

Negation, Coordination and Monotonicity

Bas Ketsman



SOFTWARE
LANGUAGES
LAB




DATABASE
GROUP

Roadmap


- **Context**
 - From Monotone to Negation-Free Queries
 - Non-rewritable Monotone Queries
 - Coordination-Freedom and System Knowledge
 - Open Problems

CALM


Programming asynchronous
distributed systems in a declarative
language



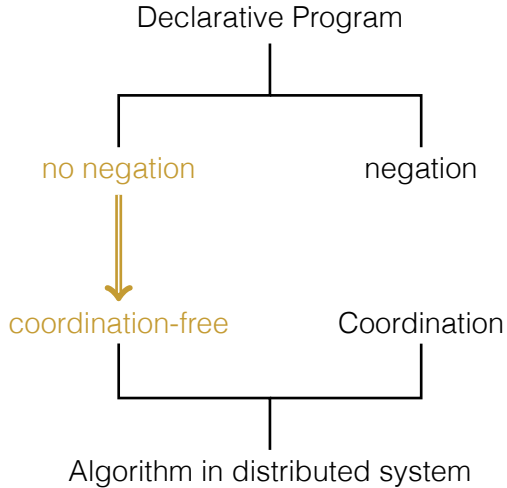
CALM



Programming asynchronous distributed systems in a declarative language

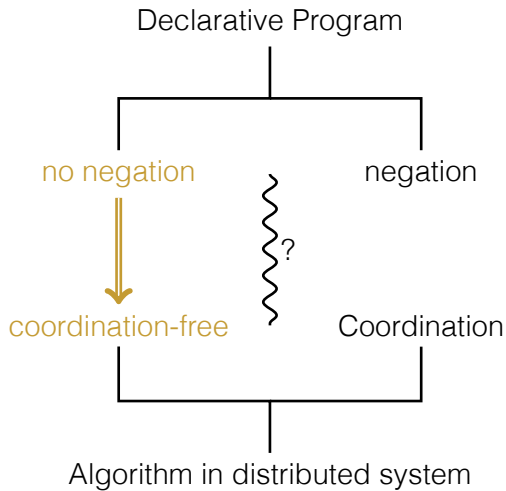


CALM



Programming asynchronous distributed systems in a declarative language

CALM



Programming asynchronous distributed systems in a declarative language

Theorem (CALM Theorem)

Monotone Programs =
Coordination-Free Programs

- Conjecture: [Hellerstein 2010]
- Proof for queries: [Ameloot, Neven, Van den Bussche 2011]


Formalization

[Ameloot, Neven, Van den Bussche 2011]

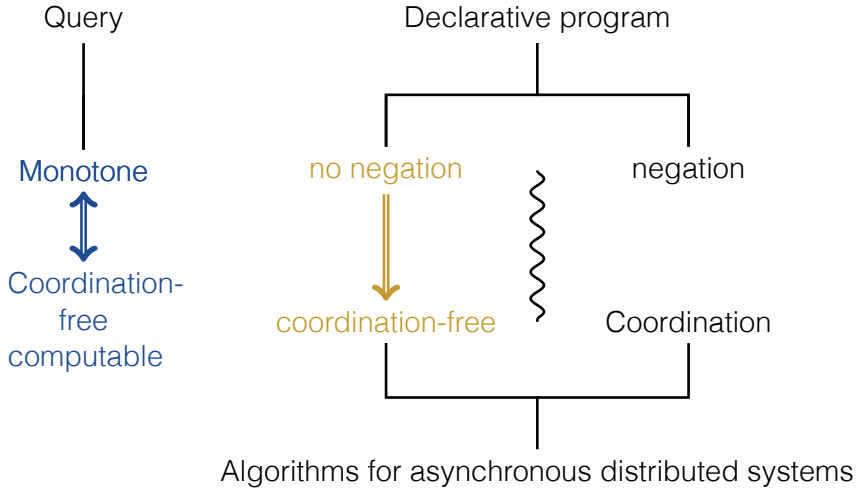
- Programs computing **queries**
 - ▶ Semantic objects
 - ▶ data-, network-, partition-independent
- **Asynchronous shared-nothing** system with at-least once message arrival guarantees
 - ▶ Write-only output relations
- An algorithm computes a query if the output relations **eventually** represent the output
- An algorithm is **coordination-free** if for every network and input database there is an “ideal” data partitioning for which the algorithm generates the query answer before any **communication** is done.

