Local verification of global proofs

and

Redundancy in distributed proofs

Laurent Feuilloley, Pierre Fraigniaud,
Juho Hirvonen, Ami Paz and Mor Perry

DISC · New Orleans · 16th October 2018

Local decision and certification

Local decision

- Setting : distributed synchronous network computing.
- Goal : check whether the network satisfies some property.
- Constraint : every node knows only its view at distance 1.
- Identifiers on $O(\log n)$ bits.

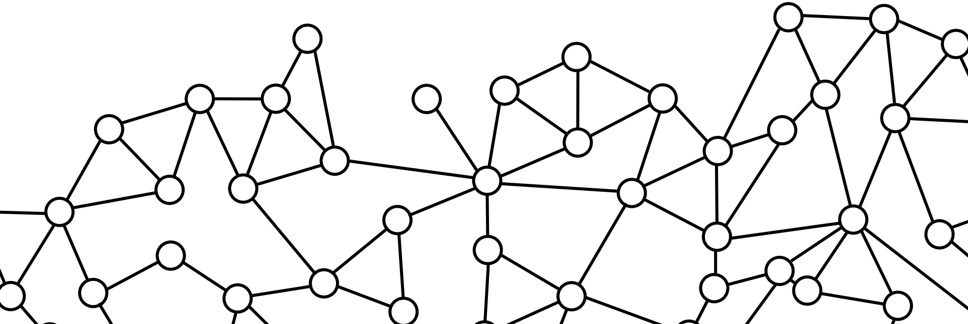

Local decision

- Setting : distributed synchronous network computing.
- Goal : check whether the network satisfies some property.
- Constraint : every node knows only its view at distance 1.
- Identifiers on $O(\log n)$ bits.

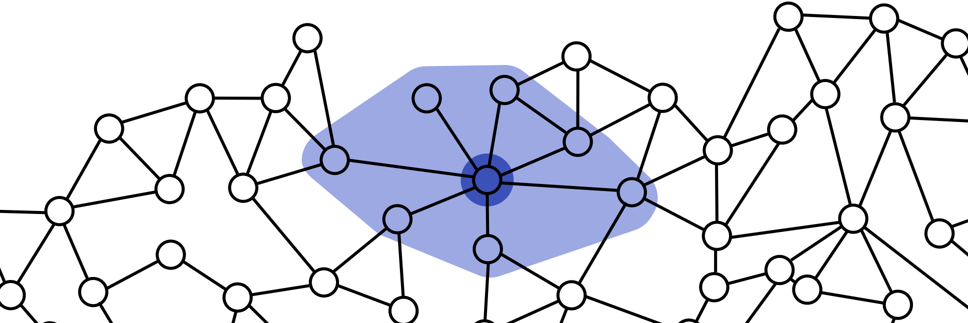

Decision rule

[Awerbuch, Patt-Shamir, Varghese 91], [Naor,Stockmeyer 93], [Itkis, Levin 94], [Afek, Kutten,Yung 97].

Decision rule :

- Every node makes one (local) decision : *accept* or *reject*.
- The configuration is accepted if and only if all the local decisions are *accept*.

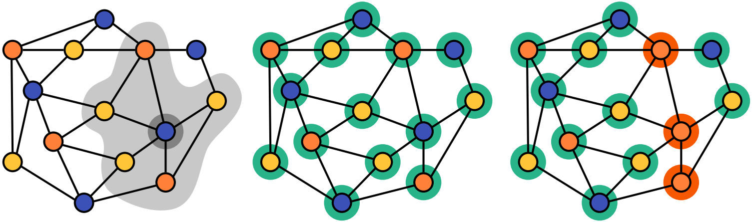

Limits of local decision

Property to check :
The marked edges form a spanning tree of the network.

Theorem [Folklore] :
There is no local decision algorithm to decide this property.

