



Rethinking Logical Interfaces to Data

Michael Benedikt

includes joint work over the past years with Pierre Bourhis, Louis Jachiet,
and Efthymia Tsamoura and joint work with Ehud Hrushovski

Great Thoughts Fridays.... Thursday warm up

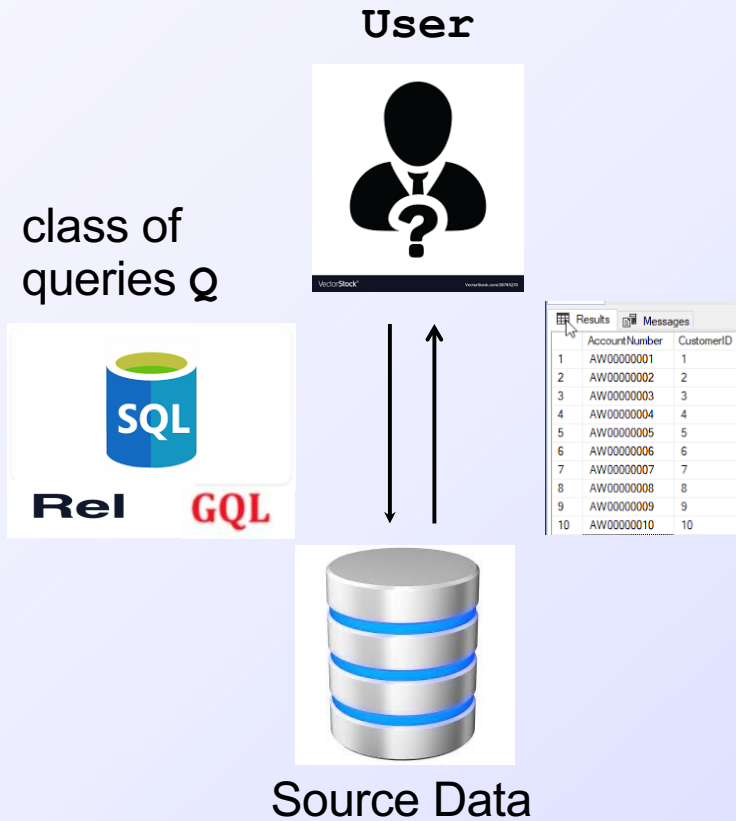
Richard Hamming



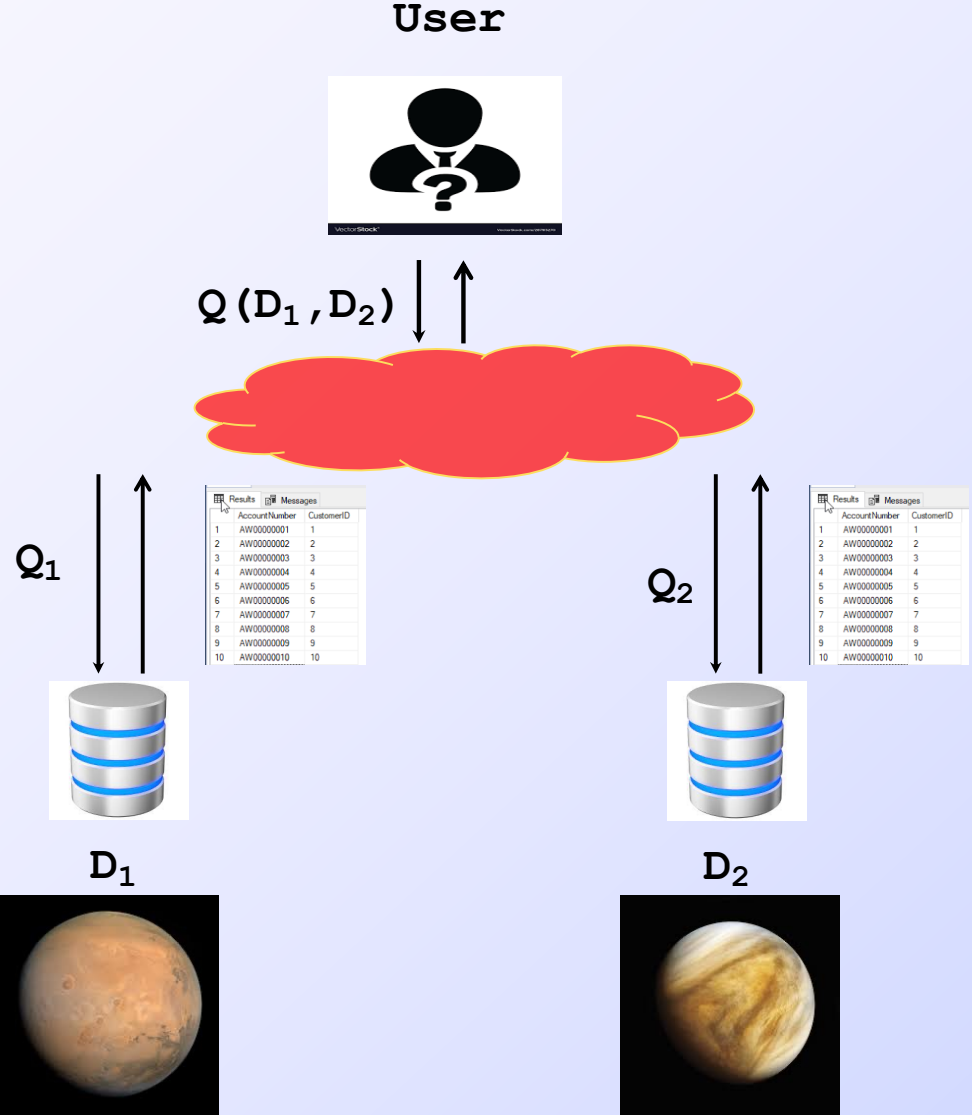
*I finally adopted what I called “Great Thoughts Time”.
When I went to lunch Friday noon, I would
only discuss great thoughts after that...*