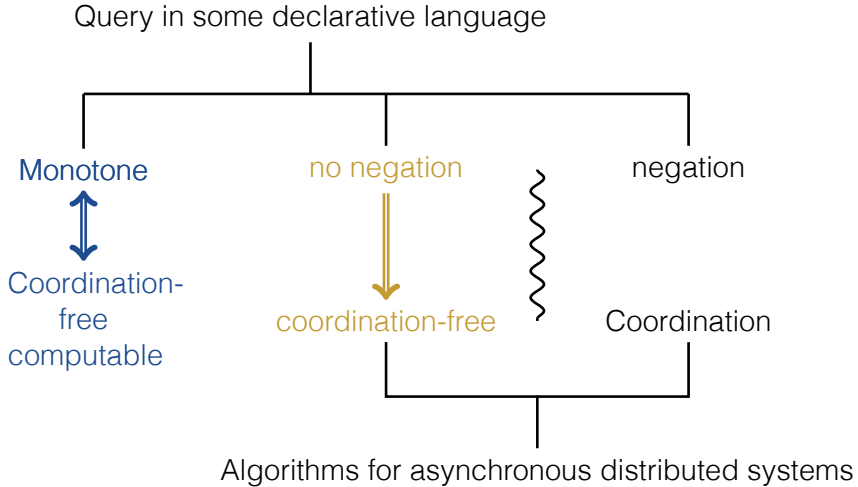
Formalization

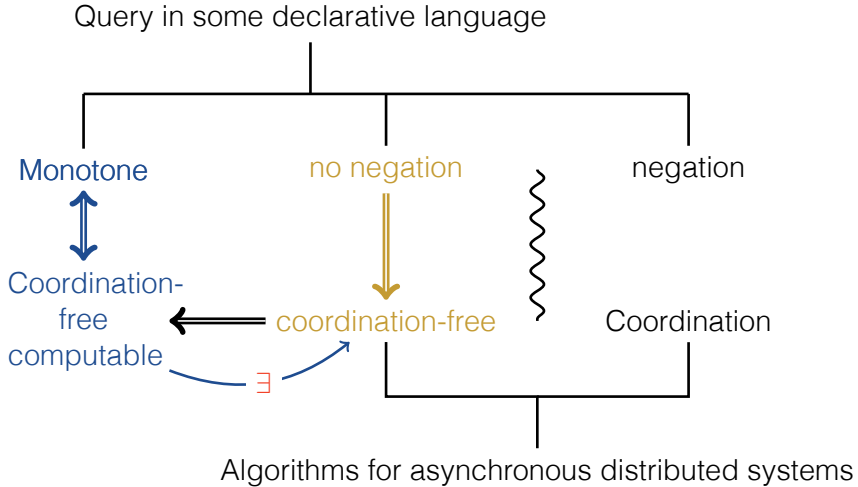
[Ameloot, Neven, Van den Bussche 2011]