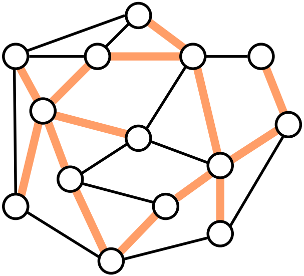

Extra information

Idea from fault-tolerance : store extra information at the nodes.

Example : for spanning trees, store root ID, and distance to root.

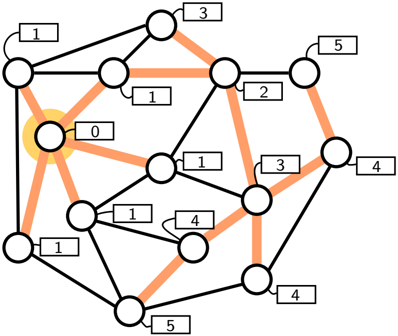

Local certification

Definition [Korman-Kutten-Peleg 05] :
A certificate (or proof) assignment is a function $V \rightarrow \{0,1\}^*$.

Story : Certificates are given by a prover, and the nodes verify.

Correctness rule :

- Good configuration $\rightarrow \exists$ certificates, $\forall v$, v accepts.
- Bad configuration, $\rightarrow \forall$ certificates, $\exists v$, v rejects.

Spanning tree scheme

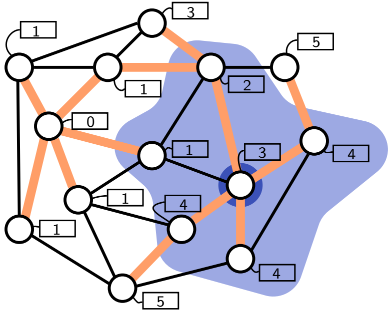

Verifier on node v :

- Check : neighbours have same root-ID.
- If $d = 0$:
  check the root-ID
  $\forall$ neighbour u, $d(u) = 1$.
- If $d > 0$:
  $\exists$ neighbour u, $d(u) = d - 1$
  $\forall$ neighbour $w \neq u$,
  $d(w) = d + 1$

Theorem [Itkis-Levin 94] : The spanning tree scheme is a correct.

Spanning tree scheme

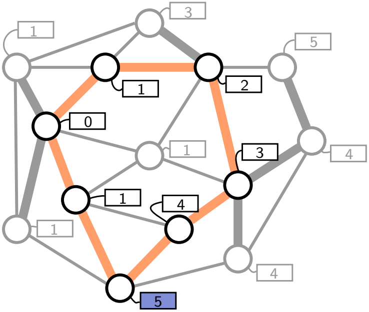

Verifier on node v :

- Check : neighbours have same root-ID.
- If $d = 0$:
  check the root-ID
  $\forall$ neighbour u, $d(u) = 1$.
- If $d > 0$:
  $\exists$ neighbour u, $d(u) = d - 1$
  $\forall$ neighbour $w \neq u$,
  $d(w) = d + 1$

Theorem [Itkis-Levin 94] : The spanning tree scheme is a correct.

Certificate size

0	$\Theta(\log n)$	$\Theta(\log^2 n)$	$\Theta(n^2)$

- [Naor-Stockmeyer 93] : LCL problems.
- [Korman, Kutten, Peleg 05] : formalization, $\Omega(\log n)$ for spanning tree, universal $O(n^2)$ scheme.
- [Korman, Kutten 06] $\Omega(\log^2 n)$ for minimum spanning tree.
- [Göös, Suomela 11] general model.

First paper

Local verification of global proofs

a.k.a.

uniformity and mixed schemes

Extra information

Idea from fault-tolerance : store extra information at the nodes.

Example : for spanning trees, store root ID, and distance to root.

