CS 284r: Incentives and Information in Networks Fall 2013

Lecture 5 — Wednesday, September 18th, 2013

Prof. Yaron Singer *Scribe: Korey Tucker and Kendall Sherman*

1 Overview

This class covers the paper "Finding Red Balloons with Split Contracts: Robustness to Individuals' Selfishness". Previous work by Kleinberg described fixed-payment contracts, in which a node strategically chooses how much to offer its child. However, the MIT Media Laboratory team that won the DARPA Network Challenge utilized a split-contract scheme. This paper models such a contract, and is able to determine that

- Their model can explain recruiting of new nodes (doesn't assume the network is pre-built)
- Split contracts are robust to individual selfishness
- The investment required to discover the information is linear in the expected distance to the information as long as $b > 1$ (logarithmic in the rarity of the answer)
- When $b = 1$, the investment required to find the answer is $O(n^2)$

Unlike Kleinberg and Raghavan's "Query Incentive Network" paper, which only shows a reward logarithmic in the rarity when the branching factor is greater than 2, this paper shows it whenever the branching factor is greater than 1.

2 DARPA Red Balloon Challenge

The Defense Advanced Research Projects Agency designed a "Network Challenge" to determine how quickly they could find a rare answer with a limited budget. Some strategies utilized were charity, celebrity status, etc. However, the MIT team split the $40,000 prize into 10 $4,000 prizes for each balloon. The first person who discovered the balloon received $2,000. The person who recruited that person received $1,000. The person who'd recruited that person received $500, and so on. Any leftover money was given to charity. The MIT team was able to determine the location of all ten red balloons within 10 hours, even though DARPA expected it would take up to a week.

3 Setup

Every node can offer a p-split contract to its children. This means that if a child has the answer, it receives a payment of r^* but has to return a fraction p of that reward to its parent. The child who finds the answer receives final reward $r^*(1-p) - 1$ (since there is a fixed cost to send it up). Each node then has function $f_v(p(v)) = p_{v+1}$. p is some number on the discretized range $(0, p_1]$.

The root, then, receives payoff $r^* \prod_{i<\text{level with answer}} p_i$. The root always selects the closest answer for payoff. This makes sense, since it will in general have to pay less, and best models the MIT team's strategy.

For any constant $\epsilon > 0$, an investment of $r^* = O(\log n)$ is required to retreive the answer with probability $1-\epsilon-\zeta$ (where ζ is the extinction probability of the process). The probability generation function is denoted by ψ and the branching factor of the process is $b = \psi'(1)$. A given node has the answer with probability $\frac{1}{n}$.

Let $h(\epsilon, n)$ be the minimum level i such that the probability that no node has the answer (ϕ_i) is less than ϵ. Then, we know that r^* must be high enough to reach the level h. We merely have to prove that the reward is linear in h to prove our final result.

4 Nash Equilibrium

A set of strategies f is a NE if a node cannot increase its payoff by deviating from its own $f(p)$. We can treat each node on the same level as symmetric - they all face the same expected payoffs. A node only knows about the payoffs offered to it, but its payoff is only impacted by the strategies of its children.

Each node wants its query to be propagated only to a certain level. We define a strategy p_i^l as the split offered to level i that will propagate the query down to level l. If it propagates too far, then the node will end up farther from the answer than it wants, and it will get a lower reward but be forced to forward the answer! If a p-split contract causes propagation down i levels, and a p'-split contract propagates the query down $i+1$ levels, then $p > p'$. Therefore, the p-splits offered increase as you go further down the tree. We show this intuitively for just the first two levels.

To describe the N.E., we need to inductively define a sequence of functions $e_i^l(\rho)$ $(1 \le l \le h, 0 \le i \le h-l)$ and thresholds ρ_i^l $(1 \le \ l \le h, 1 \le i \le h-l+1)$:

$$e_0^l(\rho) = 0,\ \rho_1^l = 1 - \frac{1}{r^*},$$

$$e_i^l(\rho) = \sum_{j=1}^{i} \gamma_j^l [r^*(1-\rho) \prod_{t=0}^{j-1} \rho_{i-t}^{l+t+1} - 1],\ \text{and}$$

$$\rho_{i+1}^l = \max\{\rho : e_i^l(\rho) \ge e_{i-1}^l(\rho)\} \text{ or undefined if no such } \rho \in [0,1] \text{ exists.}$$

Notice that $e_i^l(\rho)$ is determined by ρ_j^l with $j \le i$, whereas ρ_{i+1}^l is determined by $e_i^l(\rho)$ and $e_{i-1}^l(\rho)$, so the inductive step is well-defined.