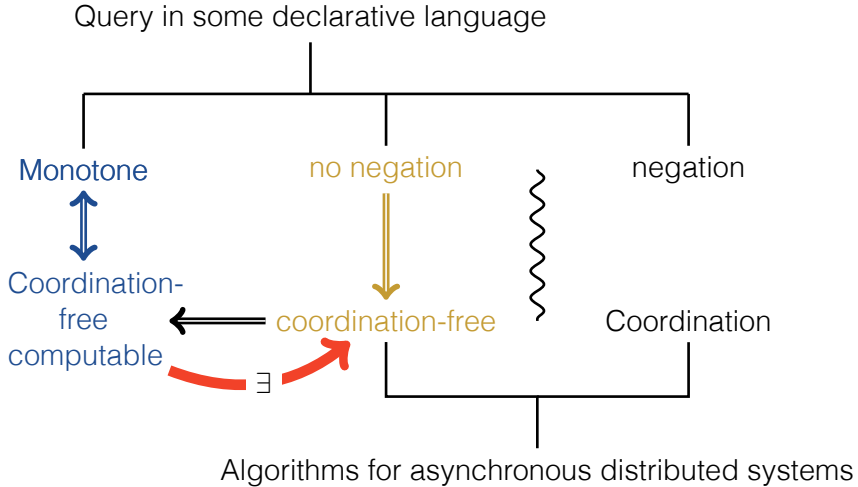


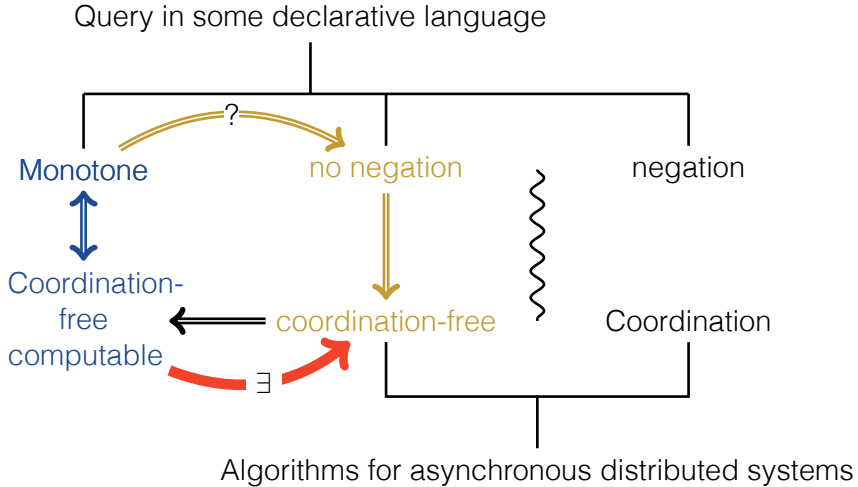
- Programs computing **queries**
 - ▶ Semantic objects
 - ▶ data-, network-, partition-independent
- **Asynchronous shared-nothing** system with at-least once message arrival guarantees
 - ▶ Write-only output relations
- An algorithm computes a query if the output relations **eventually** represent the output
- An algorithm is **coordination-free** if for every network and input database there is an “ideal” data partitioning for which the algorithm generates the query answer before any **communication** is done.





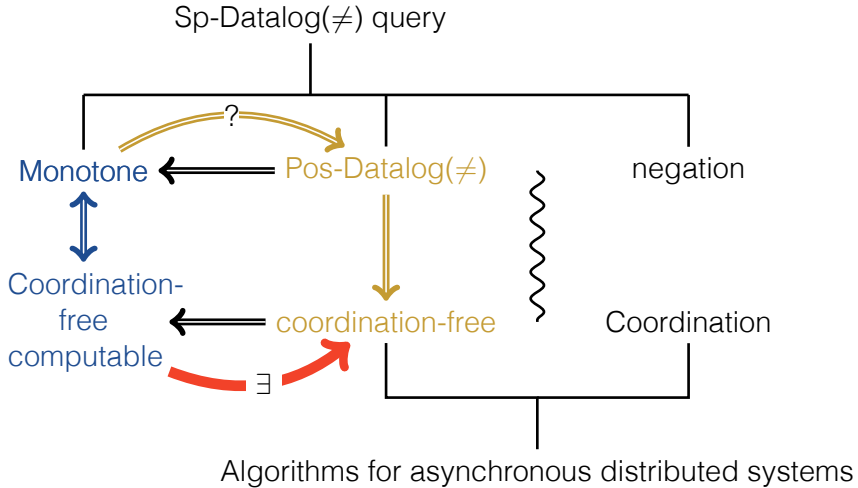






Roadmap

- Context
- **From Monotone to Negation-Free Queries**
- Non-rewritable Monotone Queries
- Coordination-Freedom and System Knowledge
- Open Problems



Textbook Datalog

[Afrati, Cosmadakis, Yannakakis 1995]

Positive Datalog Program

Monotone

Pos-Datalog

Preserved under Homomorphisms

Pos-Datalog(\neq)

Monotone

Datalog with Negation

Not necessarily Monotone

Sp-Datalog

Semi-Monotone / Preserved under Extensions

Sp-Datalog(\neq)

Semi-Monotone / Preserved under Extensions

Str-Datalog

A Monotone Sp-Datalog(\neq) Example

Input: Graph with green (R), blue (S), and black (T) edges

Sp-Datalog Program P :

$$T_1(x, y) \leftarrow R(x, y), \neg S(y, z), T(z, x).$$


$$T_2(x, y) \leftarrow R(x, y), S(y, z), \neg T(z, x).$$

$$O() \leftarrow T_1(x, y), T_2(x, y).$$

$$O() \leftarrow R(x, y), S(y, z), T(z, x), x \neq y.$$

Output: Binary output relation O

A Monotone Sp-Datalog(\neq) Example



Input: Graph with green (R), blue (S), and black (T) edges

Sp-Datalog Program P :

$$T_1(x, y) \leftarrow R(x, y), \neg S(y, z), T(z, x).$$


$$T_2(x, y) \leftarrow R(x, y), S(y, z), \neg T(z, x).$$

$$O() \leftarrow T_1(x, y), T_2(x, y).$$

$$O() \leftarrow R(x, y), S(y, z), T(z, x), x \neq y.$$

Output: Binary output relation O

Naive Rewriting: Remove Negated Atoms



Input: Graph with green (R), blue (S), and black (T) edges

Sp-Datalog Program P :

$$T_1(x, y) \leftarrow R(x, y), \neg S(y, z), T(z, x).$$

$$T_2(x, y) \leftarrow R(x, y), S(y, z), \neg T(z, x).$$


$$O() \leftarrow T_1(x, y), T_2(x, y).$$

$$O() \leftarrow R(x, y), S(y, z), T(z, x), x \neq y.$$

Output: Binary output relation O

Candidate Proof-Trees with Fringe-Conflicts



(Proof-trees decoupled from a particular database instance)



