Compact Forbidden-set Routing

Bruno Courcelle
LaBRI, CNRS
Bordeaux
bruno.courcelle@labri.fr

Andrew Twigg
Cambridge University /
Thomson Research, Paris
andrew.twigg@cl.cam.ac.uk

STACS, Feburary 2007

Background: distance labelling

Distance labelling problem: given graph G, compute labels $J(x)$ for $x \in V$ s.t. given labels $J(x), J(y)$, we can compute the distance $d(x, y)$ in G.

Main results in the area:

- ★ Exact distance labeling for general graphs: $\tilde{\Theta}(n)$ bits [Peleg, Gavoille, ...]
- ★ Stretch-3 scheme for general graphs using $\tilde{O}(n^{1/2})$ bits [Thorup]
- ★ For treewidth k graphs, exact scheme using $\Theta(k \log^2 n)$ bits.
- ★ For graphs excluding a fixed minor, $\tilde{O}(1)$ bits [Thorup, Gavoille, ...]

Most of these can be obtained as *compact routing* schemes: given $J(x), J(y)$, determine the *next-hop* on the path $x \to y$.

"Distance labelling with obstacles"

Forbidden-set distance labelling problem: given labels $\{J(x) : x \in X\}$ for $X \subseteq V$, what is the distance with no intermediate nodes in X?
i.e $d_{G\setminus(X\setminus\{x,y\})}(x, y)$ for $x, y \in X$?

Theorem 1 *For graphs of cliquewidth and treewidth $\leq k$, we can use labels of size $O(k^2 \log^2 n)$ bits. Only a factor k larger than pure distance labelling.*

- ★ WHY?? E.g. Internet routers can specify routing policies; not known if they can be satisfied using small collections of trees.

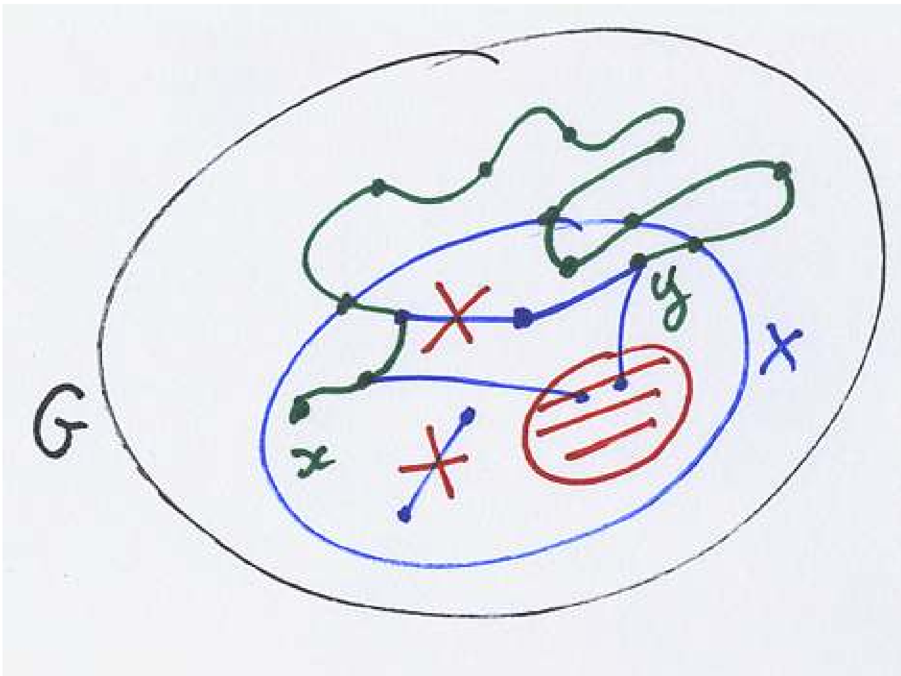

Shortcut graph $G^+[X]$

Given the labels $\{J(x) : x \in X\}$, we can compute the graph

$$G^+[X] = G[X] + \{\text{shortest detours outside } X \text{ between } x, y \in X\}.$$

Then we can answer distance queries with forbidden vertices and edges in X.