To prepare for great thoughts tomorrow, I'll present some ideas on more flexible/general interfaces to data.

Interface to data: single source



Interface to data: distributed setting



Traditional views

Traditional views are an interface based on making available derived data.

In the distributed setting, for each source s , a view-based interface is a function F_s that takes as input a database instance for the schema of s and produces derived data; a **combination function** stitches derived data from different sources to partially/totally answer the user query.

We call the functions on the local sources ***distributed views (d-views)***.

Traditional views

Traditional views are an interface based on making available derived data.

In the distributed setting, for each source s , a view-based interface is a function \mathbb{F}_s that takes as input a database instance for the schema of s and produces derived data; a **combination function** stitches derived data from different sources to partially/totally answer the user query.

We call the functions on the local sources ***distributed views (d-views)***.

Views are often **definable** by logical formulas or a given query language.

Traditional views

Traditional views are an interface based on making available derived data.

In the distributed setting, for each source s , a view-based interface is a function F_s that takes as input a database instance for the schema of s and produces derived data; a **combination function** stitches derived data from different sources to partially/totally answer the user query.

We call the functions on the local sources ***distributed views (d-views)***.

Views are often **definable** by logical formulas or a given query language.
E.g. a **conjunctive query view** over source s :

$$\{ \mathbf{x}_1 \dots \mathbf{x}_m \mid \exists \mathbf{y}_1 \dots \mathbf{y}_m \mathbf{A}_1(\mathbf{x}_1 \dots \mathbf{y}_1 \dots) \wedge \dots \wedge \mathbf{A}_m(\dots) \}$$

where \mathbf{A}_i are atoms over the relations in the local source s .

Prior interfaces beyond views

Other ways to provide a restricted interface to centralized or integrated data:

- **Access patterns:** allow access to source data, but require certain values to be specified [Chen Li and Edward Chang; Alan Nash et al.; Deutsch, Nash, Ludascher 2007]
- **Views with access patterns** [Deutsch, Nash, Ludascher 2007; Romero, Preda, Amarilli, Suchanek 2020]
- **Data Exchange/Virtual Data Integration** [Halevy; Lenzerini; Fagin, Kolaitis, Miller, Popa 2005]
- **Specification of allowed queries via automata** [Cautis, Deutsch, Onose, TOCS 2010]

Prior interfaces beyond views

Other ways to provide a restricted interface to centralized or integrated data:

- **Access patterns:** allow access to source data, but require certain values to be specified [Chen Li and Edward Chang; Alan Nash et al.; Deutsch, Nash, Ludascher 2007]
 - **Views with access patterns** [Deutsch, Nash, Ludascher 2007; Romero, Preda, Amarilli, Suchanek 2020]
 - **Data Exchange/Virtual Data Integration** [Halevy; Lenzerini; Fagin, Kolaitis, Miller, Popa 2005]
 - **Specification of allowed queries via automata** [Cautis, Deutsch, Onose, TOCS 2010]
- **Minimal information to support a target query** [B., Bourhis, Jachiet, Tsamoura KR 2020/TODS 2022]
 - **Generalizing views via indistinguishability** [B. & Hrushovski 2023]

Minimally informative query answering

We specify a set of queries Q_1, \dots, Q_k (“utility queries”) that we want to support, and ask for the **minimally informative views** (within a class) that **support these queries**.