Proof-tree for P^+

Candidate Proof-Trees with Fringe-Conflicts


(Proof-trees decoupled from a particular database instance)




Proof-tree for P^+

Candidate Proof-Trees with Fringe-Conflicts

(Proof-trees decoupled from a particular database instance)



Proof-tree for P^+



Candidate proof-tree^a for P
with fringe conflicts

(^aIn paper: candidate proof-tree without inequality conflicts)

Proof-Tree Perspective

Theorem

For monotone Sp-Datalog_{conflict-free}(\neq) programs whose candidate proof-trees are without fringe conflicts, negated atoms can be left away without influencing the program semantics.

Making Programs Conflict-Free


Crux: Pass information about forbidden facts to other rules via IDBs and exclude conflicts via disequalities

$$T_1(x, y) \leftarrow R(x, y), \neg S(y, z), T(z, x).$$

$$T_2(x, y) \leftarrow R(x, y), S(y, z), \neg T(z, x).$$

$$O() \leftarrow T_1(x, y), T_2(x, y).$$

$$O() \leftarrow R(x, y), S(y, z), T(z, x), x \neq y.$$



Making Programs Conflict-Free


Crux: Pass information about forbidden facts to other rules via IDBs and exclude conflicts via disequalities

$$T_1(x, y, \textcolor{blue}{y}, z) \leftarrow R(x, y), \neg S(\textcolor{blue}{y}, z), T(z, x).$$

$$T_2(x, y) \leftarrow R(x, y), S(y, z), \neg T(z, x).$$

$$O() \leftarrow T_1(x, y, \textcolor{blue}{y}', z'), T_2(x, y).$$

$$O() \leftarrow R(x, y), S(y, z), T(z, x), x \neq y.$$



Making Programs Conflict-Free

Crux: Pass information about forbidden facts to other rules via IDBs and exclude conflicts via disequalities


$\text{Adom}(x) \leftarrow$ “active domain of R , S , and T ”

$T_1(x, y, y, z) \leftarrow R(x, y), \neg S(y, z), T(z, x).$

$T_2(x, y, y', z') \leftarrow R(x, y), S(y, z), \neg T(z, x), \text{Adom}(y'), \text{Adom}(z').$

$O() \leftarrow T_1(x, y, y', z'), T_2(x, y, y', z').$

$O() \leftarrow R(x, y), S(y, z), T(z, x), x \neq y.$



Making Programs Conflict-Free

Crux: Pass information about forbidden facts to other rules via IDBs and exclude conflicts via disequalities

$\text{Adom}(x) \leftarrow$ “active domain of R , S , and T ”


$T_1(x, y, y, z) \leftarrow R(x, y), \neg S(y, z), T(z, x).$

$T_2(x, y, y', z') \leftarrow R(x, y), S(y, z), \neg T(z, x), \text{Adom}(y'), \text{Adom}(z'), y \neq y'.$

$T_2(x, y, y', z') \leftarrow R(x, y), S(y, z), \neg T(z, x), \text{Adom}(y'), \text{Adom}(z'), z \neq z'.$

$O() \leftarrow T_1(x, y, y', z'), T_2(x, y, y', z').$

$O() \leftarrow R(x, y), S(y, z), T(z, x), x \neq y.$



Making Programs Conflict-Free

Crux: Pass information about forbidden facts to other rules via IDBs and exclude conflicts via disequalities

$\text{Adom}(x) \leftarrow$ “active domain of R , S , and T ”


$T_1(x, y, y, z) \leftarrow R(x, y), \neg S(y, z), T(z, x).$

$T_2(x, y, y', z') \leftarrow R(x, y), S(y, z), \neg T(z, x), \text{Adom}(y'), \text{Adom}(z'), y \neq y'.$

$T_2(x, y, y', z') \leftarrow R(x, y), S(y, z), \neg T(z, x), \text{Adom}(y'), \text{Adom}(z'), z \neq z'.$

$O() \leftarrow T_1(x, y, y', z'), T_2(x, y, y', z').$

$O() \leftarrow R(x, y), S(y, z), T(z, x), x \neq y.$



Making Programs Conflict-Free

Crux: Pass information about forbidden facts to other rules via IDBs and exclude conflicts through disequalities

$\text{Adom}(x) \leftarrow$ “active domain of R , S , and T ”

$T_1(x, y, y, z, z'', x'') \leftarrow R(x, y), \neg S(y, z), T(z, x), \text{Adom}(z''), \text{Adom}(x''), z \neq z''.$