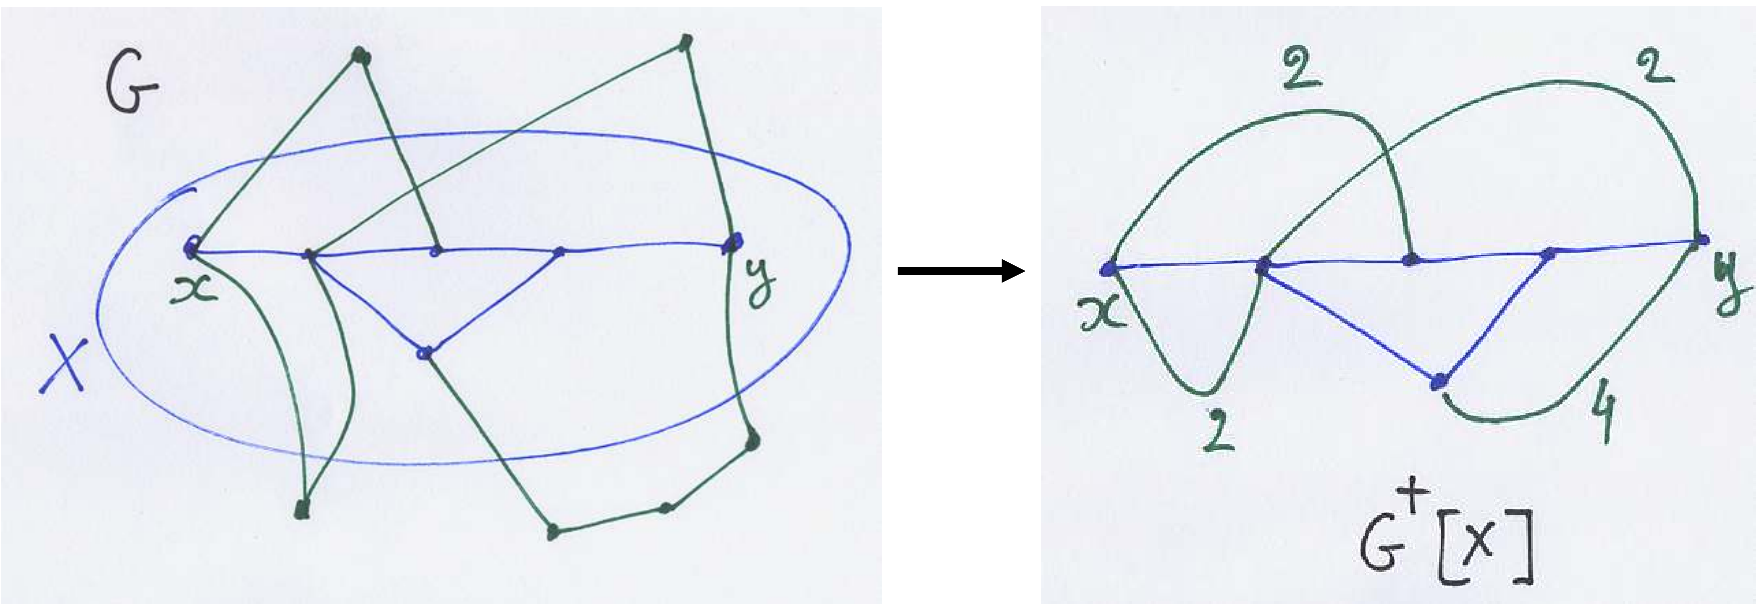

Algebraic representations: Cographs

Graphs defined by terms on two binary operations:

- ★ $G \oplus H$: disjoint union of G, H
- ★ $G \otimes H = G \oplus H$ plus all edges between G, H

For example, *adjacency labels*:
$A(x) = \otimes 1 \oplus 1 \otimes 2$
$A(y) = \otimes 2 \otimes 1 \oplus 2 \otimes 1 \otimes 1$

$G[X]$ can be constructed from $\{A(u) : u \in X\}$.

Fact 1 $|A(u)| \leq 2.ht(t)$
More powerful graph operations to be defined next.

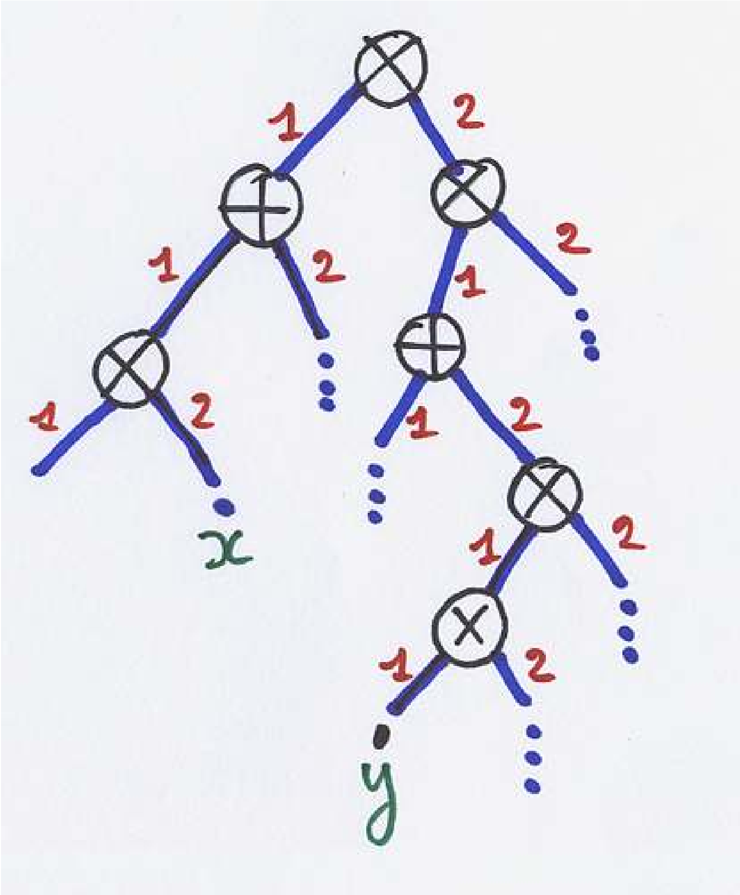

Useful class of graphs: mcwd

- A variant of clique width; each node can have several colours, drawn from a set L of k colours.
- Graph G can be represented by a term with constants (leaves) in bijection to vertices of G and internal nodes are binary operations $\otimes_{R,g,h}$:

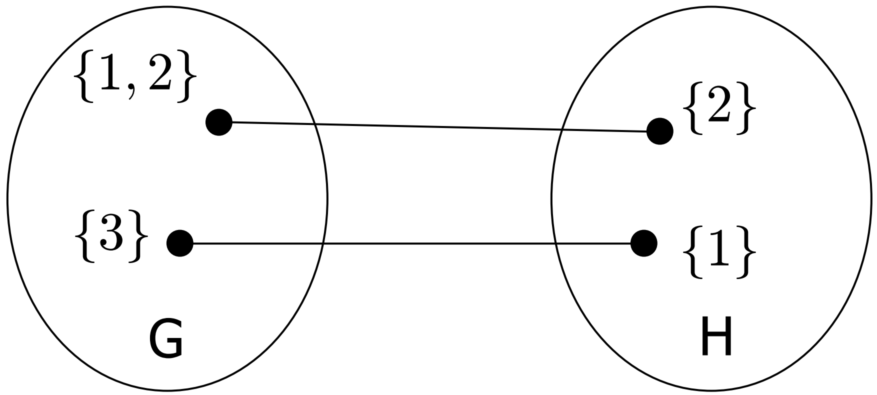

- $R = \{(1,2),(3,1)\}$ is a set of pairs of colours
- No edges are created inside G, H
- g, h are *independent* recolourings of G, H: $g, h : [k] \to \mathcal{P}([k])$

Representation of mcwd terms

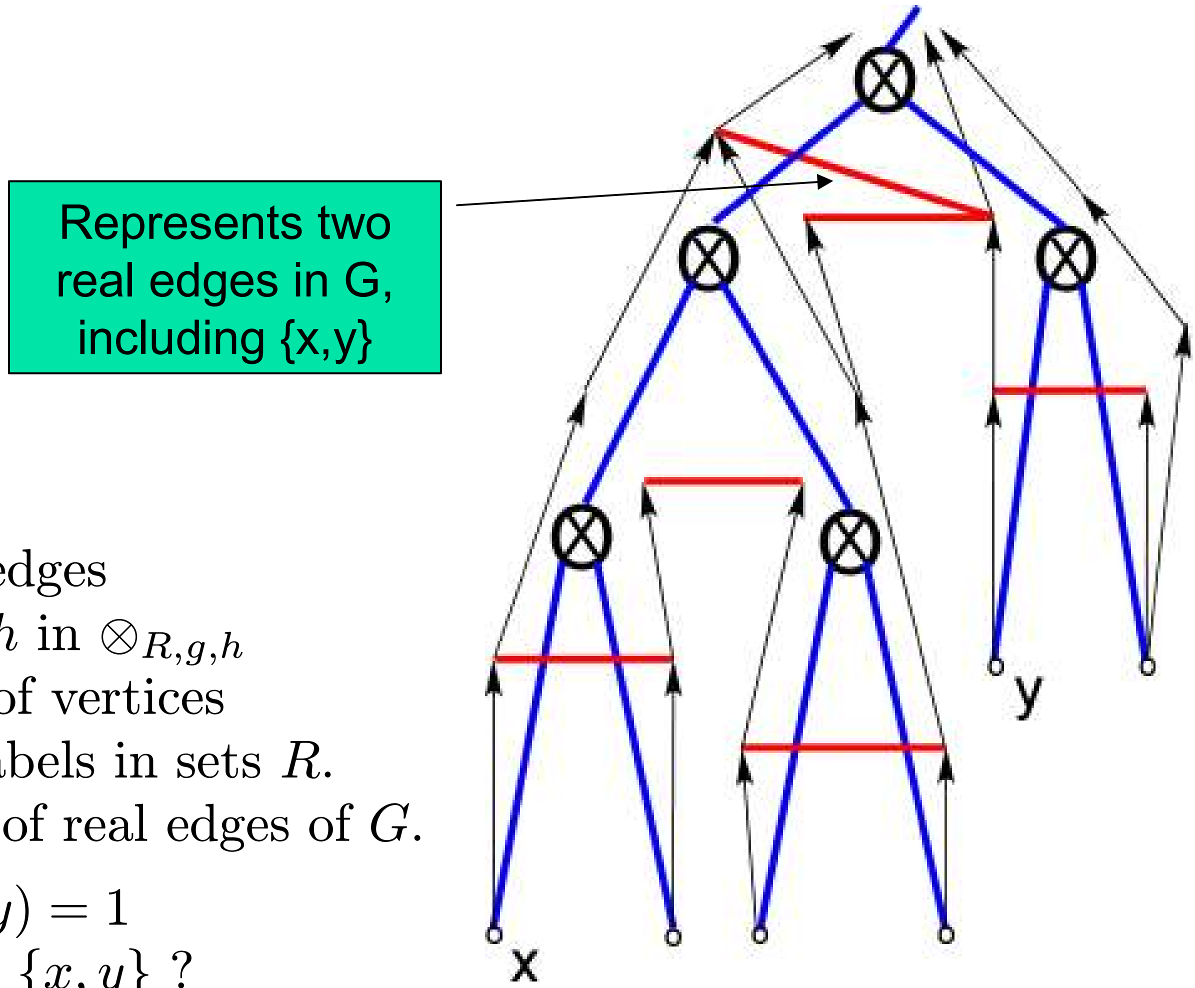

- ★ term t: $\otimes$ and blue edges
- ★ $\uparrow$: relabellings by g, h in $\otimes_{R,g,h}$ and initial labelling of vertices
- ★ red edges: pairs of labels in sets R. Each produces a set of real edges of G.

e.g. Can see that $d(x, y) = 1$
What if we delete edge $\{x, y\}$?