Example



Dagstuhl and the Simons Institute want to support access to their independent datastores

Simons

`SimonsParticipant(name, program, year)`

Dagstuhl

`DagstuhlParticipant(name, program, year)`

They want the interface to support answering some queries that span sources, like asking if there are researchers attending programs at both venues the same year.

$Q = \exists \text{program}_1 \exists \text{program}_2 \exists \text{name} \exists \text{year} \text{SimonsParticipant}(\text{name}, \text{program}_1, \text{year}) \wedge \text{DagstuhlParticipant}(\text{name}, \text{program}_2, \text{year})$

Example



Dagstuhl and the Simons Institute want to support access to their independent datastores

Simons

`SimonsParticipant(name, program, year)`

Dagstuhl

`DagstuhlParticipant(name, program, year)`


They want the interface to support answering some queries that span sources, like asking if there are researchers attending programs at both venues the same year.

$Q = \exists \text{program}_1 \exists \text{program}_2 \exists \text{name} \exists \text{year} \text{SimonsParticipant}(\text{name}, \text{program}_1, \text{year}) \wedge \text{DagstuhlParticipant}(\text{name}, \text{program}_2, \text{year})$

The sources should support this query, and give out the **minimal information** among \mathcal{d} -views supporting this join query Q .

How do we formalize the notion of support and minimal information?

Example



Dagstuhl and the Simons Institute want to support access to their independent datastores

Simons

`SimonsParticipant(name, program, year)`

Dagstuhl


`DagstuhlParticipant(name, program, year)`

They want the interface to support answering some queries that span sources, like asking if there are researchers attending programs at both venues the same year.

$Q = \exists \text{program}_1 \exists \text{program}_2 \exists \text{name} \exists \text{year} \text{SimonsParticipant}(\text{name}, \text{program}_1, \text{year}) \wedge \text{DagstuhlParticipant}(\text{name}, \text{program}_2, \text{year})$

The sources should support this query, and give out the **minimal information** among d -views supporting this join query Q .


How do we formalize the notion of support and minimal information?



Intuitively Understanding
The Shannon Entropy

$$H(X) = - \sum_T P(x_i) \log P(x_i)$$

Example



Dagstuhl and the Simons Institute want to support access to their independent datastores

Simons

`SimonsParticipant(name, program, year)`

Dagstuhl

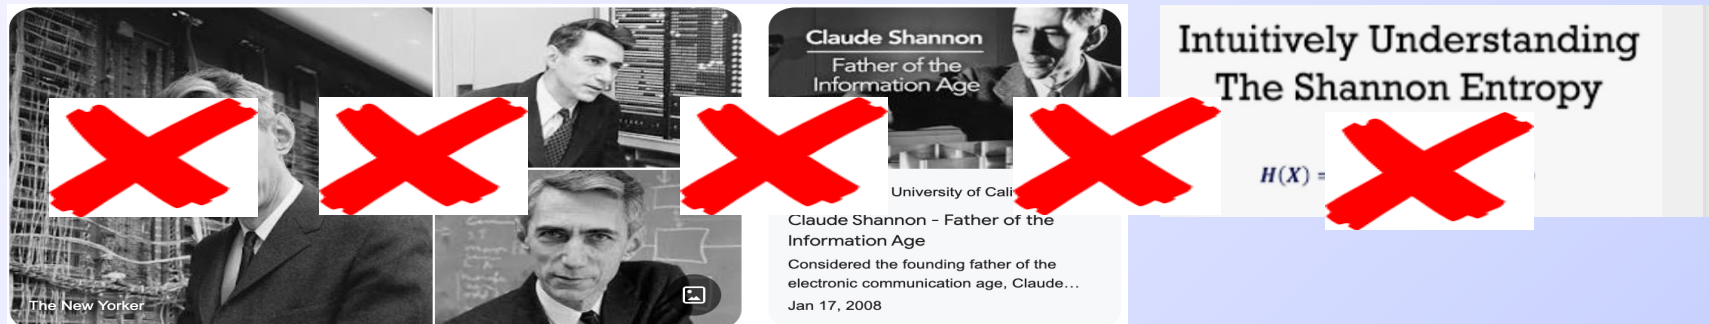
`DagstuhlParticipant(name, program, year)`

They want the interface to support answering some queries that span sources, like asking if there are researchers attending programs at both venues the same year.

$Q = \exists \text{program}_1 \exists \text{program}_2 \exists \text{name} \exists \text{year} \text{SimonsParticipant}(\text{name}, \text{program}_1, \text{year}) \wedge \text{DagstuhlParticipant}(\text{name}, \text{program}_2, \text{year})$

The sources should support this query, and give out the **minimal information** among \mathcal{d} -views supporting this join query Q .

How do we formalize the notion of support and minimal information?



Formalization: supporting a query



What formalize the notion that the views support queries $Q_1 \dots Q_u$ using Segoufin and Vianu's notion of determinacy.