Intuitively, if we consider a fixed l and a fixed node v at level l, then $e_i^l(\rho)$ represents the expected reward that v gets from its descendants, when the query is propagated exactly i levels down $v's$ subtree and when v is given a ρ-split by its parent. On the other hand, ρ_{i+1}^l represents the cheapest split (higher ρ is better for the parent and worse for the child) that $v's$ parent can offer to v in order to incentivize it to propagate the query i levels down its subtree, instead of $i-1$ levels.

Now recall that the root node seeks to retrieve the answer with high constant probability, which is equivalent to propagating the query to level h. Therefore the node v at level l has to be most willing to propagate the query $h-l$ levels further, rather than any other smaller number. This motivates the following condition:

Definition (h-consistency). *We say that h-consistency holds if for all $1 \leq l \leq h$ and $2 \leq i \leq h-l+1$, the threshold ρ_i^l is defined, and $\rho_i^l < \rho_{i-1}^l$.*

In plain words, what "h-consistency" mandates is that the incentive needed for v to be willing to propagate i levels down its subtree should be increasing in i, a phenomenon that also appeared in the fixed-payment model we studied. Also similar is the following lemma:

Lemma. *Assume h-consistency. Then, for every $1 \leq l \leq h, 0 \leq i \leq h-l-1$ and $\rho_{i+2}^l < \rho < \rho_{i+1}^l$, we have that*

$$e_i^l(\rho) > e_{i+1}^l(\rho) > \cdots > e_{h-l}^l(\rho), \text{ and}$$
$$e_i^l(\rho) > e_{i-1}^l(\rho) > \cdots > e_0^l(\rho).$$

Proof: The definition of ρ_{i+2}^l implies that for all $\rho > \rho_{i+2}^l$, $e_{i+1}^l(\rho) < e_i^l(\rho)$. But h-consistency says that $\rho(> \rho_{i+2}^l) > \rho_{i+3}^l$, hence the same argument implies $e_{i+2}^l(\rho) < e_{i+1}^l(\rho)$, so on and so forth. This proves the first half. For the other half, we note that $e_i^l(\rho)$ is linear in ρ, so the definition of ρ_{i+1}^l implies that for all $\rho \leq \rho_{i+1}^l$, $e_i^l(\rho) \geq e_{i-1}^l(\rho)$. Again h-consistency says $\rho \leq \rho_i^l$, and so $e_{i-1}^l(\rho) \geq e_{i-2}^l(\rho)$. So on and so forth, proving the second half.

We now define the Nash Equilibrium strategy:

Strategy g. *Assume h-consistency. For each $1 \leq l \leq h$, define $g^l(\rho) = \rho_i^{l+1}$ for the unique i $(0 \leq i \leq h-l)$ such that $\rho_{i+2}^l < \rho \leq \rho_{i+1}^l$, where we assume $\rho_0^{l+1} = 1$ and $\rho_{h-l+2}^l = 0$. This induces leveled strategies $g_v(\rho) = g^l(\rho)$ for every node v at level l.*

Theorem. *Assume h-consistency. Then the set of strategies $\{g_v\}$ is a Nash Equilibrium.*

Proof: Consider a node v at level l that is offered a ρ-split by its parent, where $\rho_{i+2}^l < \rho \leq \rho_{i+1}^l$. Its strategy $g_v(\rho)$ will be taken as the input $\hat{\rho}$ in the strategy $g_w(\hat{\rho})$ of its child w. But in $v's$ consideration, $w's$ strategy is as described above (N.E. imposes no deviation only when taking others' strategies as fixed). In particular, v thinks w only cares about which interval $(\rho_{k+2}^{l+1}, \rho_{k+1}^{l+1}]$ contains $\hat{\rho}$, rather than the exact value of $\hat{\rho}$. And because $w's$ strategy choice given $\hat{\rho}$ uniquely determines the strategy choice of the entire subtree under v, thus $v's$ expected payoff, it is in $v's$ best interest to offer the cheapest in the interval, or $\hat{\rho} = \rho_{k+1}^{l+1}$ for some k.

