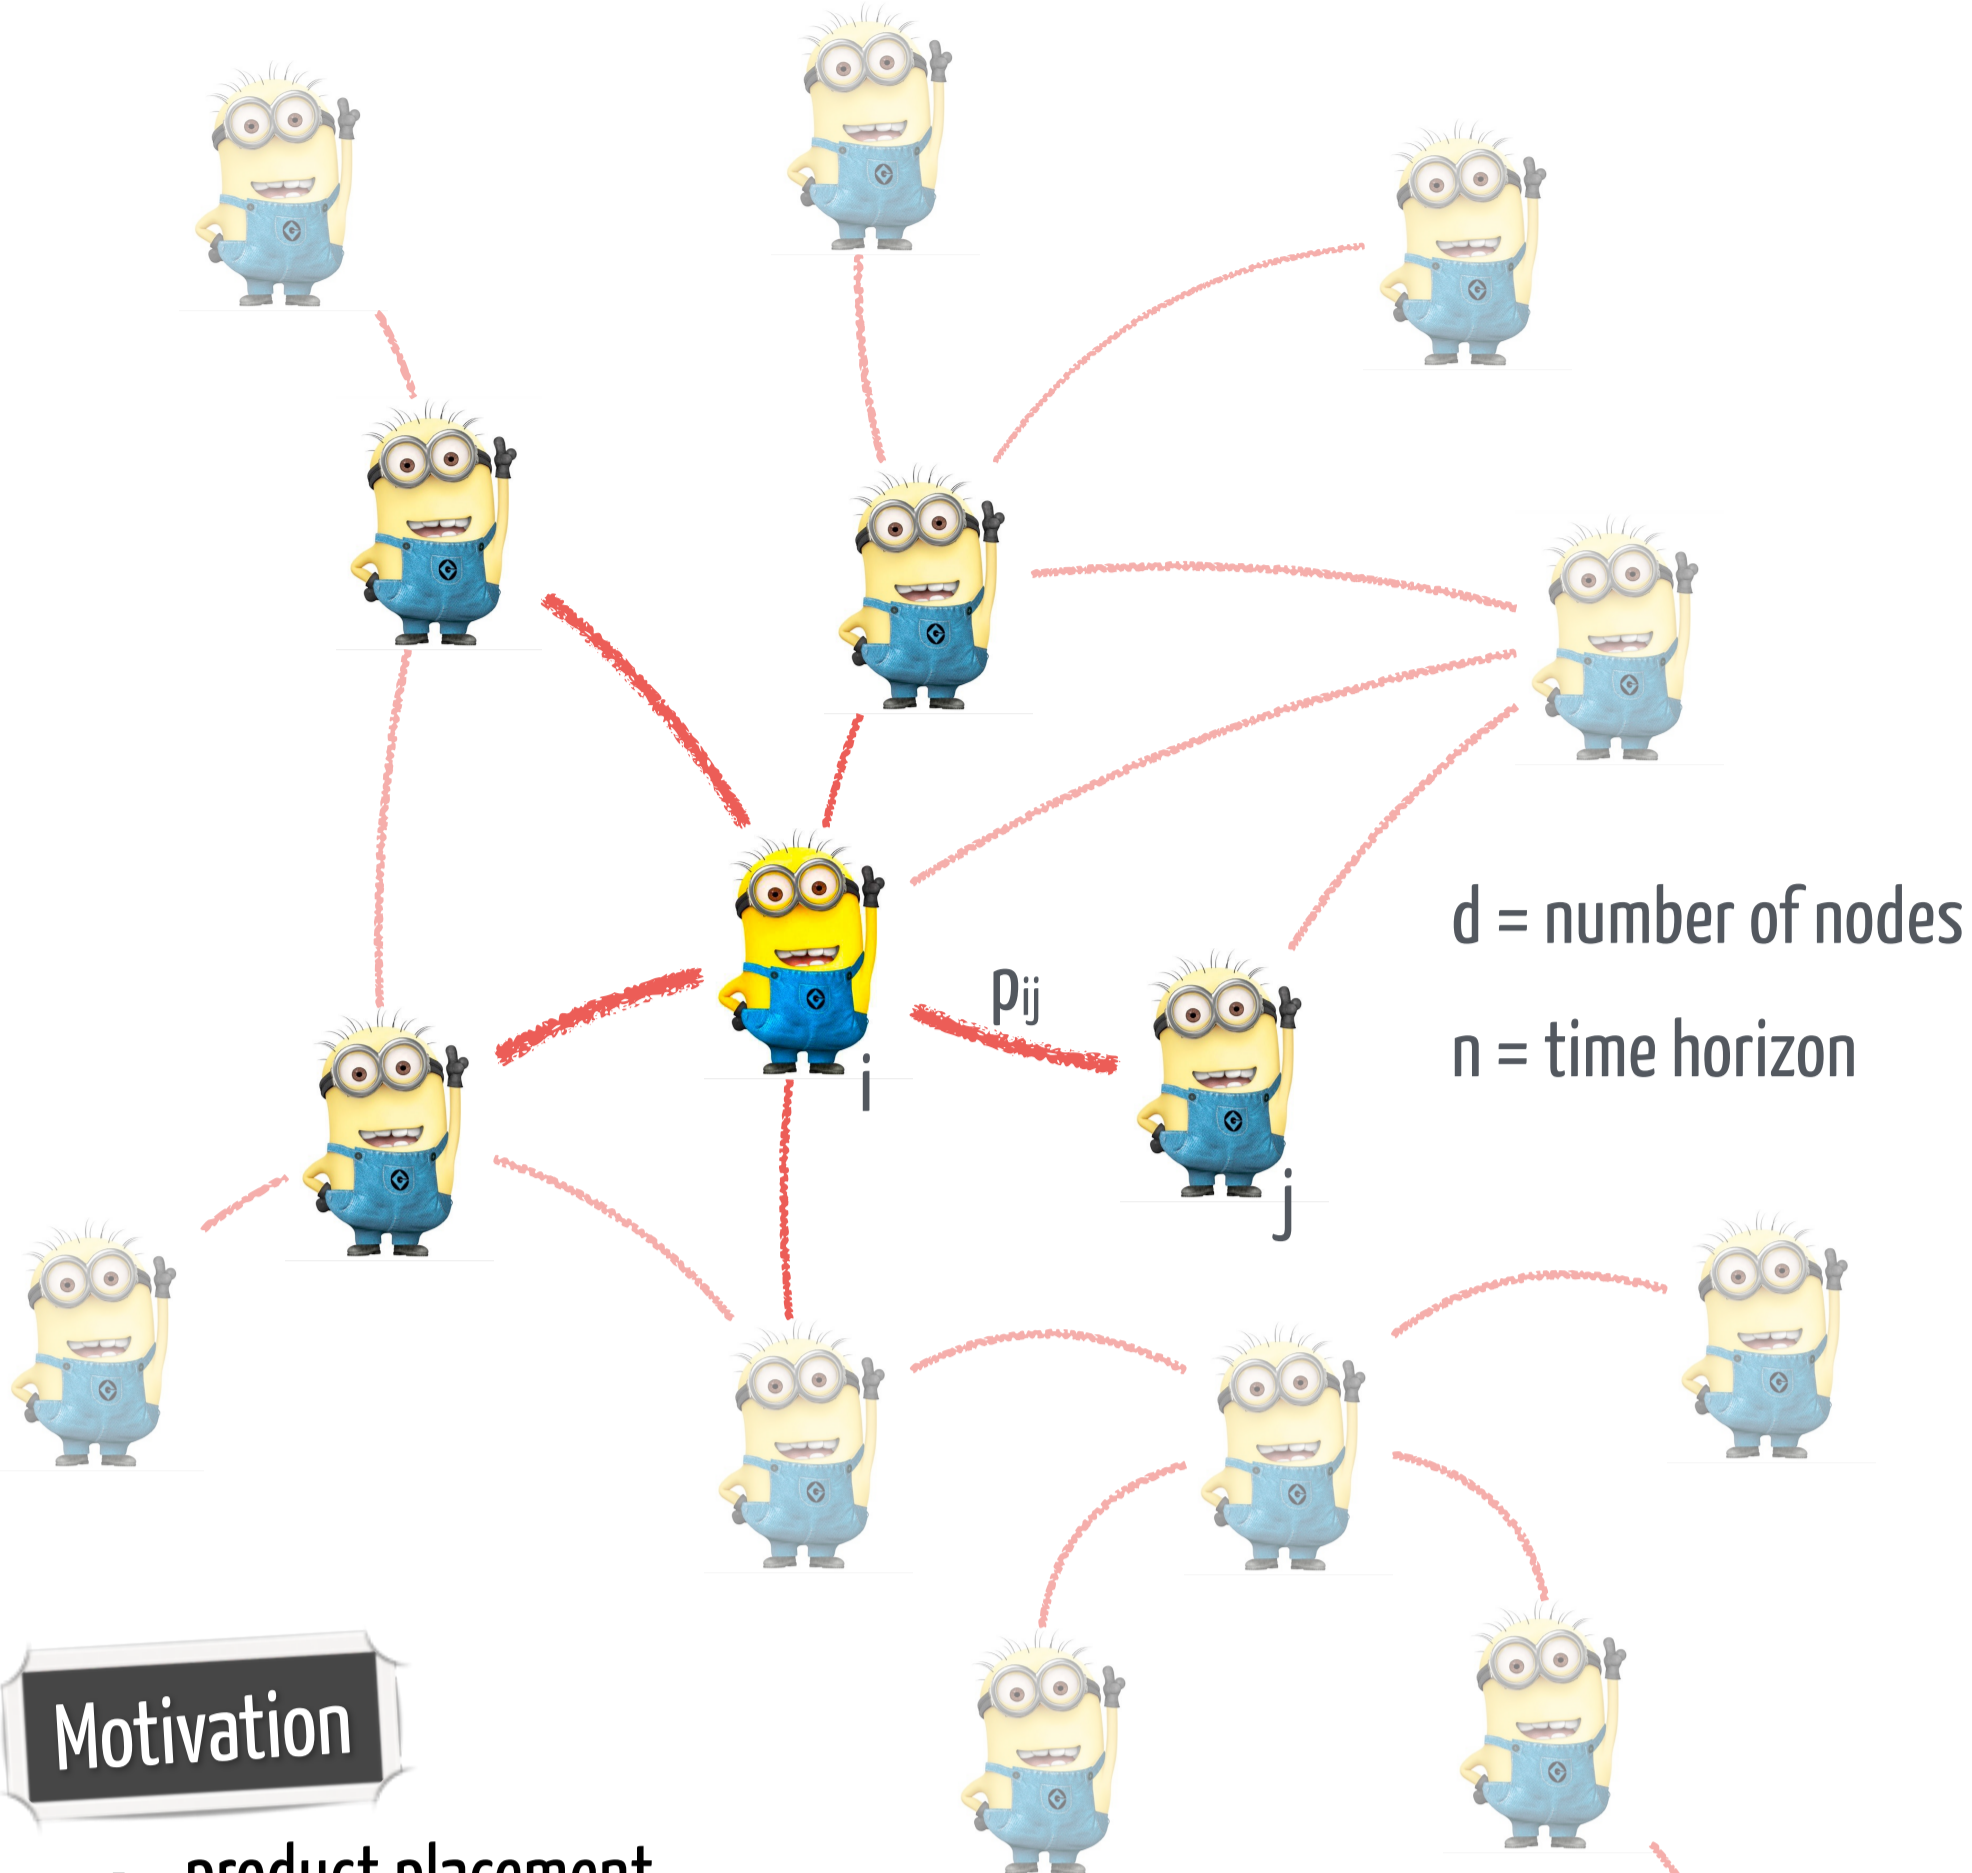



Revealing graph bandits for maximizing local influence

ALEXANDRA CARPENTIER and MICHAL VALKO
carpentier@math.uni-potsdam.de and michal.valko@inria.fr