Formalization: supporting a query



What formalize the notion that the views support queries $Q_1 \dots Q_u$ using Segoufin and Vianu's notion of determinacy.

Given a query Q and views $v_1 \dots v_t$ we say Q is **determined** by $v_1 \dots v_t$ if:

for all input D, D' with $v_1(D) = v_1(D'), \dots, v_t(D) = v_t(D')$ we have $Q(D) = Q(D')$

We say that d -view $v_1 \dots v_t$ **supports** Q if Q is determined by $v_1 \dots v_t$.

Read as " $v_1 \dots v_t$ contains all the information needed for Q "

Formalization: minimal information

We formalize the notion that the views are minimally informative using

Formalization: minimal information

We formalize the notion that the views are minimally informative using
Segoufin and Vianu's notion of determinacy.

Formalization: minimal information

We formalize the notion that the views are minimally informative using
Segoufin and Vianu's notion of determinacy.

We say a \mathcal{d} -view \mathcal{V} is a ***minimally informative supportive \mathcal{d} -view*** for query Q **within a class of queries \mathcal{C}** if:

- \mathcal{V} supports Q
- \mathcal{V} is based on queries in \mathcal{C} and for every other \mathcal{d} -view \mathcal{V}' using queries from \mathcal{C} that supports Q , we have \mathcal{V}' determines each view in \mathcal{V}

Example



Dagstuhl and Simons want to support access to their independent datastores

Simons

`SimonsParticipant(name, program, year)`

Dagstuhl

`DagstuhlParticipant(name, program, year)`

They want the interface to support answering some queries that span sources, like asking if there are researchers who attended programs at both venues in the same year.

$Q = \exists \text{program}_1 \exists \text{program}_2 \exists \text{name} \exists \text{year} \text{SimonsParticipant}(\text{name}, \text{program}_1, \text{year}) \wedge \text{DagstuhlParticipant}(\text{name}, \text{program}_2, \text{year})$

Example



Dagstuhl and Simons want to support access to their independent datastores

Simons

`SimonsParticipant(name, program, year)`

Dagstuhl

`DagstuhlParticipant(name, program, year)`

They want the interface to support answering some queries that span sources, like asking if there are researchers who attended programs at both venues in the same year.

$Q = \exists \text{program}_1 \exists \text{program}_2 \exists \text{name} \exists \text{year} \text{SimonsParticipant}(\text{name}, \text{program}_1, \text{year}) \wedge \text{DagstuhlParticipant}(\text{name}, \text{program}_2, \text{year})$

The minimal information \mathbf{d} -views that support this query are the **obvious** ones:

Simons should publish the view:

$\exists \text{program} \text{SimonsParticipant}(\text{name}, \text{program}, \text{year})$

While Dagstuhl should publish the view:

$\exists \text{program} \text{DagstuhlParticipant}(\text{name}, \text{program}, \text{year})$

Example of our results

Theorem [B., Bourhis, Jachiet, Tsamoura] For **any** utility queries, minimally informative \mathfrak{d} -views exist, and for **CQ** utility queries they are expressible as traditional views in relational algebra. The same holds in the presence of integrity constraints on each local source.

Example of our results

Theorem [B., Bourhis, Jachiet, Tsamoura] For **any** utility queries, minimally informative \mathfrak{d} -views exist, and for **CQ** utility queries they are expressible as traditional views in relational algebra. The same holds in the presence of integrity constraints on each local source.

However, there are CQ utility queries, where the minimally informative \mathfrak{d} -views are not CQs (and in particular, are not the obvious ones).

Example of our results

Theorem [B., Bourhis, Jachiet, Tsamoura] For **any** utility queries, minimally informative \mathfrak{d} -views exist, and for **CQ** utility queries they are expressible as traditional views in relational algebra. The same holds in the presence of integrity constraints on each local source.

However, there are CQ utility queries, where the minimally informative \mathfrak{d} -views are not CQs (and in particular, are not the obvious ones).

Theorem [B., Bourhis, Jachiet, Tsamoura] For any CQ utility queries, **minimally informative CQ** \mathfrak{d} -views exist.

The same holds in the presence of integrity constraints on each local source.

Using logic-based information theory

These tools allow us to analyze trade-offs in view design.
Questions of the form “are there distributed views that support query Q but which do not reveal any information about query P ”


Using logic-based information theory

These tools allow us to analyze trade-offs in view design.
Questions of the form “are there distributed views that support query Q but which do not reveal any information about query p ”

R



S, T



$$Q = \exists \mathbf{x} \exists \mathbf{y} R(\mathbf{x}, \mathbf{y}) \wedge S(\mathbf{x}, \mathbf{y}) \wedge T(\mathbf{x}, \mathbf{y})$$

Clearly, we can design views at each source to answer Q :
each source just exports its data.

