Mechanism Design for Resource Critical Task Execution via Crowdsourcing[1]

Swaprava Nath

Department of Computer Science and Automation
Indian Institute of Science, Bangalore

Joint Work with Pankaj Dayama, Dinesh Garg, Y. Narahari, and James Zou

Topics in Game Theory, Fall 2012

October 4, 2012

[1] To appear in Workshop on Internet and Network Economics (WINE), Liverpool, UK, December 9 - 12, 2012

Outline of Talk

1. Motivation
   - The DARPA Network Challenge
2. Potential Dangers in Strategic Setting
   - Sybil Attack
   - Node Collapse Attack
   - Design Desiderata
3. Main Results
   - Impossibility and Possibility Results
   - Approximate Versions of Desirable Properties
   - Incentives for Task Forwarding and Execution
   - Cost Critical Setting
   - Time Critical Setting
4. Summary and Future Work

Outline of Talk

1. Motivation
   - The DARPA Network Challenge
2. Potential Dangers in Strategic Setting
   - Sybil Attack
   - Node Collapse Attack
   - Design Desiderata
3. Main Results
   - Impossibility and Possibility Results
   - Approximate Versions of Desirable Properties
   - Incentives for Task Forwarding and Execution
   - Cost Critical Setting
   - Time Critical Setting
4. Summary and Future Work

DARPA Network Challenge, 2009

- The challenge is to identify the locations of 10 balloons
- Whoever locates all of them in the shortest time will get a reward of \$40,000
- Balloons are spread across the continental USA
    - Impossible for any individual to travel to all the places
    - Time-critical competition
- **Crowdsourcing** with some help from modern technology is a natural approach

The Winning Solution: MIT Media Lab

- Winning solution: MIT Media Lab [2]
- Efficiently harnesses the collective intelligence and collaborative effort of a social network
- Incentive scheme is a *geometric* reward mechanism, decreasing from leaf to root

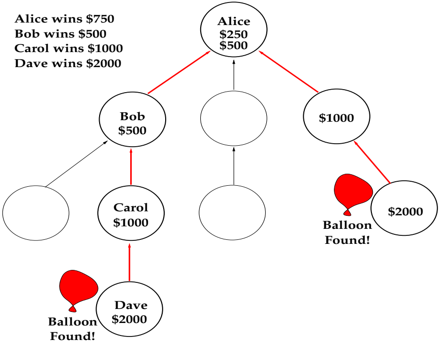

[2] G. Pickard, W. Pan, I. Rahwan, M. Cebrian, R. Crane, A. Madan, and A. Pentland. Time-Critical Social Mobilization. Science, 334(6055):509-512, October 2011

Outline of Talk

1. Motivation
   - The DARPA Network Challenge
2. Potential Dangers in Strategic Setting
   - Sybil Attack
   - Node Collapse Attack
   - Design Desiderata
3. Main Results
   - Impossibility and Possibility Results
   - Approximate Versions of Desirable Properties
   - Incentives for Task Forwarding and Execution
   - Cost Critical Setting
   - Time Critical Setting
4. Summary and Future Work

Potential Dangers in Strategic Setting

- We are considering **Atomic Tasks**
  Indivisible tasks, accomplished by a single individual

Potential Dangers in Strategic Setting

- We are considering **Atomic Tasks**
  Indivisible tasks, accomplished by a single individual

- Human participants of the social network are **strategic**.
- Can manipulate the mechanism in order to maximize their own **payoff**
- Two major problems with the incentive mechanism: **sybil attack** and **node collapse attack**.