$T_1(x, y, y, z, z'', x'') \leftarrow R(x, y), \neg S(y, z), T(z, x), \text{Adom}(z''), \text{Adom}(x''), x \neq x''.$

$T_2(x, y, y', z', z, x) \leftarrow R(x, y), S(y, z), \neg T(z, x), \text{Adom}(y'), \text{Adom}(z'), y \neq y'.$

$T_2(x, y, y', z', z, x) \leftarrow R(x, y), S(y, z), \neg T(z, x), \text{Adom}(y'), \text{Adom}(z'), z \neq z'.$

$O() \leftarrow T_1(x, y, y', z', z'', x''), T_2(x, y, y', z', z'', x'').$

$O() \leftarrow R(x, y), S(y, z), T(z, x), x \neq y.$

Equivalent to original program and candidate proof-trees have no fringe conflicts

Making Programs Conflict-Free

Crux: Pass information about forbidden facts to other rules via IDBs and exclude conflicts with disequalities

$\text{Adom}(x) \leftarrow$ “active domain of R , S , and T ”

$$T_1(x, y, y, z, z'', x'') \leftarrow R(x, y), \neg S(y, z), T(z, x), \text{Adom}(z''), \text{Adom}(x''), z \neq z''.$$

$$T_1(x, y, y, z, z'', x'') \leftarrow R(x, y), \neg S(y, z), T(z, x), \text{Adom}(z''), \text{Adom}(x''), x \neq x''.$$

$$T_2(x, y, y', z', z, x) \leftarrow R(x, y), S(y, z), \neg T(z, x), \text{Adom}(y'), \text{Adom}(z'), y \neq y'.$$

$$T_2(x, y, y', z', z, x) \leftarrow R(x, y), S(y, z), \neg T(z, x), \text{Adom}(y'), \text{Adom}(z'), z \neq z'.$$

$$O() \leftarrow T_1(x, y, y', z', z'', x''), T_2(x, y, y', z', z'', x'').$$

$$O() \leftarrow R(x, y), S(y, z), T(z, x), x \neq y.$$

Equivalent to original program, even without negated atoms

Negation-Bounded Datalog

When does such a rewriting exist?

Negation-Bounded Datalog

When does such a rewriting exist?

- For all programs with a **bounded number of forbidden facts** in their proof-tree fringes;
 - [K, Koch, 2020]
- For all programs with a **bounded number of required facts** in their proof-tree fringes.
 - Symmetric case
- For all **programs without non-equalities**
 - ▶ Because $\text{Sp-Datalog} \cap \text{Monotone} = \text{Sp-Datalog} \cap \text{Hom}$
 - ▶ $\text{Sp-Datalog} \cap \text{Hom} = \text{Pos-Datalog}$ [Feder, Vardi, 2003]

Roadmap

- Context
- From Monotone to Negation-Free Queries
- **Non-rewritable Monotone Queries**
- Coordination-Freedom and System Knowledge
- Open Problems

Counterexample Ingredients

Heavily based on: [Rudolph, Thomazo 2015]

Definition

Given a graph, the **perfect matching problem** asks if there is a subset M of its edges such that every vertex of the graph is incident to **precisely one** edge in M .

- A polynomial-time computable query.
- All polynomial-time computable queries are expressible in order-invariant SP-datalog [Papadimitriou 1985]
 - ▶ Succ[2] encoding some linear order over adom;
 - ▶ Min[1] and Max[1] denoting minimal and maximal element.

Separating Query

Input: instance over $\{\text{Edge}[2], \text{Succ}[2], \text{Min}[1], \text{Max}[1]\}$.

Output:

- **True** if Succ, Min, Max encode an **inconsistency**
- **True** if Succ, Min, Max **contain** a complete total order and **the graph induced by the active domain elements between Min and Max in Next has a perfect matching**.
- **False** (\emptyset) otherwise.

Sp-Datalog(\neq) query | monotone

Separating Query

Input: instance over $\{\text{Edge}[2], \text{Succ}[2], \text{Min}[1], \text{Max}[1]\}$.

Output:

- **True** if Succ, Min, Max encode an **inconsistency**
 - ▶ Some negation-free rules
- **True** if Succ, Min, Max **contain** a complete total order and **the graph induced by the active domain elements between Min and Max in Next has a perfect matching**.
 - ▶ Some negation-free rules + order invariant SP-Datalog program with Succ, Min, Max given as part of input
- **False** (\emptyset) otherwise.

Sp-Datalog(\neq) query | monotone

Separating Query

Input: instance over $\{\text{Edge}[2], \text{Succ}[2], \text{Min}[1], \text{Max}[1]\}$.