But suppose we want to keep the following query private:

$$p = \exists \mathbf{x} R(\mathbf{x}, \mathbf{x})$$

Using logic-based information theory

These tools allow us to analyze trade-offs in view design.
Questions of the form “are there distributed views that support query Q but which do not reveal any information about query p ”

R



S, T



$$Q = \exists x \exists y R(x, y) \wedge S(x, y) \wedge T(x, y)$$

Clearly, we can design views at each source to answer Q :
each source just exports its data.

But suppose we want to keep the following query private:

$$p = \exists x R(x, x)$$

Intuitively, any views (no matter what query language) that allow Q to be answered must disclose p on some instance.

Using the prior theorem, we can prove this.

Using logic-based information theory

These tools allow us to answer questions of the form “are there distributed views that support query Q but which do not reveal any information about query P ”

R, S



S, T



There is a partial synchronization between Simons and Dagstuhl:
s is replicated between the two sources.

Using logic-based information theory

These tools allow us to answer questions of the form “are there distributed views that support query Q but which do not reveal any information about query p ”

R, S



S, T



There is a partial synchronization between Simons and Dagstuhl:
 s is replicated between the two sources.

$$Q = \exists x \exists y R(x, y) \wedge S(x, y) \wedge T(x, y)$$

Clearly, we can design views at each source to answer Q :
each source just exports its data.

But suppose we want to keep the following query private:

$$p = \exists x R(x, x)$$

Using logic-based information theory

These tools allow us to answer questions of the form “are there distributed views that support query Q but which do not reveal any information about query p ”

R, S



S, T



There is a partial synchronization between Simons and Dagstuhl:
 s is replicated between the two sources.

$$Q = \exists x \exists y R(x, y) \wedge S(x, y) \wedge T(x, y)$$

Clearly, we can design views at each source to answer Q :
each source just exports its data.

But suppose we want to keep the following query private:


$$p = \exists x R(x, x)$$

It **is** possible to support Q without revealing p .

But we will need an interface mechanism beyond relational algebra views.


Tentative moral on minimal information querying

- Compare the expressiveness of different interface mechanisms.
- Develop the notion of determinacy from Segoufin and Vianu as a metric to perform this comparison.



Tentative moral on minimal information querying

- Compare the expressiveness of different interface mechanisms.
- Develop the notion of determinacy from Segoufin and Vianu as a metric to perform this comparison.



Also used in query pricing [Koutris, Upadhyaya, Howe, Balazinska, Suciu JACM 2015] and in other work on information disclosure [B., Bourhis, ten Cate, Puppis, Vanden Boom TOCL 2021; B., Cuenca Grau, Kostylev JAIR 2018]


Interfaces beyond views

Other ways to provide restricted access to centralized or integrated data:

- **Access patterns:** allow access to source data, but require certain values to be specified [Chen Li and Edward Chang; Alan Nash et al.' Deutsch, Nash, Ludascher 2007]
 - **Views with access patterns** [Deutsch, Nash, Ludascher 2007; Romero, Preda, Amarilli, Suchanek 2020]
 - **Data Exchange/Virtual Data Integration** [Halevy; Lenzerini; Fagin, Kolaitis, Miller, Popa 2005]
 - **Specification of allowed queries via automata** [Cautis, Deutsch, Onose, TOCS 2010]
- **Minimal information views that support a query** [B., Bourhis, Jachiet, Tsamoura KR 2020/TODS 2022]
 - **Generalizing views via indistinguishability** [B. & Hrushovski 2023]


Generalizing views via database indistinguishability

An **indistinguishability relation** is an equivalence relation on databases. This defines an interface.