Outline of Talk

Sybil Attack

Nodes can multiply their identities by creating clones or fake nodes below themselves in the referral tree. Example: Carol can create two fake nodes to earn $750 more in the MIT scheme

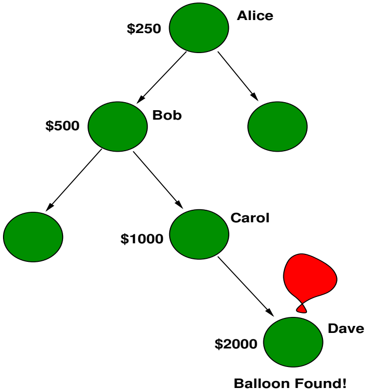

Sybil Attack

Nodes can multiply their identities by creating clones or fake nodes below themselves in the referral tree. Example: Carol can create two fake nodes to earn $750 more in the MIT scheme

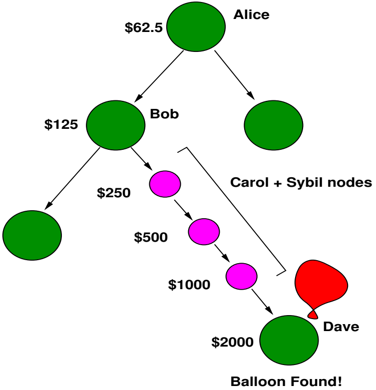

Sybil attack is undesirable because,

- Increases the expenditure of the task owner, as the sybils are getting paid.
- Reduces the reward of the ancestors of the sybil-creating nodes.

Outline of Talk

1. Motivation
   - The DARPA Network Challenge
2. Potential Dangers in Strategic Setting
   - Sybil Attack
   - Node Collapse Attack
   - Design Desiderata
3. Main Results
   - Impossibility and Possibility Results
   - Approximate Versions of Desirable Properties
   - Incentives for Task Forwarding and Execution
   - Cost Critical Setting
   - Time Critical Setting
4. Summary and Future Work

Node Collapse Attack

- To combat the sybil attack, one can think of a naïve reward scheme.
- TOP-DOWN: if the number of nodes in the winning chain (call this 'length') is t, node at depth d gets $\$4000/2^{d+t}$.
- This could lead to a different problem: node collapse problem

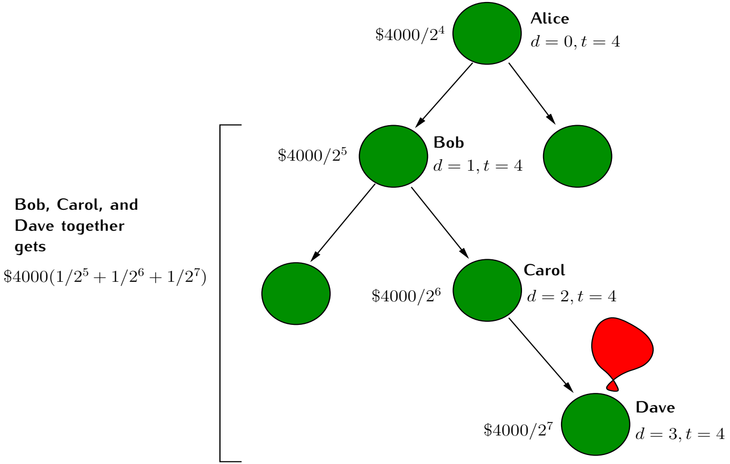

Node Collapse Attack

- To combat the sybil attack, one can think of a naïve reward scheme.
- TOP-DOWN: if the number of nodes in the winning chain (call this 'length') is $\mathtt{t}$, node at depth $\mathtt{d}$ gets $\$4000/2^{\mathtt{d}+\mathtt{t}}$.
- This could lead to a different problem: node collapse problem

Bob, Carol, and Dave together gets
$\$4000/2^3$
which is larger than earlier
$\$4000(1/2^5 + 1/2^6 + 1/2^7)$

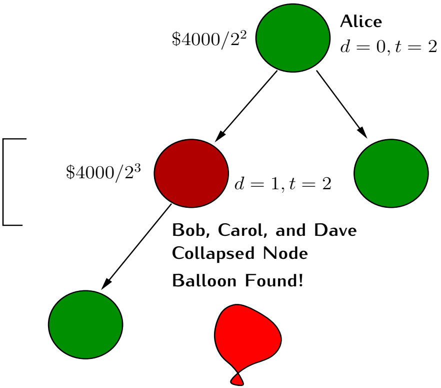

Node Collapse Attack (Contd.)

