_i - 1} \cup S}\right)\right)$$

Solution: Layered Shapley Value

$$\hat{\phi}_i = \sum_{S \subseteq L_{l_i} - \{i\}} \frac{|S|!(|L_{l_i}| - |S| - 1)!}{|L_{l_i}|!} \cdot \left(v\left(D'_{L^*_{l_i-1} \cup S \cup \{i\}} \right) - v\left(D'_{L^*_{l_i-1} \cup S} \right) \right)$$

Properties:

- participants are incentivized to invite more participants (new participants do not compete with them)
- the requester does not need to pay for redundant data

Summary

Mechanism Design Powered by Social Interactions

- Diffusion Mechanism for Resource Allocation (competitive environment)
  - for selling single item
  - for selling multiple items
- Diffusion Mechanism for Task Allocation (both competitive and collaborative)
  - crowdsourcing data acquisition

http://dengji-zhao.net/ijcai19.html

Related Work

1. Mahajan, V. and R.A. Peterson, Models for innovation diffusion. Beverly Hills California Sage Publications, 1985.
2. Kempe, D. and J. Kleinberg, Eva Tardos: Maximizing the spread of influence through a social network. Kdd, 2003: p. 137–146.
3. Jackson, M.O., Social and Economic Networks. 2008: Princeton University Press. 44-74(31).
4. Singer, Y. How to win friends and influence people, truthfully:influence maximization mechanisms for social networks. in ACM International Conference on Web Search and Data Mining. 2012.
5. Leskovec, J., L.A. Adamic, and B.A. Huberman, The Dynamics of Viral Marketing. 2005. 1(1): p. 228-237.
6. Emek, Y., et al. Mechanisms for multi-level marketing. in ACM Conference on Electronic Commerce. 2011.
7. Iribarren, J.L. and E. Moro, Branching dynamics of viral information spreading. Physical Review E Statistical Nonlinear & Soft Matter Physics, 2011. 84(4 Pt 2): p. 046116.

Related Work

8. Condorelli, D., A. Galeotti, and V. Skreta, Selling Through Referrals. Working Papers, 2013.
9. Raghavan, P. Query Incentive Networks. in Asian Computing Science Conference. 2005.
10. Babaioff, M., et al., On Bitcoin and red balloons. Acm Sigecom Exchanges, 2011. 10(3): p. 5-9.
11. Pickard, G., et al., Time-Critical Social Mobilization. Science, 2011. 334(6055): p. 509-12.
12. Cebrian, M., et al. Finding red balloons with split contracts:robustness to individuals' selfishness. in Forty-Fourth ACM Symposium on Theory of Computing. 2012.
13. Wei, H.Y. and R.D. Gitlin, Incentive Mechanism Design for Selfish Hybrid Wireless Relay Networks. Mobile Networks & Applications, 2005. 10(6): p. 929-937.
14. Li, C., B. Yu, and K. Sycara. An incentive mechanism for message relaying in unstructured peer-to-peer systems. 2007.