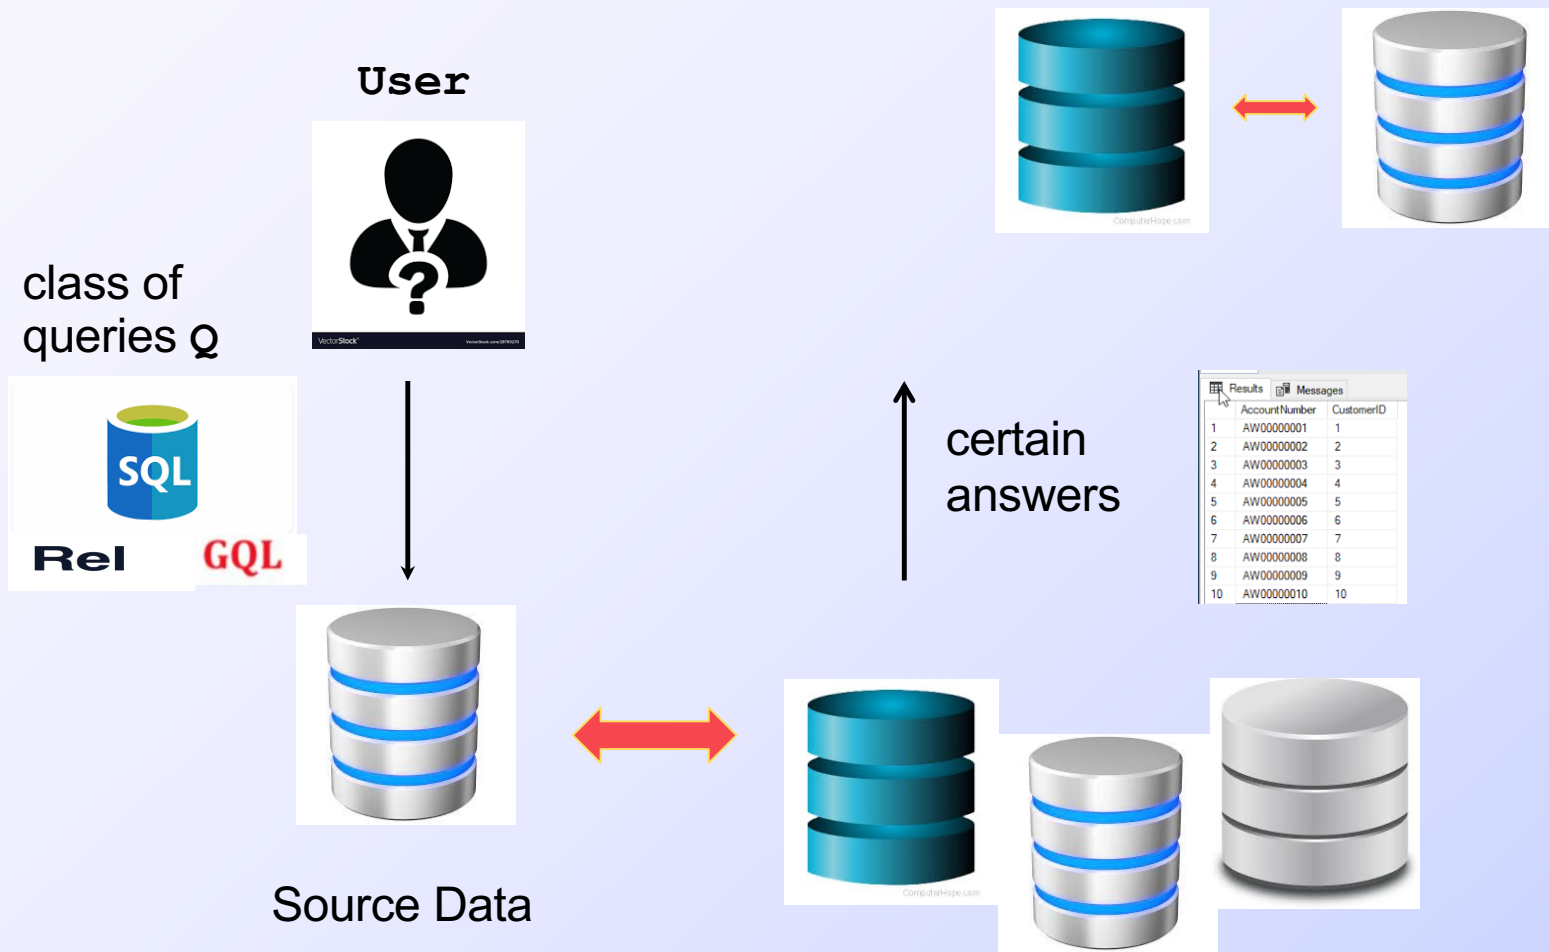
Generalizing views via database indistinguishability

An **indistinguishability relation** is an equivalence relation on databases. This defines an interface.



Generalizing views via database indistinguishability

An **indistinguishability relation** is an equivalence relation on databases. This defines an interface.



Defining indistinguishability with logic

An **indistinguishability relation** is an equivalence relation on databases. It can be thought of as an “abstract view”: we are exporting the equivalence class of a database.

Defining indistinguishability with logic

An **indistinguishability relation** is an equivalence relation on databases. It can be thought of as an “abstract view”: we are exporting the equivalence class of a database.