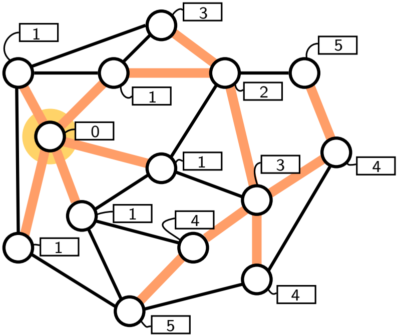

Mixed schemes

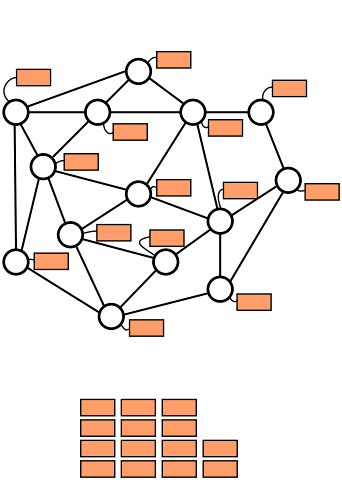

Uniformity

Definition :
The uniformity of a property, is the ratio of optimal proof sizes :

$$\frac{\sum_v |c(v)| \text{ (Classic scheme)}}{|c_{Glob}| + \sum_v |c_{Loc}(v)| \text{ (Mixed scheme)}}$$

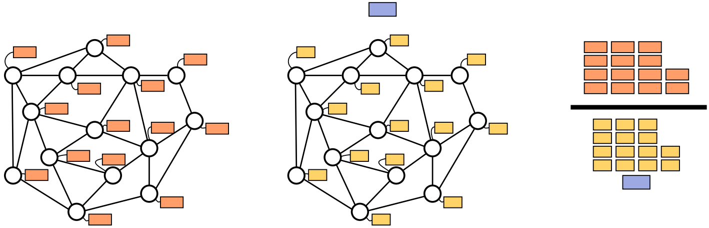

Alternative : the uniformity is a *price of locality*.

Bounds

Theorem : The uniformity is between 1 and n, for all properties.

Proof :

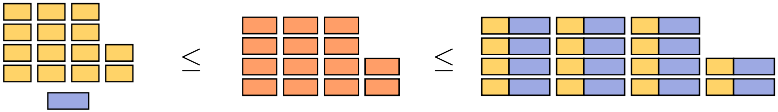

Uniformity ratios

Theorems :

- Uniform schemes have uniformity n.
- Minimum spanning tree has uniformity $\Theta(\log n)$.
- "At least one object" has uniformity $O(1)$.
  $\rightarrow$ Spanning tree, non-bipartiteness, leader election...

Proof :

- Uniform scheme already use global proofs.
- There exists a global proof of size $O(n \log n)$.
- "At least one node selected" has mixed certificates in $\Omega(n \log n)$.

Bridges thanks to global proofs

- Global proofs are more common than local proofs
  - property testing
  - communication complexity
- Crossing borders :
  - In [F., Fraigniaud, Hirvonen 16], a local hierarchy is defined.
  - We do not know if this hierarchy is infinite.
  - The classic lower bound technique somehow uses global proofs.
  - Theorem : If this technique works then it would settle the same question in a hierarchy in communication complexity (open since [Babai, Frankl, Simon 86]).

Open problem

Can purely global proofs be worse than local proofs ?

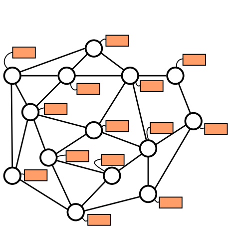

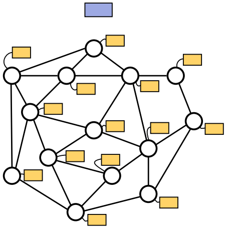

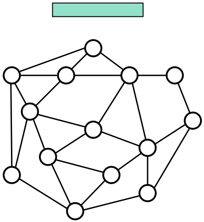

$\hookrightarrow$ A good starting point : bipartiteness.