Output:

- **True** if Succ, Min, Max encode an **inconsistency**
 - ▶ Some negation-free rules
 - ▶ We cannot fix such an inconsistency by adding more facts
- **True** if Succ, Min, Max **contain** a complete total order and **the graph induced by the active domain elements between Min and Max in Next has a perfect matching**.
 - ▶ Some negation-free rules + order invariant SP-Datalog program with Succ, Min, Max given as part of input
 - ▶ We cannot undo existence of a perfect matching by adding more facts
- **False** (\emptyset) otherwise.

Sp-Datalog(\neq) query | monotone

Separating Query $\not\in$ Pos-Datalog(\neq)

Theorem ([Razborov 85])

No family of monotone Boolean circuits exists that answer the perfect matching problem and has circuits of polynomial size in the number of input gates.

Claim: If there is a Pos-Datalog(\neq) program P expressing the separating query, then there is a Boolean circuit contradicting the above result.

Proof by construction

For a graph with v vertices, consider a Boolean circuit with v^2 input gates, one per possible edge. The gate is set to 1 iff the edge is in the graph.

1. Consider all groundings of rules in P over domain $\{1, \dots, v\}$.
2. Assume natural order over these values as interpretation for **First**, **Last**, **Next** and remove all grounded rules that contradict with it.
3. Remove all facts over interpreted relation names.
4. Make copies of rules organized in strata. First strata copies all EDBS to index 0. For next stratum all atoms in body have index i and those in heads have index $i + 1$. We need only as many strata as we need rounds to finish programs over domain $\{1 \dots v\}$.
5. EDB atoms = input gates, rules are AND gates whose inputs are the gates whose output represents body atoms; IDBs become OR gates whose inputs are the gates represented by rules that produce the ground atoms.

Roadmap

- Context
- From Monotone to Negation-Free Queries
- Non-rewritable Monotone Queries
- **Coordination-Freedom and System Knowledge**
- Open Problems

A More Fine-Grained Answer to the CALM Conjecture

Systems in which data is **arbitrarily** partitioned.

Theorem ([Ameloot, Neven, Van den Bussche, 2011])

$$\mathcal{C}_q(\text{Id} + \text{A11}) \cap \text{Coordination-Free} = \text{Monotone}$$

A More Fine-Grained Answer to the CALM Conjecture

Systems in which data is **arbitrarily** partitioned.

Theorem ([Ameloot, Neven, Van den Bussche, 2011])

$$\mathcal{C}_q(\text{Id} + \text{All}) \cap \text{Coordination-Free} = \text{Monotone}$$

Systems in which data is partitioned using a **tuple-based** policy.

- For every possible tuple in the database there is a node that knows it would store that tuple if it were part of the database

Theorem ([Ameloot, K, Neven, Zinn, 2014])

$$\mathcal{C}_q(\text{Id} + \text{All} + \text{Policy}) \cap \text{Coordination-Free} = \text{Semi-Monotone}$$

A More Fine-Grained Answer to the CALM Conjecture

Systems in which data is **arbitrarily** partitioned.

Theorem ([Ameloot, Neven, Van den Bussche, 2011])

$$\mathcal{C}_q(\text{Id} + \text{All}) \cap \text{Coordination-Free} = \text{Monotone}$$

Theorem ([Ameloot, K, Neven, Zinn, 2014])

$$\mathcal{C}_q(\text{Id} + \text{All} + \text{Policy}) \cap \text{Coordination-Free} = \text{Semi-Monotone}$$

Systems in which data is partitioned using a **value-based** strategy.

- For every possible data value in the database there is a node that knows it would store all tuples in the database containing that data value.

Theorem ([Ameloot, K, Neven, Zinn, 2014])

$$\mathcal{C}_q(\text{Id} + \text{All} + \text{AdomPolicy}) \cap \text{Coordination-Free} = \text{Disjoint-Monotone}$$

From Queries to Behaviors


Definition

A **distributed behavior** is a **non-deterministic** mapping from distributed instances to distributed instances

From Queries to Behaviors

Definition

A **distributed behavior** is a **non-deterministic** mapping from distributed instances to distributed instances



(Queries are specific types of distributed behaviors)

Configuration Constraints

Information made available to nodes in the form of interpreted relations or functions defined as a **behavior**.

Configuration Constraints

Information made available to nodes in the form of interpreted relations or functions defined as a **behavior**.

[Ameloot, Neven, Van den Bussche 2011]

- **Id**: relation containing the unique ID of the node at hand;
- **All**: relation containing all IDs of nodes in the network;

Configuration Constraints

Information made available to nodes in the form of interpreted relations or functions defined as a **behavior**.