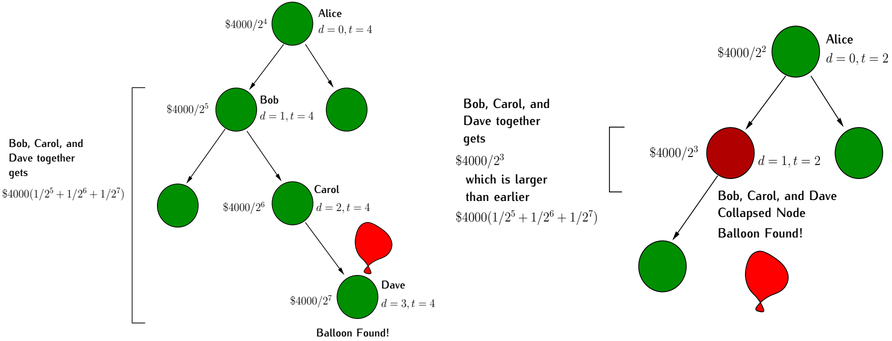

Node collapse is undesirable:

- Costs more to the social planner
- Sharing of this surplus could lead to bargaining among the agents
- Hides the structure of the actual network, which could otherwise be used for different purposes.

Outline of Talk

Desirable Properties

Definition (Downstream Sybil-Proofness (DSP))

Given the depth k of a node in a recruitment tree, a reward mechanism R is called *downstream sybilproof*, if the node cannot gain by adding fake nodes below itself in the current subtree. Formally,

$$R(k,t) \geqslant \sum_{i=0}^{n} R(k+i, t+n) \quad \forall k \leqslant t, \forall t, n.$$

Definition (Collapse-Proofness (CP))

Given a depth k in a winning chain, a reward mechanism R is called *collapse-proof*, if the user in the subchain of length p lying beneath k collectively cannot gain by collapsing to depth k. Mathematically,

$$\sum_{i=0}^{p} R(k+i, t) \geqslant R(k, t-p) \quad \forall k+p \leqslant t, \forall t.$$

Desirable Properties

Definition (Downstream Sybil-Proofness (DSP))

Given the depth k of a node in a recruitment tree, a reward mechanism R is called *downstream sybilproof*, if the node cannot gain by adding fake nodes below itself in the current subtree. Formally,

$$R(k,t) \geqslant \sum_{i=0}^{n} R(k+i, t+n) \quad \forall k \leqslant t, \forall t, n.$$

Definition (Collapse-Proofness (CP))

Given a depth k in a winning chain, a reward mechanism R is called *collapse-proof*, if the user in the subchain of length p lying beneath k collectively cannot gain by collapsing to depth k. Mathematically,

$$\sum_{i=0}^{p} R(k+i, t) \geqslant R(k, t-p) \quad \forall k+p \leqslant t, \forall t.$$

- **This asks for a Dominant Strategy implementation**

Desirable Properties (Contd.)

Definition (Strict Contribution Rationality (SCR))

This ensures a positive payoff to the nodes belonging to the winning chain. For all $t \geqslant 1$:

$$R(k, t) > 0, \quad \forall k \leqslant t, \text{ if } t \text{ is the length of the winning chain.}$$

Definition (Weak Contribution Rationality (WCR))

This ensures a non-negative payoff to the nodes in the winning chain. For all $t \geqslant 1$:

$$R(k, t) \geqslant 0, \quad \forall k \leqslant t - 1, \text{if } t \text{ is the length of the winning chain.}$$
$$R(t, t) > 0, \qquad \text{winner gets positive reward.}$$

Desirable Properties (Contd.)

Definition (Budget Balance (BB))

Suppose the maximum budget allocated by the planner for executing a task is R_{max}. Then, a mechanism R is budget balanced if,

$$\sum_{k=1}^{t} R(k, t) \leqslant R_{max}, \quad \forall t.$$

Outline of Talk

Outline of Talk

Impossibility and Possibility Results

Not all these properties are simultaneously satisfiable.