Second paper

Redundancy in distributed proofs

Local decision

- Setting : distributed network computing.
- Goal : check whether the network satisfies some property.
- Constraint : every node knows only its view at distance 1.
- Identifiers on $O(\log n)$ bits.

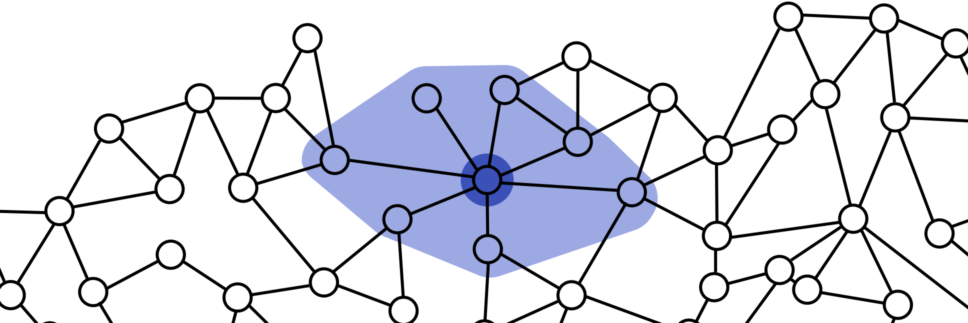

(Slightly less) Local decision

- Setting : distributed network computing.
- Goal : check whether the network satisfies some property.
- Constraint : every node knows only its view at distance **r**.
- Identifiers on $O(\log n)$ bits.

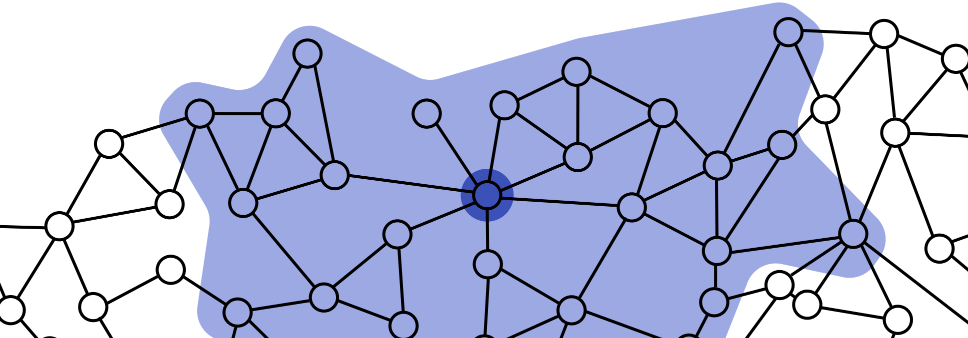

Distance-r certification

- [Korman, Kutten, Masuzawa 11]
  $(\log n, \log n)$-scheme for minimum spanning tree.
  $\rightarrow$ certificate of size $\frac{\log^2 n}{r}$, with radius $r = \log n$.

- [Ostrovsky, Perry, Rosenbaum 17]
  Linear scaling :
  $$\text{proof-size}(r) \leq \frac{\text{proof-size}(r=1)}{r}$$
  proved for the universal scheme and spanning trees.

Scaling

Definition : The *scaling* of a language is a function $f(r)$ s.t. :

$$\text{proof-size}(r) \leq \frac{\text{proof-size}(r=1)}{f(r)}$$

Two main scenarios :

- Linear scaling : $f(r)$ is $\Theta(r)$.
- Maximum scaling : $f(r)$ is $\Theta(b(r))$,
  $b(r)=$ minimum number of nodes in a ball of radius r.

Point of view : a measure of redundancy.

Technique 1 : sampling

→ Pick the good bits of information in each certificate.

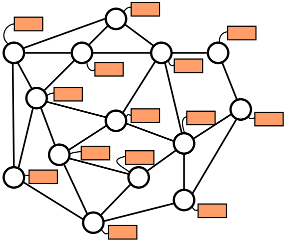

Theorem :

- Uniform schemes have maximum scaling.
- Distances scheme have linear scaling.