Motivation

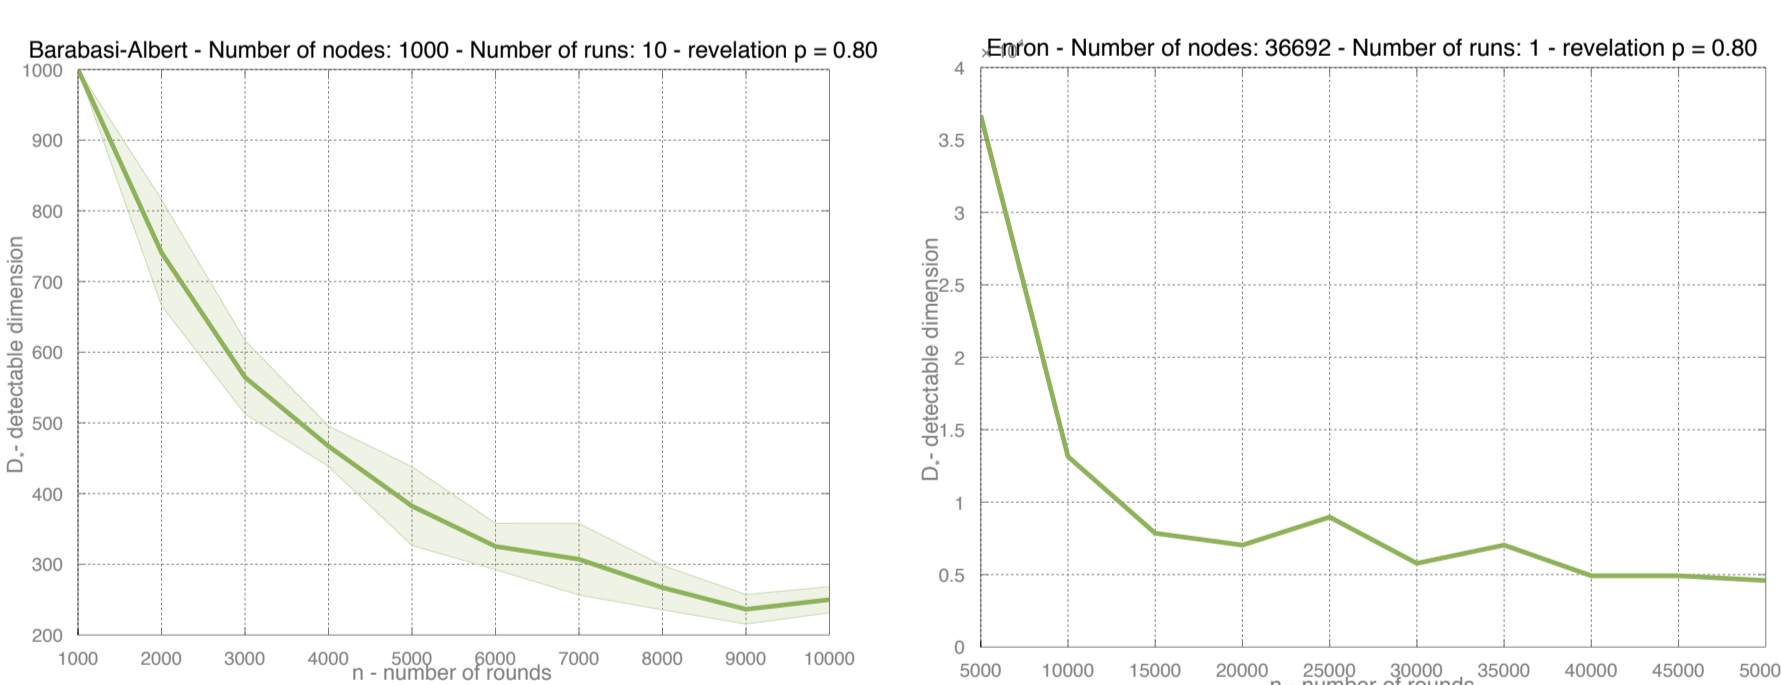
- product placement
- information campaign
- local model of influence and unknown graphs
- information gathered through "likes" or promotional codes

Detectable dimension

- number of nodes we can efficiently extract in less than n rounds
- function D controls number of nodes given a gap
- $D(\Delta) \stackrel{\text{def}}{=} |\{i \leq d : r_i^\circ - r_i^\circ \leq \Delta\}|$
- $D(r) = d$ for $r \geq r_*$ and $D(0) =$ number of most influenced nodes
- Detectable dimension $D_* = D(\Delta_*)$
- Detectable gap Δ_* constants coming from the analysis and the Bernstein inequality

$$\Delta_* \stackrel{\text{def}}{=} 16 \sqrt{\frac{r_*^\circ d \log(nd)}{T_*}} + \frac{80d \log(nd)}{T_*}$$

- Detectable horizon T_* , smallest integer s.t. $T_* r_*^\circ \geq \sqrt{D_* n r_*^\circ}$
- Equivalently: D_* corresponding to smallest T_* such that
- $T_* r_*^\circ \geq \sqrt{D \left(16 \sqrt{\frac{r_*^\circ d \log(nd)}{T_*}} + \frac{80d \log(nd)}{T_*} \right) n r_*^\circ}$
- For (easy, structured) star graphs $D_* = 1$ even for small n (big gain)
- For (difficult) empty graphs $D_* = d$ even for large n (no gain)
- In general: D_* roughly decreases with n and it is small when D decreases quickly
- For n large enough D_* is the number of the most influences nodes
- Example: D_* for Barabási-Albert model & Ernon graph as a function of n



Setting

Unknown $(p_{ij})_{ij}$ — (symmetric) probability of influences
In each time step $t = 1, \dots, n$
learner picks a node k_t
environment reveals the set of influenced node S_{k_t}
Select influential people = Find the strategy maximising
$$L_n = \sum_{t=1}^n |S_{k_t, t}|$$