Theorem (Impossibility Result)

For $\texttt{t} \geqslant 3$*, no reward mechanism can simultaneously satisfy DSP, SCR, and CP.*

Impossibility and Possibility Results

Not all these properties are simultaneously satisfiable.

Theorem (Impossibility Result)

For $\mathsf{t} \geqslant 3$*, no reward mechanism can simultaneously satisfy DSP, SCR, and CP.*

Theorem (Possibility Result A)

For $\mathsf{t} \geqslant 3$*, a mechanism satisfies DSP, WCR, CP, and BB iff it is a* Winner Takes All *(WTA) mechanism. A reward mechanism* R *is called WTA if* $\mathsf{R}_{\max} \geqslant \mathsf{R}(\mathsf{t}, \mathsf{t}) > 0$*, and* $\mathsf{R}(\mathsf{k}, \mathsf{t}) = 0, \ \forall \mathsf{k} < \mathsf{t}$.

Outline of Talk

Approximate Sybil-proofness

Potential way outs:

- **Relax the equilibrium**: Nash implementation[3]
- **Relax the properties**: equilibrium in dominant strategies (this talk)

[3] M. Babaioff, S. Dobzinski, S. Oren, and A. Zohar. On Bitcoin and Red Balloons. In Proceedings of ACM Electronic Commerce, 2012.

Approximate Sybil-proofness

Potential way outs:

- **Relax the equilibrium**: Nash implementation[3]
- **Relax the properties**: equilibrium in dominant strategies (this talk)

Definition (ϵ-Downstream Sybil-Proofness (ϵ-DSP))

A reward mechanism R is called ϵ - *DSP*, if no node can gain by more than a factor of $(1+\epsilon)$ by adding fake nodes below herself in the current subtree. Mathematically,

$$(1+\epsilon) \cdot R(k,t) \geqslant \sum_{i=0}^{n} R(k+i, t+n) \quad \forall k \leqslant t, \forall t, n.$$

[3] M. Babaioff, S. Dobzinski, S. Oren, and A. Zohar. On Bitcoin and Red Balloons. In Proceedings of ACM Electronic Commerce, 2012.

A Possibility Result

Question: Can we design mechanisms with limited sybil attacks?

A Possibility Result

Question: Can we design mechanisms with limited sybil attacks?
Answer: Yes!

Theorem (Possibility Result B)

For all $\epsilon > 0$*, there exists a mechanism that is* ϵ*-DSP, CP, BB, and SCR.*

A Possibility Result

Question: Can we design mechanisms with limited sybil attacks?
Answer: Yes!

Theorem (Possibility Result B)

For all $\epsilon > 0$, *there exists a mechanism that is* ϵ*-DSP, CP, BB, and SCR.*

- $R(t,t) = (1-\delta).R_{max} \quad \forall t$ where $\delta \leqslant \frac{\epsilon}{1+\epsilon}$
- $R(k,t) = \delta.R(k+1,t) \;\; \forall \;\; k,t$

The Mechanism Design Space

- **Theorem 1:** Impossibility Result
  DSP, CP, and SCR is impossible

The Mechanism Design Space

- **Theorem 1:** Impossibility Result
  DSP, CP, and SCR is impossible
  Corollary: DSP, CP, SCR, and BB is impossible

The Mechanism Design Space

- **Theorem 1:** Impossibility Result
  DSP, CP, and SCR is impossible
  Corollary: DSP, CP, SCR, and BB is impossible

- **Theorem 2:** Possibility Result A
  DSP, CP, WCR, and BB $\Leftrightarrow$ WTA mechanism

The Mechanism Design Space

- **Theorem 1:** Impossibility Result
  DSP, CP, and SCR is impossible
  Corollary: DSP, CP, SCR, and BB is impossible

- **Theorem 2:** Possibility Result A
  DSP, CP, WCR, and BB $\Leftrightarrow$ WTA mechanism

- **Theorem 3:** Possibility Result B
  ϵ-DSP, CP, BB, and SCR is possible, for all $\epsilon > 0$

Outline of Talk

Incentive for Task Forwarding

- Not all mechanisms in the non-empty space would be interesting
- Leads us to define two additional fairness criteria.