If $\hat{\rho} = \rho_{k+1}^{l+1}$, then w offers to its child ρ_k^{l+2} according to "strategy g", who in turn offers ρ_{k-1}^{l+3}, so on and so forth, until the penultimate node offers ρ_1^{k+l+1}. Hence the query will be propagated exactly $k+1$ levels down the subtree under v. The expected payoff of v from its descendants is therefore $e_{k+1}^l(\rho)$. But the previous lemma in particular shows that $e_{k+1}^l(\rho)$ is maximized at $k+1 = i$, hence it is optimal for v to set $\hat{\rho} = g_v(\rho) = e_i^{l+1}$, which is exactly what strategy g prescribes. This completes the proof.

Remark. *The previous proof not only shows that strategy g is a Nash Equilibrium, but shows something stronger. That is, strategy g is a Subgame-Perfect Nash Equilibrium (SPNE). This means that the node v is acting optimally even given an offer that is not encountered along the equilibrium path (i.e. $\rho \neq \rho_k^l$ for any k). This is a much stronger equilibrium notion because it assumes that players do not employ unreasonable "threats" in an unrealized subgame. The authors showed in the paper that strategy g is essentially the only SPNE of this game, whose proof we omit.*

Guaranteeing h-consistency

In this section we give and prove a lower bound on the initial investment that guarantees h-consistency. First recall that γ_i^l is the probability that some node i levels down from v (a fixed node at level l) will be the answer-holder with smallest depth. We define $\Gamma_i^l = \frac{1}{\gamma_i^l}\sum_{j=1}^{i-1}\gamma_j^l$.

Theorem. *Suppose $r^* \geq 4h\max\{1, \Gamma_i^l\}$ where the maximum is taken over all indices $1 \leq l \leq h, 1 \leq i \leq h-l$. Then for all $1 \leq l \leq h, 1 \leq i \leq h-l+1$, ρ_i^l is defined and satisfies*

$$1 - \frac{1}{r^* - i} < \rho_i^l \leq 1 - \frac{1}{r^* - i + 1}.$$

In particular, h-consistency holds.

Proof: We prove by induction: when $l = 1$, ρ_1^l is defined to be $1 - \frac{1}{r^*}$ so we are good. Suppose the statement is true for $1, \ldots, i$, then we recall that ρ_{i+1}^{l-1} is defined as

$$\rho_{i+1}^{l-1} = \max\{\rho : e_i^{l-1}(\rho) \geq e_{i-1}^{l-1}(\rho)\}, \text{ where}$$

$$e_i^{l-1}(\rho) = \sum_{j=1}^{i}\gamma_j^{l-1}[r^*(1-\rho)\prod_{t=0}^{j-1}\rho_{i-t}^{l-1+t+1} - 1], \text{ and}$$

$$e_{i-1}^{l-1}(\rho) = \sum_{j=1}^{i-1}\gamma_j^{l-1}[r^*(1-\rho)\prod_{t=0}^{j-1}\rho_{i-1-t}^{l-1+t+1} - 1].$$

It follows that $\rho_{i+1}^{l-1} = 1 - \frac{1}{r^*\Delta}$, where

$$\Delta = \prod_{t=0}^{i-1}\rho_{i-t}^{l+t} - \sum_{j=1}^{i-1}\frac{\gamma_j^{l-1}}{\gamma_i^{l-1}}\left(\prod_{t=0}^{j-1}\rho_{i-t-1}^{l+t} - \prod_{t=0}^{j-1}\rho_{i-t}^{l+t}\right). \tag{1}$$

By the induction hypothesis, the above parenthesized term is non-negative. It follows that

$$\Delta \leq \prod_{t=0}^{i-1}\rho_{i-t}^{l+t} \leq \prod_{t=0}^{i-1}(1 - \frac{1}{r^* - i + t + 1}) = \frac{r^* - i}{r^*}$$

which shows that $\rho_{i+1}^{l-1} = 1 - \frac{1}{r^*\Delta} \leq 1 - \frac{1}{r^* - i}$, giving the desired upper bound for $i+1$ (we will show $\Delta > 0$ below).