Example:

Declare graph databases G and G' indistinguishable if they have the same triangles:

$\forall \mathbf{x}_1 \mathbf{x}_2 \mathbf{x}_3$

$$[(G(\mathbf{x}_1, \mathbf{x}_2) \wedge G(\mathbf{x}_2, \mathbf{x}_3) \wedge G(\mathbf{x}_3, \mathbf{x}_1)) \leftrightarrow (G'(\mathbf{x}_1, \mathbf{x}_2) \wedge G'(\mathbf{x}_2, \mathbf{x}_3) \wedge G'(\mathbf{x}_3, \mathbf{x}_1))]$$

Defining indistinguishability with logic

An **indistinguishability relation** is an equivalence relation on databases. It can be thought of as an “abstract view”: we are exporting the equivalence class of a database.

Example:

Declare graph databases G and G' indistinguishable if they have the same triangles:

$\forall \mathbf{x}_1 \mathbf{x}_2 \mathbf{x}_3$

$$[(G(\mathbf{x}_1, \mathbf{x}_2) \wedge G(\mathbf{x}_2, \mathbf{x}_3) \wedge G(\mathbf{x}_1, \mathbf{x}_3)) \leftrightarrow (G'(\mathbf{x}_1, \mathbf{x}_2) \wedge G'(\mathbf{x}_2, \mathbf{x}_3) \wedge G'(\mathbf{x}_1, \mathbf{x}_3))]$$

A **first order definable indistinguishability relation** is given by φ a first order sentence in the language of **two copies of the schema**.

Thus φ defines a collection of pairs of databases. **If** φ defines an equivalence relation, then φ provides an indistinguishability relation.

Note: a typical first order φ will **not** define an equivalence relation on databases. For example, transitivity will fail.

Defining indistinguishability with logic

An **indistinguishability relation** is an equivalence relation on databases.

A **first order definable indistinguishability relation** is given by φ a first order sentence in the language of **two copies of the schema** which happens to define an equivalence relation.

Class of examples of FO indistinguishability:

Traditional relational algebra views $v_1 \dots v_k$ give a first order indistinguishability relation:

$$\forall \mathbf{x}_1 \dots \mathbf{x}_j \ [\bigwedge_{i \leq k} v_i(\mathbf{x}_1 \dots \mathbf{x}_j) \leftrightarrow v'_i(\mathbf{x}_1 \dots \mathbf{x}_j)]$$

where v'_i is a copy of v_i on the primed signature.

Recall: building interfaces beyond traditional views

Theorem [B., Bourhis, Jachiet, Tsamoura] For any utility queries, minimally informative \mathcal{d} -views exist **as an indistinguishability relation.**

For CQ utility queries they are expressible as traditional views in relational algebra. The same holds in the presence of integrity constraints on each local source.

Recall: building interfaces beyond views

These tools allow us to analyze questions of the form “are there distributed views that support query Q but which do not reveal any information about query p ”

R, S



S, T



There is a partial synchronization between Simons and Dagstuhl:
 s is replicated between the two sources.

$$Q = \exists x y R(x, y) \wedge S(x, y) \wedge T(x, y)$$

Clearly, we can design views at each source to answer Q :
each source just exports its data.

But suppose we want to keep the following query private:

$$p = \exists x R(x, x)$$

It **is** possible to support Q without revealing p .

But we will need an interface mechanism beyond relational algebra views

Recall: building interfaces beyond views

These tools allow us to analyze questions of the form “are there distributed views that support query Q but which do not reveal any information about query p ”

R, S



S, T



There is a partial synchronization between Simons and Dagstuhl:
 s is replicated between the two sources.

$$Q = \exists x y R(x, y) \wedge S(x, y) \wedge T(x, y)$$

Clearly, we can design views at each source to answer Q :
each source just exports its data.

But suppose we want to keep the following query private:

$$p = \exists x R(x, x)$$

It **is** possible to support Q without revealing p .

But we will need an interface mechanism beyond relational algebra views

- namely, an indistinguishability relation.

Super-generalizing views via database indistinguishability

An **indistinguishability relation** is an equivalence relation on databases. It can be thought of as an “abstract view”: we are exporting the equivalence class of a database.

This is a super-general notion.

In current work with Hrushovski we study it primarily in the setting of **classical model theory**: indistinguishability relations over **infinite structures**, focusing on relations definable in **first order and infinitary logic**.

Motivated by classification theory, descriptive set theory, model theory for topology and analysis.

But there are some results for first order indistinguishability relations on databases/finite models.

Defining indistinguishability with logic

A **first order definable indistinguishability relation** is given by φ a first order sentence in the language of **two copies of the schema** which happens to define an equivalence relation.

Recall:

Traditional relational algebra views $v_1 \dots v_k$ give a first order indistinguishability relation:

$$\forall \mathbf{x}_1 \dots \mathbf{x}_j [\bigwedge_{i \leq k} v_i(\mathbf{x}_1 \dots \mathbf{x}_j) \leftrightarrow v'_i(\mathbf{x}_1 \dots \mathbf{x}_j)]$$

Indistinguishability versus query-based views

A **first order definable indistinguishability relation** is given by φ a first order sentence in the language of **two copies of the schema** which happens to define an equivalence relation.