Incentive for Task Forwarding

Definition (δ - Strict Contribution Rationality (δ-SCR))

This property ensures that a node in the winning chain gets at least $\delta \in (0, 1)$ fraction of her successor. Also the winner gets a positive reward. For all $t \geqslant 1$,

$$R(k, t) \geqslant \delta R(k + 1, t), \forall k \leqslant t - 1, \text{ if } t \text{ is the length of the winning chain.}$$
$$R(t, t) > 0, \qquad \text{winner gets positive reward.}$$

Incentive for Task Execution

Definition (Winner's γ Security, γ-SEC)

This property ensures that payoff to the winning node is at least γ fraction $(0 < \gamma < 1)$ of the total available budget.

$$R(t, t) \geqslant \gamma \cdot R_{\max}, \quad t \text{ is the length of the winning chain}$$

Incentive for Task Execution

Definition (Winner's γ Security, γ-SEC)

This property ensures that payoff to the winning node is at least γ fraction $(0 < \gamma < 1)$ of the total available budget.

$$R(t, t) \geqslant \gamma \cdot R_{max}, \quad t \text{ is the length of the winning chain}$$

- Properties ϵ-DSP, δ-SCR, and γ-SEC, parametrized by ϵ, δ, and γ, ensure fairness to the participants and limit the spread of fake nodes.
- We characterize the space of mechanisms that satisfy this set of properties.

Outline of Talk

Cost Critical Setting

Goal: Accomplishing the task at minimum cost
Note: γ-SEC property is essential, otherwise the solution would be all-zero.

Definition (MINCOST over $\mathscr{C}$)

A reward mechanism R is called *MINCOST* over a class of mechanisms $\mathscr{C}$, if it minimizes the total reward distributed to the participants in the winning chain. That is, R is *MINCOST* over $\mathscr{C}$, if

$$\mathsf{R} \in \arg\min_{\mathsf{R}' \in \mathscr{C}} \sum_{k=1}^{t} \mathsf{R}'(k, t), \quad \forall t.$$

A Characterization Theorem

Let us define, $\mathscr{E} = \{(\delta, \epsilon, \gamma) : \delta \leqslant \min\{1 - \gamma, \frac{\epsilon}{1+\epsilon}\}\}$, a technical condition on the parameters

Theorem (Characterization of Cost Critical Setting)

If $(\delta, \epsilon, \gamma) \in \mathscr{E}$, a mechanism is `MINCOST` *over the class of mechanisms satisfying ϵ-DSP, δ-SCR, γ-SEC, and BB iff it is (γ, δ)-GEOM.*

A Characterization Theorem

Let us define, $\mathscr{E} = \{(\delta, \epsilon, \gamma) : \delta \leqslant \min\{1-\gamma, \frac{\epsilon}{1+\epsilon}\}\}$, a technical condition on the parameters

Theorem (Characterization of Cost Critical Setting)

If $(\delta, \epsilon, \gamma) \in \mathscr{E}$, a mechanism is `MINCOST` *over the class of mechanisms satisfying ϵ-DSP, δ-SCR, γ-SEC, and BB iff it is (γ, δ)-GEOM.*

(γ, δ)-Geometric Mechanism ((γ, δ)-GEOM)

This mechanism gives γ fraction of the total reward to the winner and geometrically decreases the rewards from leaf towards root by a factor δ. For all t,

$$R(t, t) = \gamma \cdot R_{\max}$$
$$R(k, t) = \delta^{t-k} \cdot \gamma R_{\max}, \ k \leqslant t-1$$