Technique 1 : sampling

→ Pick the good bits of information in each certificate.

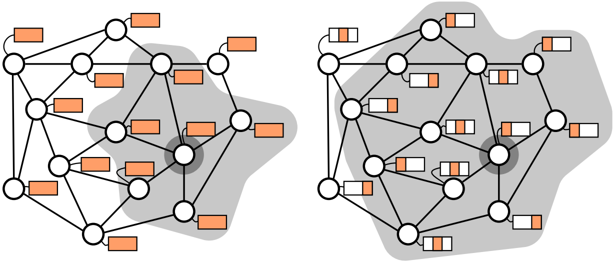

Theorem :

- Uniform schemes have maximum scaling.
- Distances scheme have linear scaling.

Technique 2 : sparsify-spread

→ Erase most of the certificates, spread the rest.

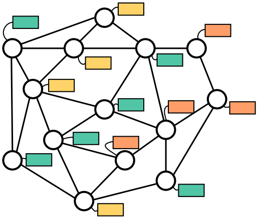

Theorem :

- Minimum spanning tree has linear scaling.
- In paths, cycles, trees, grids, *any property* has a linear scaling.

Technique 2 : sparsify-spread

→ Erase most of the certificates, spread the rest.

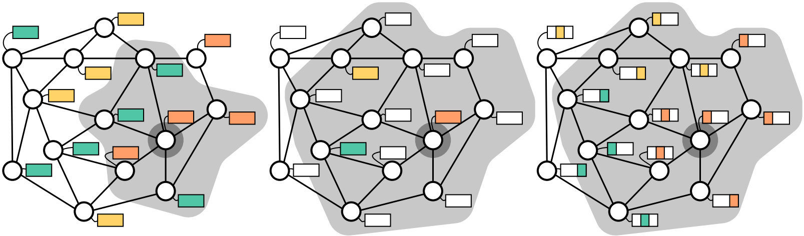

Theorem :

- Minimum spanning tree has linear scaling.
- In paths, cycles, trees, grids, *any property* has a linear scaling.

Technique 2 : sparsify-spread

Focus : any language on cycles.

Start : A distance-1 scheme.

Sparsify : Make well separated zones of unlabeled nodes, with diameter $\approx r$.

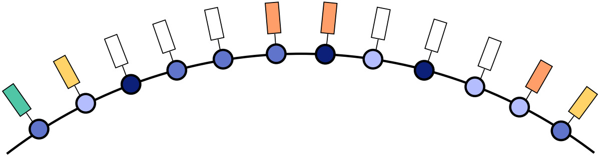

Spread : Spread the labels in layers

Lower bound

Theorem : A $\tilde{\Omega}\left(\frac{n}{r}\right)$ lower bound for diameter.

$\rightarrow$ Non-deterministic communication complexity reduction.

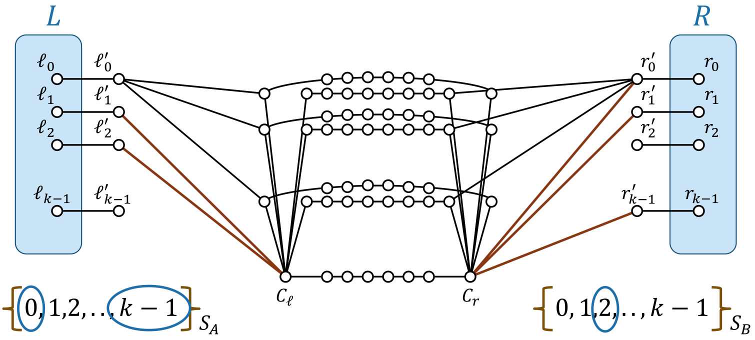

Open problems

General problem : how does certification scale ?

More concrete problem : Does every property scales linearly ?

Even more concrete : Does $k(n)$-colourability always scales linearly ?