[Ameloot, Neven, Van den Bussche 2011]

- **Id**: relation containing the unique ID of the node at hand;
- **All**: relation containing all IDs of nodes in the network;

[Ameloot, K, Neven, Zinn 2014]

- **Policy**: function mapping a fact to true iff the node is responsible for it;
- **AdomPolicy**: similar as previous but based on data values of tuples;

Configuration Constraints

Information made available to nodes in the form of interpreted relations or functions defined as a **behavior**.

[Ameloot, Neven, Van den Bussche 2011]

- **Id**: relation containing the unique ID of the node at hand;
- **All**: relation containing all IDs of nodes in the network;

[Ameloot, K, Neven, Zinn 2014]

- **Policy**: function mapping a fact to true iff the node is responsible for it;
- **AdomPolicy**: similar as previous but based on data values of tuples;

Many other options:


- **L**: relation containing the ID of some node considered leader in the network;
- **Order**: function mapping two data values a, b on true iff $a < b$ according to some coordinated order.
- ...

(can be combined)

Computing Behaviors

Is a certain behavior computable in our model for a specific set of constraints?

Is leader election computable without configuration constraints?



Computing Behaviors

Is a certain behavior computable in our model for a specific set of constraints?


Is leader election computable with Id ?



Computing Behaviors

Is a certain behavior computable in our model for a specific set of constraints?

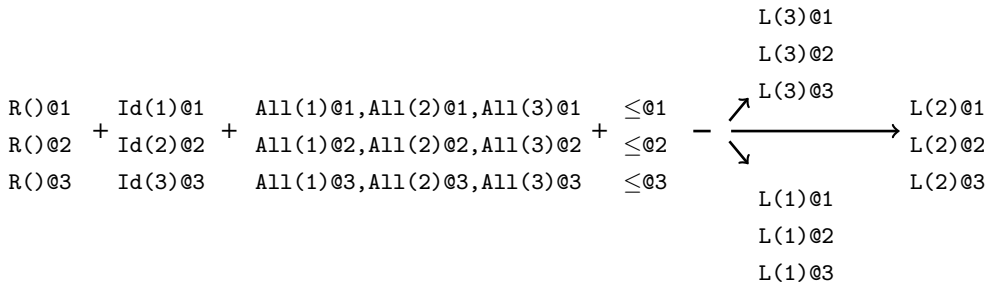
Is leader election computable with Id + All?



Computing Behaviors

Is a certain behavior computable in our model for a specific set of constraints?

Is leader election computable with Id + All + Order?



Computing Behaviors

Definition

$\mathcal{C}(\chi)$ = the set of all behaviors computable in the model with constraint χ .

- $\mathcal{C}(\text{Id} + \text{All})$ contains all (computable) **queries**;
- $\mathcal{C}(\text{Id} + \text{All} + \text{Order})$ equals all (computable) **behaviors**;

Coordination-Freedom for Behaviors

Definition

An algorithm is **coordination-free** if for some (ideal) partitioning of the input data, the algorithm finds the correct output before doing any **communication**.

Coordination-Freedom for Behaviors

Definition

An algorithm is **coordination-free** if for some (ideal) partitioning of the input data, the algorithm finds the correct output before doing any **communication**.

This definition only makes sense for behaviors that are data-partitioning independent**

Semantics Versus Syntactic Transducers

Coordination-Free computability = Computability in restricted variant of the model.

Theorem ([Ameloot, Neven, Van den Bussche, 2011])

$$- \mathcal{C}_{\text{queries}}(\text{Id} + \text{A11}) \cap \text{Coordination-Free} = \mathcal{C}_{\text{queries}}(\text{Id})$$

Coordination Free = Ability to compute without access to A11.

Semantics Versus Syntactic Transducers

Coordination-Free computability = Computability in restricted variant of the model.

Theorem ([Ameloot, Neven, Van den Bussche, 2011])

- $\mathcal{C}_{\text{queries}}(\text{Id} + \text{All}) \cap \text{Coordination-Free} = \mathcal{C}_{\text{queries}}(\text{Id})$

Theorem ([Ameloot, K, Neven, Zinn, 2014])

- $\mathcal{C}_q(\text{Id} + \text{All} + \text{Policy}) \cap \text{Coordination-Free} = \mathcal{C}_q(\text{Id} + \text{Policy})$
- $\mathcal{C}_q(\text{Id} + \text{All} + \text{AdomPolicy}) \cap \text{Coordination-Free} = \mathcal{C}_q(\text{Id} + \text{AdomPolicy})$

Coordination Free = Ability to compute without access to All.

Coordination / Computability

Definition

We call a behavior **coordinating-free** if it is in the set $\mathcal{C}(\text{Id} + \text{Ord})$

Definition

We call a behavior **weakly coordinating** if it is in a set $\mathcal{C}(\text{Id} + \text{Something})$ with property that there is no termination-aware algorithm to compute **All**.

$$\mathcal{C}(\text{Id} + \text{All} + \text{Order}) = \text{All behaviors}$$

Monotone Behaviors?