Forbidden edges

For shortest paths leaving a set X, we replace their corresponding subterms with *shortcut* edges in the graph representation of a term, and stored in labels assigned to vertices of G.

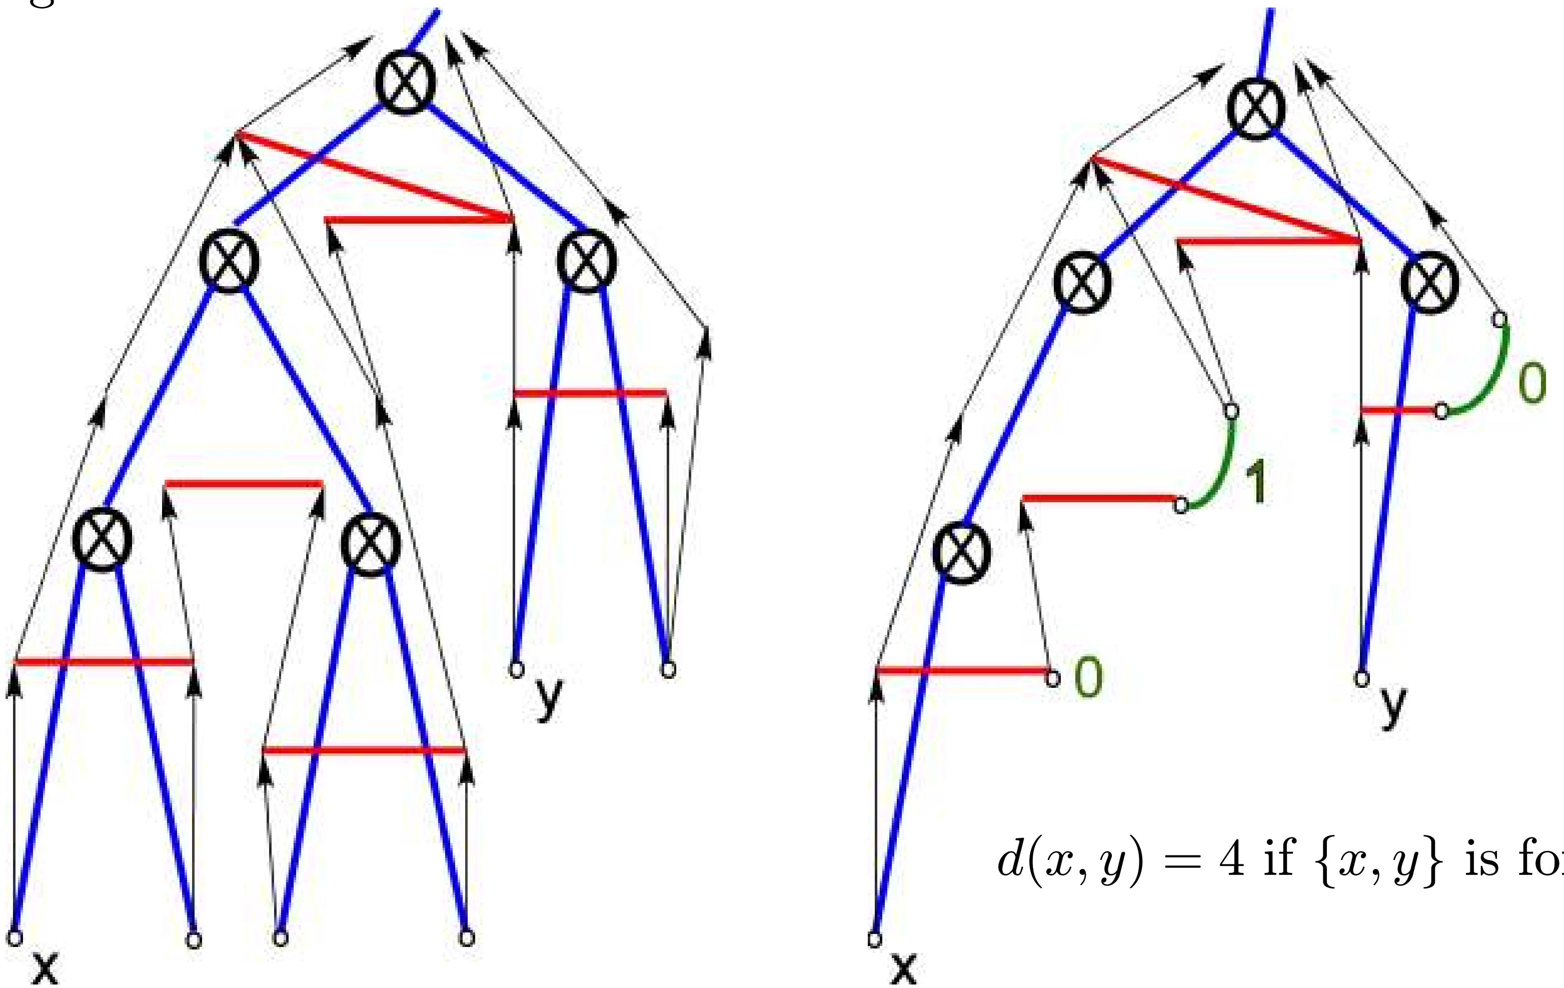

$d(x, y) = 4$ if $\{x, y\}$ is forbidden

Forbidden vertices

We can do a similar thing for deleted vertices. The ellipse touches all edges of G that are adjacent to vertex y of G. E.g. x, y and y, z are adjacent but $d(x, z) = 3$ if y is forbidden.