On the other hand, the parenthesized term in (1) is at most

$$\begin{aligned}
&\prod_{t=0}^{j-1}(1-\frac{1}{r^*-i+t+2})-\prod_{t=0}^{i-1}(1-\frac{1}{r^*-i+t})\\
=&\frac{r^*-i+1}{r^*-i+j-1}-\frac{r^*-i-1}{r^*-i+j-1}\\
=&\frac{2j}{(r^*-i+j+1)(r^*-i-j-1)}\\
<&\frac{2i}{(r^*)^2}
\end{aligned}$$

Plugging this into (1) gives $\Delta > \frac{r^*-i-1}{r^*-1} - \frac{2i}{(r^*)^2}\Gamma_i^{l-1} \geq \frac{r^*-i-1}{r^*-1} - \frac{1}{2r^*}$ by using the assumption on r^* (also note $i \leq h$). Hence $r^*\Delta > r^* - i - 1 + \frac{r^*-i-1}{r^*-1} - \frac{1}{2} > r^* - i - 1$ as $r^* \geq 4h \geq 4i$. This gives the desired lower bound and completes the inductions step.

Remark. *The preceding result implies that the equilibrium outcome of strategy g is such that the nodes at level l receive a split ρ_{h-l}^l, which is close to $1 - \frac{1}{r^*-(h-l)}$. In particular, those at level h get roughly 1, which is just enough to cover their relaying cost.*

Upper Bound for Γ_i^l

In this final section we heuristically show that for any $b > 1$, Γ_i^l is bounded by an absolute constant (depending on b, ϵ but not n), a rather technical result in the paper. Note that the number of nodes i levels down from v is roughly b^i, whereas the number of nodes with level less than $l+i$ is roughly $\frac{b^{l+i}}{b-1}$. Hence the probability γ_i^l is roughly $(1-\frac{1}{n})^{\frac{b^{l+i}}{b-1}}\left(1-(1-\frac{1}{n})^{b^i}\right)$. For $j < i$, the ratio $\frac{\gamma_j^l}{\gamma_i^l}$ is roughly

$$\frac{(1-\frac{1}{n})^{\frac{b^{l+j}}{b-1}}}{(1-\frac{1}{n})^{\frac{b^{l+i}}{b-1}}} * \frac{1-(1-\frac{1}{n})^{b^j}}{1-(1-\frac{1}{n})^{b^i}} < \frac{(1-\frac{1}{n})^{\frac{b^{l+j}}{b-1}}}{(1-\frac{1}{n})^{\frac{b^{l+i}}{b-1}}} * \frac{1-(1-\frac{1}{n})^{b^{l+j}}}{1-(1-\frac{1}{n})^{b^{l+i}}}$$

because $\frac{1-x}{1-x^t}$ is decreasing in x whenever $t = b^{i-j} > 1$.

The RHS evaluates simply to $\frac{(1-\frac{1}{n})^{\frac{b^{l+j}}{b-1}}-(1-\frac{1}{n})^{\frac{b^{l+j+1}}{b-1}}}{(1-\frac{1}{n})^{\frac{b^{l+i}}{b-1}}-(1-\frac{1}{n})^{\frac{b^{l+i+1}}{b-1}}}$. Summing over j then gives, roughly:

$$\Gamma_i^l \leq \frac{(1-\frac{1}{n})^{\frac{b^l}{b-1}}-(1-\frac{1}{n})^{\frac{b^{l+i}}{b-1}}}{(1-\frac{1}{n})^{\frac{b^{l+i}}{b-1}}-(1-\frac{1}{n})^{\frac{b^{l+i+1}}{b-1}}}$$

The numerator is obviously less than 1, so we only need to give a lower bound to the denominator. Since the function $x - x^b$ is decreasing in x, a lower bound for $(1-\frac{1}{n})^{\frac{b^{l+i}}{b-1}}$ suffices. But this is roughly the probability that no node in the first $l+i-1 < h$ levels holds the answer, which by the choice of h is greater than ϵ, so we are done (a rough upper bound for Γ_i^l would be $\frac{1}{\epsilon - \epsilon^b}$).

Recall the previous theorem with the condition $r^* \geq 4h \max\{1, \Gamma_i^l\}$. We have thus shown that an initial investment of $\Theta(h) = \Theta(\log(n))$ is sufficient, for any fixed $b > 1$. This contrast with Kleinberg and Raghavan is the main contribution of the paper.