Theorem

$\mathcal{C}(\text{Id} + \text{Ord} + \chi) = \text{all } \chi\text{-monotone behaviors.}$

Details in [Baccaert, K, 2023]

Monotone Behaviors?

Example:

- Consider the query that takes a graph as input and asks if the graph as an isolated edge.
- Consider $\chi = \text{Id} + \text{Policy} + \text{Min} + \text{Max} + \text{Ord}$ as configuration constraints

Monotone Behaviors?

Example:

- Consider the query that takes a graph as input and asks if the graph as an isolated edge.
- Consider $\chi = \text{Id} + \text{Policy} + \text{Min} + \text{Max} + \text{Ord}$ as configuration constraints

$E(1, 2)$

$E(4, 5)$

Output: **True**

Monotone Behaviors?

Example:

- Consider the query that takes a graph as input and asks if the graph as an isolated edge.
- Consider $\chi = \text{Id} + \text{Policy} + \text{Min} + \text{Max} + \text{Ord}$ as configuration constraints

$E(1, 2)$
 $\text{id}(1)$
 $\{(1, 2), (2, 1), (1, 3), (3, 1),$
 $(1, 4), (4, 1), (1, 5), (5, 1), \dots\}$
 $\mathbf{1 \leq 2 \leq 3 \leq 4 \leq 5}$

$E(4, 5)$
 $\text{id}(3)$
 $\{(2, 4), (4, 2), (2, 5), (5, 2)$
 $(4, 5), (5, 4), \dots\}$
 $\mathbf{1 \leq 2 \leq 3 \leq 4 \leq 5}$

Output: **True**

Monotone Behaviors?

Example:

- Consider the query that takes a graph as input and asks if the graph as an isolated edge.
- Consider $\chi = \text{Id} + \text{Policy} + \text{Min} + \text{Max} + \text{Ord}$ as configuration constraints

$E(1, 2)$	$E(2, 3)E(3, 4)$	$E(4, 5)$
$\text{id}(1)$	$\text{id}(2)$	$\text{id}(3)$
$\{(1, 2), (2, 1), (1, 3), (3, 1),$ $(1, 4), (4, 1), (1, 5), (5, 1), \dots\}$	$\{(2, 3), (3, 2), (3, 4), (4, 3)$ $(3, 5), (5, 3), \dots\}$	$\{(2, 4), (4, 2), (2, 5), (5, 2)$ $(4, 5), (5, 4), \dots\}$
$\mathbf{1 \leq 2 \leq 3 \leq 4 \leq 5}$	$\mathbf{1 \leq 2 \leq 3 \leq 4 \leq 5}$	$\mathbf{1 \leq 2 \leq 3 \leq 4 \leq 5}$

Output: **False**

Monotone Behaviors?

Example:

- Consider the query that takes a graph as input and asks if the graph as an isolated edge.
- Consider $\chi = \text{Id} + \text{Policy} + \text{Min} + \text{Max} + \text{Ord}$ as configuration constraints

$E(1, 2)$	$E(2, 3)E(3, 4)$	$E(4, 5)$
$\text{id}(1)$	$\text{id}(2)$	$\text{id}(3)$
$\{(1, 2), (2, 1), (1, 3), (3, 1),$ $(1, 4), (4, 1), (1, 5), (5, 1), \dots\}$	$\{(2, 3), (3, 2), (3, 4), (4, 3)$ $(3, 5), (5, 3), \dots\}$	$\{(2, 4), (4, 2), (2, 5), (5, 2)$ $(4, 5), (5, 4), \dots\}$
$\mathbf{1} \leq 2 \leq 3 \leq 4 \leq \mathbf{5}$	$\mathbf{1} \leq 2 \leq 3 \leq 4 \leq \mathbf{5}$	$\mathbf{1} \leq 2 \leq 3 \leq 4 \leq \mathbf{5}$

Output: **False**

\Rightarrow The isolated edge query is not χ -monotone \Rightarrow not in $\mathcal{C}(\chi)$

Roadmap

- Context
- From Monotone to Negation-Free Queries
- Non-rewritable Monotone Queries
- Coordination-Freedom and System Knowledge
- **Open Problems**

Open Questions

- Are there other techniques to eliminate negation in monotone Datalog programs?
- Is there an elegant semantic definition of coordination-freedom for behaviors that matches our definition based on the **A11** relation?
- The definitions of behavior and constraint are all semantic, are there suitable languages that can be used to define them in a more elegant way?

More information: bas.ketsman@vub.be