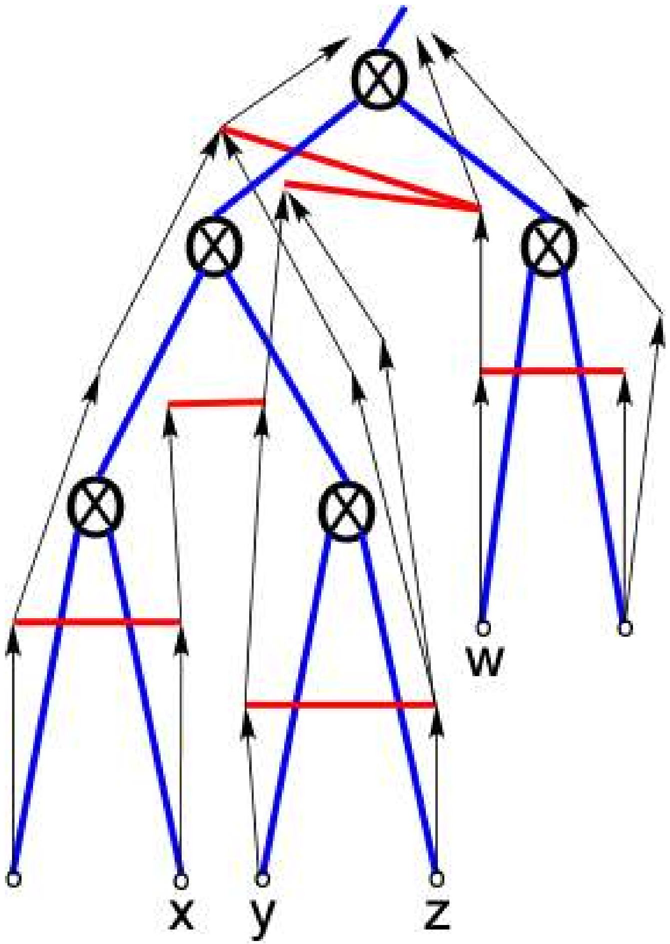

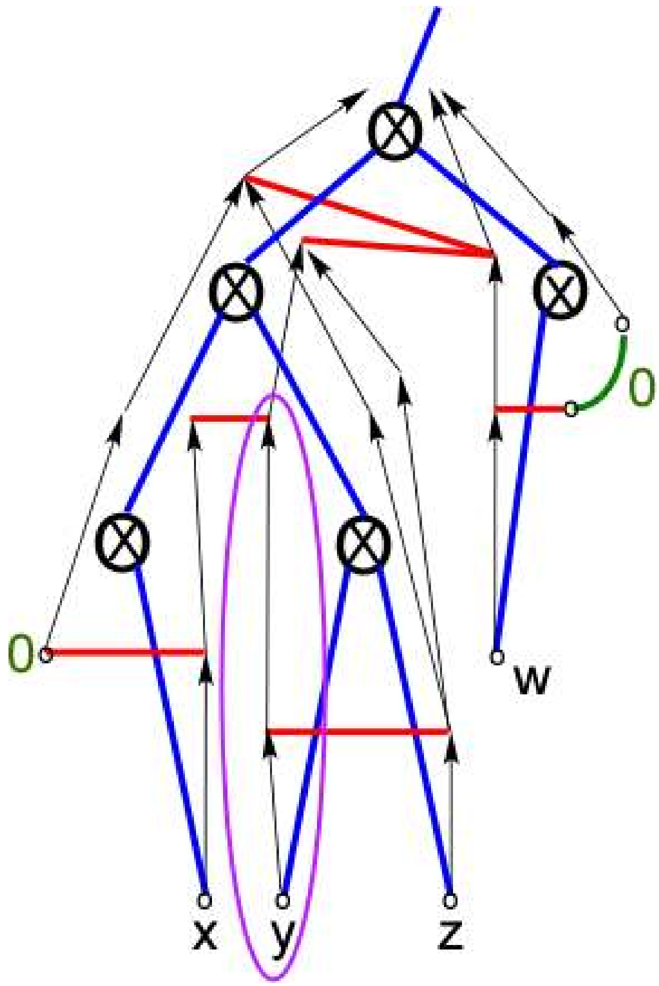

m-clique width

Definition 1 $mcwd(G) = \min\{k : G = val(t)$ *for a term* t *with colours* $\in [k]\}$

Proposition 1

$$\begin{aligned} mcwd(G) &\leq cwd(G) \leq 2^{mcwd(G)+1} \\ mcwd(G) &\leq twd(G) + 3 \end{aligned}$$

The adjacency labelling scheme for cographs extends to graphs defined by mcwd terms.