Our solution

BAndit REvelator: 2-phase algorithm

- global exploration phase
 - super-efficient exploration 🐱
 - linear regret 🐱 — needs to be short!
 - extracts D_* nodes
- bandit phase
 - uses a minimax-optimal bandit algorithm
 - GraphMOSS is a little brother of MOSS
 - has a "square root" regret on D_* nodes
- D_* realizes the optimal trade-off!
- different from exploration/exploitation tradeoff

Guarantees

- Upper bound on the regret of BARE
- $\mathbb{E}[R_n] \leq C \min(r_* n, D_* r_* + \sqrt{r_* n D_*})$
- Matching lower bound

Algorithm

GraphMOSS

Input
 d : the number of nodes
 n : time horizon
Initialization
Sample each arm twice
Update $\hat{r}_{k,2d}, \hat{\sigma}_{k,2d}$, and $T_{k,2d} \leftarrow 2$, for $\forall k \leq d$
for $t = 2d + 1, \dots, n$ do
 $C_{k,t} \leftarrow 2\hat{\sigma}_{k,t} \sqrt{\frac{\max(\log(n/(dT_{k,t})), 0)}{T_{k,t}} + \frac{2 \max(\log(n/(dT_{k,t})), 0)}{T_{k,t}}}$, for $\forall k \leq d$
 $k_t \leftarrow \arg \max_k \hat{r}_{k,t} + C_{k,t}$
Sample node k_t and receive $|S_{k_t, t}|$
Update $\hat{r}_{k_t, t+1}, \hat{\sigma}_{k_t, t+1}$, and $T_{k_t, t+1}$, for $\forall k \leq d$
end for

Definitions

The number of expected influences of node k is by definition
$$r_k = \mathbb{E}[|S_{k,t}|] = \sum_{j \leq d} p_{k,j}$$

Oracle strategy always selects the best
$$k^* = \arg \max_k \mathbb{E} \left[\sum_{t=1}^n |S_{k,t}| \right] = \arg \max_k n r_k$$

Expected regret of the oracle strategy
$$\mathbb{E}[L_n^*] = n r_*$$

Expected regret of any adaptive strategy unaware of $(p_{ij})_{ij}$
$$\mathbb{E}[R_n] = \mathbb{E}[L_n^*] - \mathbb{E}[L_n]$$

Baseline

- We only receive $|S|$ instead of S
- Can be mapped to multi-arm bandits
 - rewards are $0, \dots, d$
 - variance bounded with r_{kt}
- We adapt MOSS to GraphMOSS
- Regret upper bound of GraphMOSS
$$\mathbb{E}[R_n] \leq U \min(r_* n, r_* d + \sqrt{r_* n d})$$
- matching lower bound

Algorithm

BARE - BAndit REvelator

Input
 d : the number of nodes
 n : time horizon
Initialization
 $T_{k,t} \leftarrow 0$, for $\forall k \leq d$
 $\hat{r}_{k,t}^\circ \leftarrow 0$, for $\forall k \leq d$
 $t \leftarrow 1, \hat{T}_* \leftarrow 0, \hat{D}_{*,t} \leftarrow d, \hat{\sigma}_{*,t} \leftarrow d$
Global exploration phase
while $t (\hat{\sigma}_{*,t} - 4\sqrt{d \log(dn)/t}) \leq \sqrt{\hat{D}_{*,t} n}$ do
Influence a node at random (choose k_t uniformly at random) and get $S_{k_t, t}$ from this node
 $\hat{r}_{k_t, t+1}^\circ \leftarrow \frac{t}{t+1} \hat{r}_{k_t, t}^\circ + \frac{d}{t+1} |S_{k_t, t}|$
 $\hat{\sigma}_{*, t+1} \leftarrow \max_{k'} \sqrt{\frac{\hat{r}_{k', t+1}^\circ}{t+1} + 8d \log(dn)/(t+1)}$
 $w_{*, t+1} \leftarrow 8\hat{\sigma}_{*, t+1} \sqrt{\frac{d \log(dn)}{t+1} + \frac{24d \log(dn)}{t+1}}$
 $\hat{D}_{*, t+1} \leftarrow \left\{ k : \max_{k'} \hat{r}_{k', t+1}^\circ - \hat{r}_{k_t, t+1}^\circ \leq w_{*, t+1} \right\}$
 $t \leftarrow t + 1$
end while
 $\hat{T}_* \leftarrow t$
Bandit phase
Run minimax-optimal bandit algorithm on the \hat{D}_{*, \hat{T}_*} chosen nodes (e.g., Algorithm 1)

Experiments

