Event Extent Estimation

Marcin Bienkowski Leszek Gasieniec Marek Klonowski
Miroslaw Korzeniowski Stefan Schmid

SIROCCO 2010

June 8th, 2010

Problem description

- Given a set of nodes V connected into a graph G,
- perform some precomputations.
- An adversary chooses a set of *active nodes* $V' \subseteq V$ *which wake up*
- *Goal: at least one active node has to learn* V'
- *Constraints:*
  - *works in rounds*
  - *a message from u to v costs* 1 *unit of energy*
  - *goal: minimize time and the total used energy (not maximum)*
  - *synchronized and time of wake-up known*
  - *no interferences*

Problem on a picture

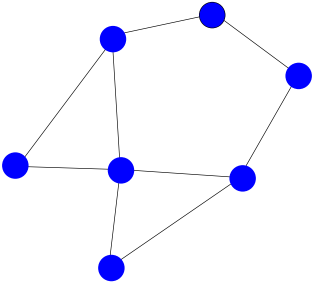

Problem on a picture

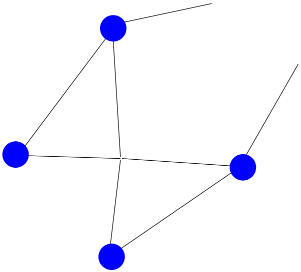

Distributed Disaster Disclosure

- B. Mans, S. Schmid, and R. Wattenhofer, "Disributed Disaster Disclosure", SWAT 2008
  - difference in model: non-active nodes (from $V \setminus V'$) can help
  - solution: clustering, leaders of clusters made responsible for their clusters
  - not possible in our setting — any node may be *inactive*

Graphs

- **Complete graphs**
- General graphs
- Planar graphs

Deterministic upper bounds

- Algorithm GROUP, example for base $k = 2$
- assume $\log_2 n$ is an integer
- partition V recursively (into halves on each level)
- two clusters, an active leader in each one, each leader pings the whole other cluster
- sequentially: the 1st leader has an opportunity to ping the other cluster, then the 2nd leader
- in parallel: they do it at once

Randomized upper bounds

- Cardinality guessing, algorithm RAND
- if $s = |V'|$ known, ping all nodes with probability $\frac{1}{s}$
- $\Pr[\mathrm{failure}] = (1 - \frac{1}{s})^s \approx \frac{1}{e}$
- approximate $1/s$: start from $1/n$ and double in each round until success or 1
- expected time: $O(\log n/s)$
- expected communication cost: $O(n)$ messages

Overview

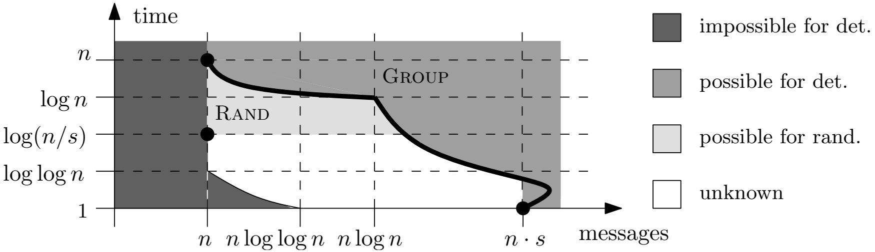
time
n
$\log n$
$\log(n/s)$
$\log \log n$
1
GROUP
RAND
n
$n \log \log n$
$n \log n$
$n \cdot s$
messages
impossible for det.
possible for det.
possible for rand.
unknown

Graphs

- Complete graphs
- **General graphs**
- Planar graphs

Arboricity

- treeishness
- the minimal number of trees needed to cover G
- or the minimal number of forests constituting a partitioning of G
- useful for active neighborhood discovery in our model:
  - 1st round: ping your parents in all trees
  - 2nd round: answer all your active children
  - total cost: $O(\alpha \cdot s)$, where *alpha* is the arboricity
- arboricity is constant for planar graphs

Searches

- DFS (Depth First Search)
- BFS (Breadth First Search)
- parallelized versions

Algorithm MinID

- neighborhood discovery
- assume s known
- starting from 1-node clusters:
  - spread your number in your cluster
  - learn neighboring numbers in neighboring clusters
  - choose minimum and propose to join it
  - if no other cluster smaller, annex recursively all smaller clusters
  - in two rounds a cluster annexes or gets annexed
- for unknown s approximate from 1 and double
- result: $O(s \log s)$ rounds and $O(\alpha \cdot s \log s)$ messages

Overview

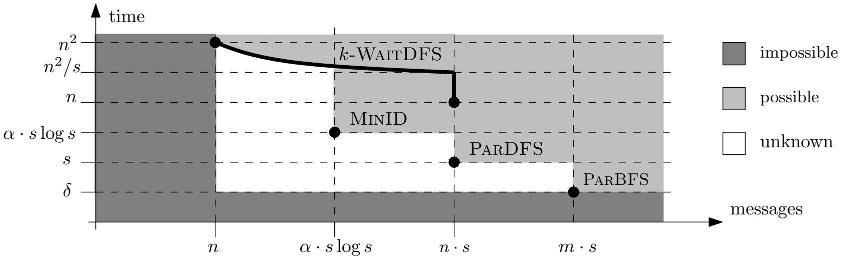
time
n^2
n^2/s
n
$\alpha \cdot s \log s$
s
δ
k-WaitDFS
MinID
ParDFS
ParBFS
messages
n
$\alpha \cdot s \log s$
$n \cdot s$
$m \cdot s$
impossible
possible
unknown

Graphs

- Complete graphs
- General graphs
- **Planar graphs**

Separation Theorem

Theorem (Lipton and Tarjan separation theorem)

For a planar graph $G = (V, E)$, where $|V| = n$, it is possible to find A, B, U such that

- $A \uplus B \uplus U = V$
- *A and B are not connected to each other*
- $|A|, |B| \leq \frac{2}{3} \cdot n$
- $|U| \in O(\sqrt{n})$

Algorithm k-SEP

- precomputing phase: hierarchical decomposition
- start from neighborhood discovery (constant arboricity)
- level sizes form a geometric sequence
- on a level of size n, $O(\sqrt{n})$ nodes start a search
- different levels of parallelism possible

Overview

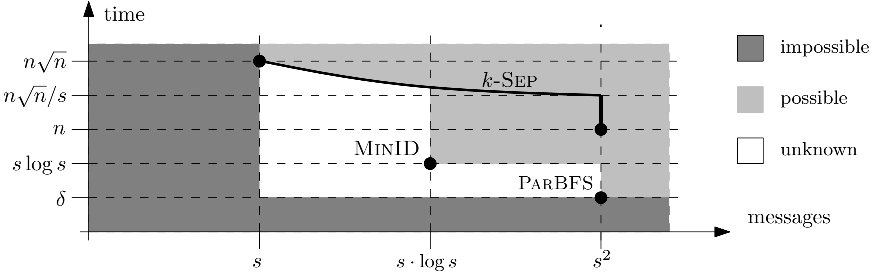
time
$n\sqrt{n}$
$n\sqrt{n}/s$
n
$s \log s$
δ
k-Sep
MinID
ParBFS
s
$s \cdot \log s$
s^2
messages
impossible
possible
unknown

Thank you

Thank You for Your Attention!