




1. Conflicts & Trust mappings in Community DBs

Background 1: Conflicting beliefs*



Background 2: Trust mappings



Recent work on community databases:

- Orchestra [SIGMOD'06, VLDB'07]
- Youtopia [VLDB'09], BeliefDB [VLDB'09]

* Current state of knowledge on the Indus Script: Rao et al., Science 324(5931):1165, May 2009

2. Stable solutions

Priority trust network (TN)


- assume a fixed key
- users (nodes): A, B, C
- values (beliefs): v, w, u
- trust mappings (arcs) from "parents"

Stable solution

- assignment of values to each node*, s.t. each belief has a "non-dominated lineage" to an explicit belief

Possible / Certain semantics

- a stable solution determines, for each node, a possible value ("poss")
- certain value ("cert") = intersection of all stable solutions, per user



$$SS1 = \{A:v, B:w, C:v, D:v, E:u\}$$


$$SS2 = \{A:v, B:w, C:w, D:w, E:u\}$$

X	poss(X)	cert(X)
A	{v}	{v}
B	{w}	{w}
C	{v,w}	∅
D	{v,w}	∅
E	{u}	{u}

* each node with at least one ancestor with explicit belief

3. Logic programs with stable model semantics

Step 1: Binarization



Step 2: Logic program

1: accept all poss of preferred parent

$$P(C,x) \leftarrow P(A,x)$$

$$F(C,B,y) \leftarrow P(B,y), P(C,x), x \neq y$$

$$P(C,y) \leftarrow P(B,y), \neg F(C,B,y)$$

2: accept poss from non-preferred parent, that are not conflicting with an existing value

$$F(C,A,y) \leftarrow P(A,y), P(C,x), x \neq y$$


$$P(C,y) \leftarrow P(A,y), \neg F(C,A,y)$$

$$F(C,B,y) \leftarrow P(B,y), P(C,x), x \neq y$$

$$P(C,y) \leftarrow P(B,y), \neg F(C,B,y)$$

4. Resolution Algorithm (1/2)


- Keep 2 sets: closed / open
- Initialize closed with explicit beliefs
- MAIN
- Step 1: if \exists preferred edges from open to closed \rightarrow follow



X	poss(X)	cert(X)
A	{v}	{v}
B	{w}	{w}
C	{u}	{u}
D	{v}	{v}
E	?	?
F	?	?
G	?	?
H	?	?
I	?	?
J	?	?
K	?	?
L	?	?

5. Resolution Algorithm (2/2)

- Keep 2 sets: closed / open
- Initialize closed with explicit beliefs
- MAIN
- Step 1: if \exists preferred edges from open to closed \rightarrow follow
- Step 2: else \rightarrow construct SCC graph of open \rightarrow resolve minimum SCCs



"Minimal SCC" can be calculated in $O(n)$

Tarjan [1972]

PTIME resolution algorithm

$O(n^2)$ worst case

$O(n)$ on reasonable graphs


X	poss(X)	cert(X)
A	{v}	{v}
B	{w}	{w}
C	{u}	{u}
D	{v}	{v}
E	{w}	{w}
F	{u}	{u}
G	{v,w}	∅
H	{v,w}	{w}
I	{v,w}	∅
J	{v,w}	∅
K	{v,w}	∅
L	?	?

7. Experiments on large network data

Calculating poss / cert for fixed key

- DLV: State-of-the-art logic programming solver
- RA: Resolution algorithm

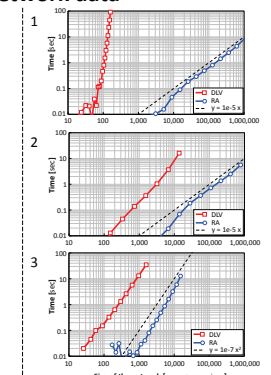

Network 1: "Oscillators"



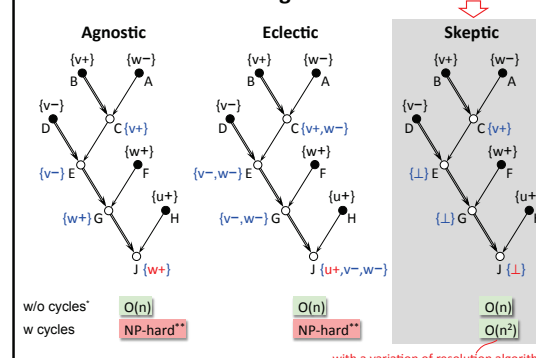
Network 2: "Web link data"

- Web data set with 5.4m links between 270k domain names. Approach:
- Sample links with increasing ratio
- Include both nodes in sample
- Assign explicit beliefs randomly

Network 3: "Worst case" $O(n^2)$



8. Three semantics for negative beliefs




* assuming total order on parents for each node

** checking if a belief is possible at a given node is NP-hard, checking if it is certain is co-NP-hard

6. Detail: Strongly Connected Components (SCCs)

For every cyclic or acyclic directed graph:

- The Strongly Connected Components graph is a DAG
- can be calculated in $O(n)$ Tarjan [1972]



9. Take-aways automatic conflict resolution

Problem

- Given explicit beliefs & trust mappings, how to assign consistent value assignment to users?

Our solution

- Stable solutions with possible/certain value semantics
- PTIME algorithm [$O(n^2)$ worst case, $O(n)$ experiments]
- Several extensions
 - negative beliefs: 3 semantics, two hard, one $O(n^2)$
 - bulk inserts
 - agreement checking
 - consensus value
 - lineage computation

Slides soon available on our project page:

<http://db.cs.washington.edu/beliefDB>