Graphical Illustration

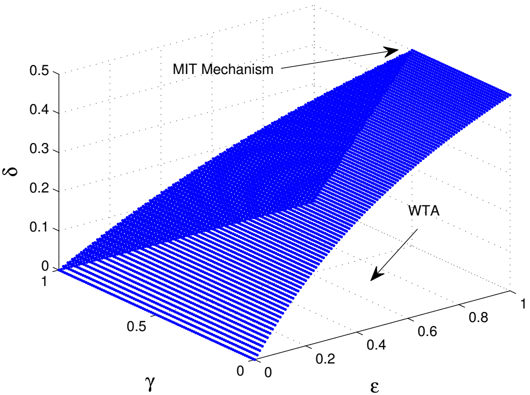

The set of $(\delta, \epsilon, \gamma)$ tuples, given by $\mathscr{E}$, for which the *MINCOST* mechanism is the (γ, δ)-GEOM mechanism, is the space below the shaded region. MIT mechanism $(\epsilon = 1, \delta = 0.5, \gamma = 0.5)$ and the WTA mechanism ($\delta = 0$, the floor of the space in the figure above) are special cases.

Outline of Talk

Time Critical Setting

Goal: Accomplishing the task at the minimum time. So, the entire budget $\mathsf{R}_{\max}$ can be exhausted to encourage faster *task execution* and *propagation.*

Definition (MAXLEAF over $\mathscr{C}$)

A reward mechanism R is called ***MAXLEAF*** over a class of mechanisms $\mathscr{C}$, if it maximizes the reward of the leaf node in the winning chain. That is, R is ***MAXLEAF*** over $\mathscr{C}$, if

$$\mathsf{R} \in \arg\max_{\mathsf{R}' \in \mathscr{C}} \mathsf{R}'(\mathsf{t}, \mathsf{t}), \quad \forall \mathsf{t}.$$

A Characterization Theorem

Theorem (A Characterization for Time Critical Setting)

If $\delta \leqslant \frac{\epsilon}{1+\epsilon}$, a mechanism is **`MAXLEAF`** *over the class of mechanisms satisfying ϵ-DSP, δ-SCR, and BB iff it is δ-GEOM mechanism.*

A Characterization Theorem

Theorem (A Characterization for Time Critical Setting)

*If $\delta \leqslant \frac{\epsilon}{1+\epsilon}$, a mechanism is **MAXLEAF** over the class of mechanisms satisfying ϵ-DSP, δ-SCR, and BB iff it is δ-GEOM mechanism.*

δ-Geometric mechanism (δ-GEOM)

This mechanism gives $\frac{1-\delta}{1-\delta^t}$ fraction of the total reward to the winner and geometrically decreases the rewards towards root with the factor δ; t is the length of the winning chain.

$$R(t,t) = \frac{1-\delta}{1-\delta^t} \cdot R_{\max}$$

$$R(k,t) = \delta \cdot R(k+1,t) = \delta^{t-k} \cdot R(t,t), \ \ k \leqslant t-1$$

A Characterization Theorem

Theorem (A Characterization for Time Critical Setting)

*If $\delta \leqslant \frac{\epsilon}{1+\epsilon}$, a mechanism is **MAXLEAF** over the class of mechanisms satisfying ϵ-DSP, δ-SCR, and BB iff it is δ-GEOM mechanism.*

δ-Geometric mechanism (δ-GEOM)

This mechanism gives $\frac{1-\delta}{1-\delta^t}$ fraction of the total reward to the winner and geometrically decreases the rewards towards root with the factor δ; t is the length of the winning chain.

$$R(t,t) = \frac{1-\delta}{1-\delta^t} \cdot R_{\max}$$

$$R(k,t) = \delta \cdot R(k+1,t) = \delta^{t-k} \cdot R(t,t), \ k \leqslant t-1$$

In addition:

- **Both (γ, δ)-GEOM and δ-GEOM are CP**

Outline of Talk

Summary and Future Work

Summary: The major contributions of this paper are

- Introducing the concept of **Collapse-Proofness**
- Exhibiting the conflict among the desirable properties
- Proposing an *approximate* **Dominant Strategy Implementation**
- Presenting a **Resource-critical Optimization** technique

Future work:

- Investigating tightness of the characterization results
- Approximating the CP property
- Extension to non-atomic tasks
- Efficiently fusing information

Thank you!