But: Labels for $x \in X$ give $G[X]$ but to construct $G^{+}[X]$ we need knowledge of paths outside X

Construction of labels

Let G have $mcwd(G) \leq k$. The shortcut edges are represented by a $(k \times k)$ matrix of integers at each occurrence in the term tree. We enrich $A(u)$ into $J(u)$ by inserting at each position corresponding to an occurrence in the term, the associated shortcut matrix.

Lemma 1 *We have forbidden-set distance labels $J(x)$ of size $O(k^2 ht(t) \log n)$ bits.*

Q: How to replace the height $ht(t)$ by $\log n$?
A: Using balanced terms

We need balanced terms

Definition 1 *A term t is a-balanced if $ht(t) \leq a \log |t|$.*

If f is associative then we can make terms balanced by simple reorganization:

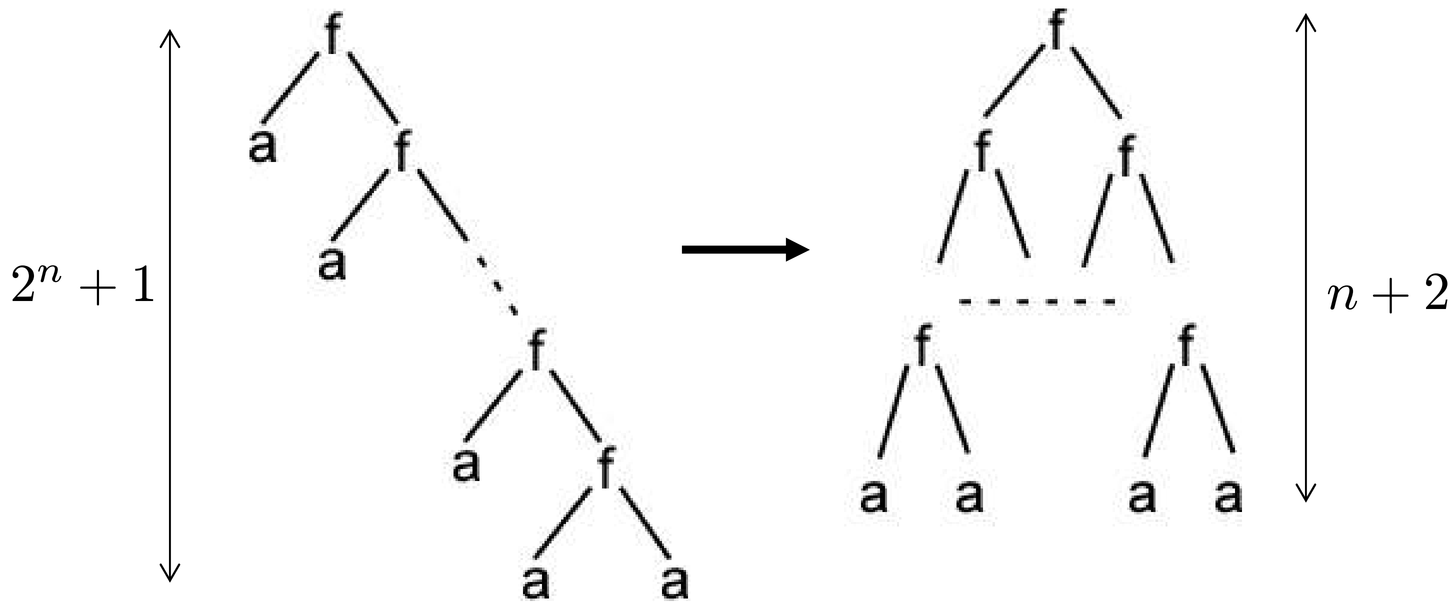

Contexts and terms

A *context* is a term with a special variable u. For nonassociative f, use an explicit substitution operation $\bullet$ on context c and term t where $c \bullet t = c[t/u]$:

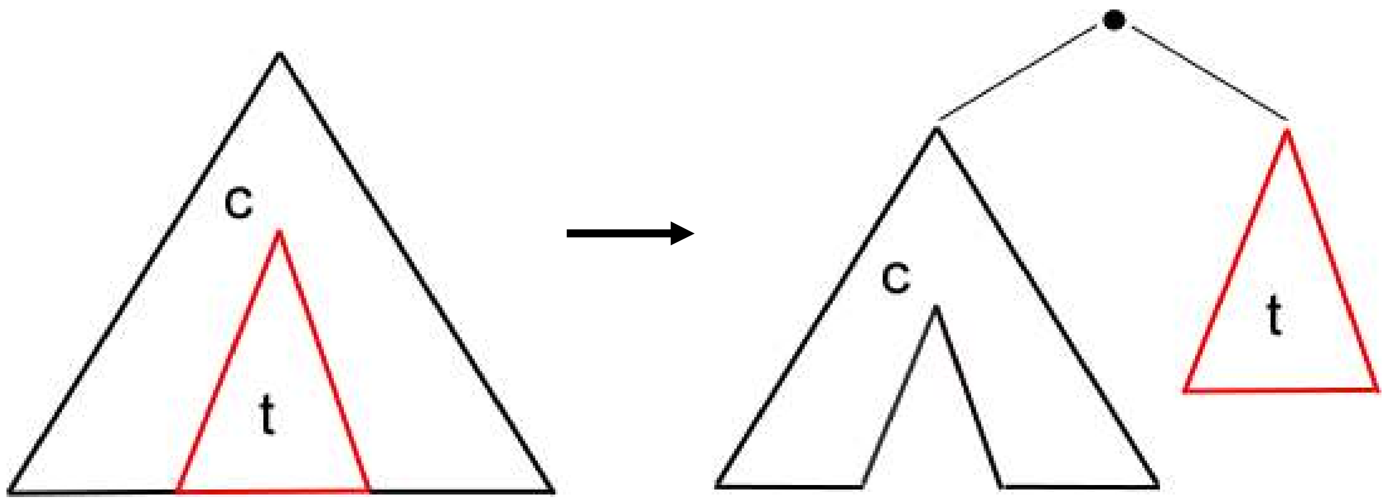

Idea: Choose c, t almost equal in size and recurse

Balanced terms II

Let $\bullet$ be substitution of terms into contexts, $\circ$ be substitution of contexts into contexts and Id the special *identity* context Id with $c \circ \mathsf{Id} = \mathsf{Id} \circ c = c$.
($\circ$ is associative with unit element Id)

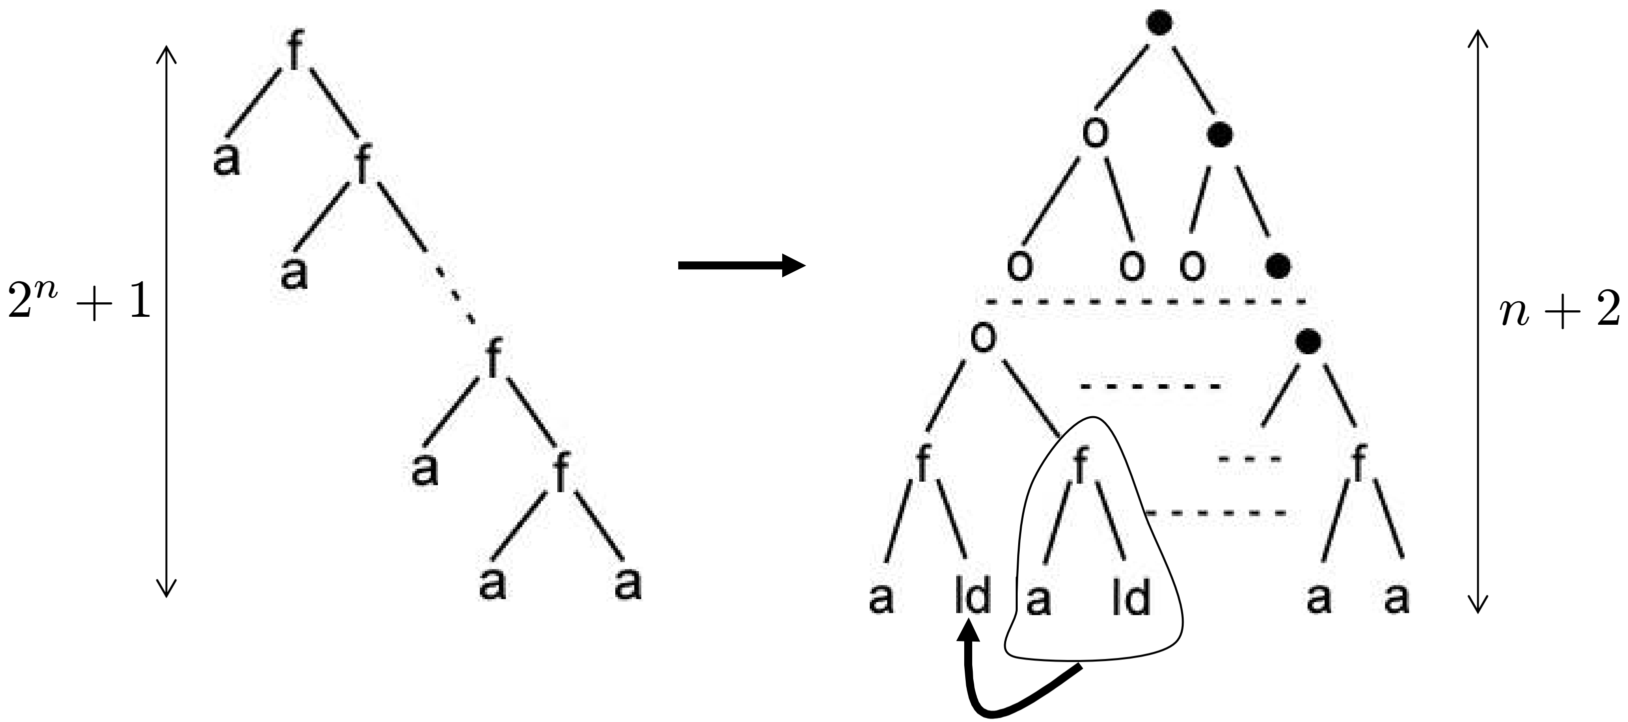

Recursively cut terms using $\bullet$ and contexts using $\circ$

Balanced terms III

Proposition 1 (Courcelle-Vanicat) *Every term in $T(F, C)$ is equivalent to a 3-balanced term in $T(F \cup \{\circ, \bullet\}, C \cup \{Id\})$.*

But this is no longer a cwd term! To get a balanced cwd term has exponential blowup in cwd. For mcwd, we get only a constant blowup:

Lemma 1 *If $mcwd(G) \leq k$ then G can be defined by a 3-balanced mcwd term of width $\leq 2k$*

Balanced terms III

Proposition 1 (Courcelle-Vanicat) *Every term in $T(F, C)$ is equivalent to a 3-balanced term in $T(F \cup \{\circ, \bullet\}, C \cup \{Id\})$.*

But this is no longer a cwd term! To get a balanced cwd term has exponential blowup in cwd. For mcwd, we get only a constant blowup:

Lemma 1 *If $mcwd(G) \leq k$ then G can be defined by a 3-balanced mcwd term of width $\leq 2k$*

Proof. Given mcwd term t of width $\leq k$, use above proposition to obtain a 3-balanced term using $\circ, \bullet$. Interpret $\circ, \bullet$ as mcwd operations on $2k$ labels:

$$c \bullet t_H \equiv G_c \otimes_{R,g,h_c} H$$

Where $t_H \in T(F_k, C_k)$ defines the graph H, G_c is a graph with labels in $[k] \cup \{k+1, \ldots, 2k\}$. Applying this recursively to c and t_H we get a 6-balanced term in $T(F_{2k}, C_{2k})$. The construction is compositional:

$$G_{c \circ d} = G_c \otimes G_d \quad \text{and} \quad h_{c \circ d} = h_c \circ h_d$$

Main result

The balancing lemma allows us to use $ht(t) = O(\log n)$, hence we get

Theorem 1 *For G having mcwd $\leq k$ and n vertices, given $\{J(x) : x \in X\}$ we can compute the shortcut graph $G^+[X]$, where $|J(x)| = O(k^2 \log^2 n)$ bits.*

Since mcwd is more powerful than cwd and twd, we get labels of size $O(k^2 \log^2 n)$ for graphs having $cwd, twd \leq k$.

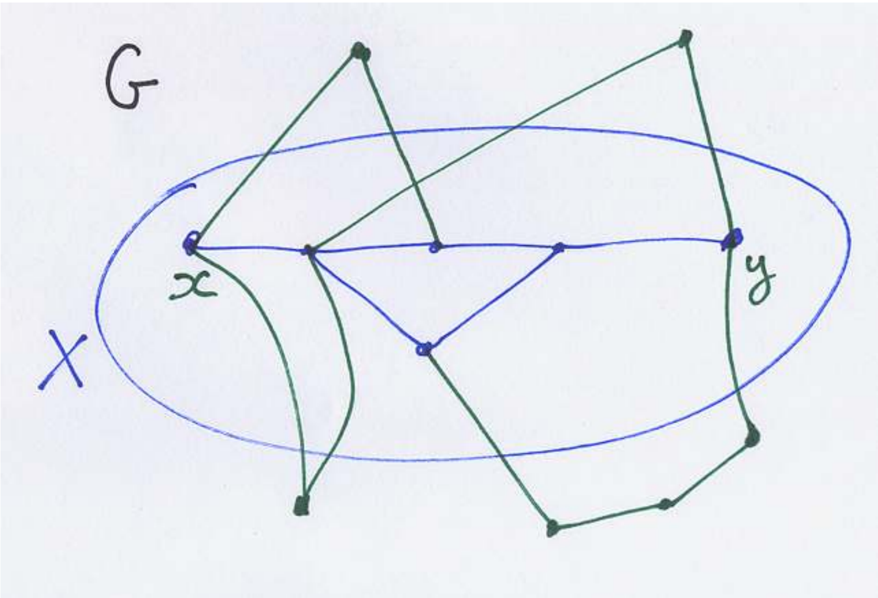

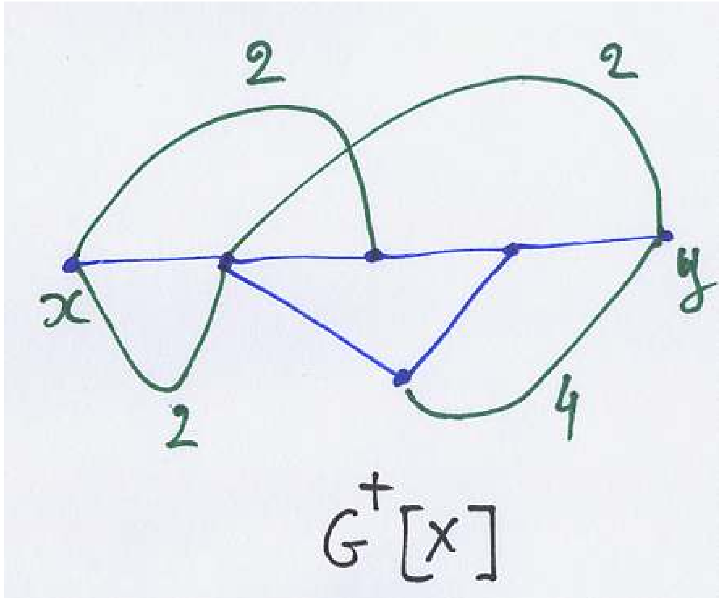

Conclusion

- Forbidden-set distance labelling for mcwd graphs
  - “Distance labelling with obstacles”
  - Can obtain a compact routing scheme
- Techniques
  - Use algebraic representation of graphs
  - Better describes *how* the graph is connected
  - Mcwd terms have good balancing properties

Open problems

  - Planar graphs??