Traditional **nested relational calculus views** $V_1 \dots V_k$ give a first order indistinguishability relation.

Example: Given binary $R(x,y)$, consider the view corresponding to the nested query $\{ \{ y \mid (x,y) \in R \} \mid x \in \pi_1(R) \}$
That is, R and R' are indistinguishable if they have the same adjacency sets of nodes.

$$\forall x \exists x' [\forall y R(x,y) \leftrightarrow R'(x',y)] \wedge \\ \forall x' \exists x [\forall y R(x,y) \leftrightarrow R'(x',y)]$$

Indistinguishability versus query-based views

Let \mathbf{E} be an “indistinguishability relation”: an equivalence relation on databases. \mathbf{E} can be thought of as an “abstract view”.

A **first order definable indistinguishability relation** is given by φ a first order sentence in the language of **two copies of the schema**. Thus φ defines a collection of pairs of databases, and we require φ to define an equivalence relation.

Traditional **nested relational calculus views** $\mathbf{V}_1 \dots \mathbf{V}_K$ give a first order indistinguishability relation.

Example: Given ternary $\mathbf{R}(\mathbf{x}, \mathbf{y}, \mathbf{z})$, consider the view corresponding to the nested query

$$\{ \{ \mathbf{z} \mid (\mathbf{x}, \mathbf{y}, \mathbf{z}) \in \mathbf{R} \} \mid \mathbf{y} \in \pi_2(\mathbf{R}) \} \mid \mathbf{x} \in \pi_1(\mathbf{R}) \}$$

\\adjacency set of x, y
\\set of adjacency sets for x
\\set of sets of adjacency sets

That is, \mathbf{R} and \mathbf{R}' are indistinguishable if they have the same sets of sets of adjacency sets of pairs.

Separation

Hierarchy Theorem [B., Hrushovski]

For every n , there are depth n nested relational views whose indistinguishability relation is not given by depth $n-1$ nested relational views.

Collapse of Nested Relational Views to Relational Views

Sparse Collapse Theorem [B., Hrushovski]

Suppose \mathfrak{E} is given by a set of nested relational views on a graph database.

Suppose \mathcal{C} is a collection of graphs that exclude a minor.

Then over \mathcal{C} , \mathfrak{E} is given by a set of relational algebra views.

Prefix Classes of FO Indistinguishability Relations

Can classify FO Indistinguishability relations by the quantifier alternation, focusing only on the quantified variables that vary over both models.

E.g. Triangle-based equivalence example is $\Pi_1 : \forall \mathbf{x}_1 \mathbf{x}_2 \mathbf{x}_3 \dots$

The first nested relational calculus example (adjacency sets) is $\Pi_3 :$

$\forall \mathbf{x}_1 \dots \exists \mathbf{y}_1 \dots \forall \mathbf{z}_1 \dots$

Prefix Classes of FO Indistinguishability Relations

Can classify FO indistinguishability relations by the quantifier alternation, focusing only on the quantified variables that vary over both models.

E.g. Triangle-based equivalence example is $\Pi_1 : \forall \mathbf{x}_1 \mathbf{x}_2 \mathbf{x}_3 \dots$

The first nested relational calculus example (adjacency sets) is $\Pi_3 :$

$\forall \mathbf{x}_1 \dots \exists \mathbf{y}_1 \dots \forall \mathbf{z}_1 \dots$

Π_2 Theorem [B., Hrushovski]

Suppose \mathbf{E} is a Π_2 indistinguishability relation: given by a $\forall \mathbf{x}_1 \dots \exists \mathbf{y}_1 \dots$ sentence in two copies of the signature, that happens to define an equivalence relation, showing here only the quantifiers of variables that span both models.

Then \mathbf{E} is a Π_1 indistinguishability relation: given by a universal sentence.

First Order Indistinguishability and Nested Relations

Question: Is every first order indistinguishability relation is given by nested relational calculus views?

Tentative Moral on Indistinguishability Relations

- Indistinguishability relations make the world of traditional view interfaces look very small
- Issue of converting between interface specifications of different natures. In this case, from a compactly-represented equivalence class to a canonical representative.

Many analogies in descriptive complexity theory and descriptive set theory.

Lead-In To Great Thoughts Friday

1970



1980



1990



2000



2010



2020

Relational databases have been around for over 50 years.

And in the first 40 years, the notion of logical interface today and notions of comparing interfaces were frequently revisited, often radically so.

Lead-In To Great Thoughts Friday

1970



A Relational Model of Data for Large Shared Data Banks

E. F. CODD
IBM Research Laboratory, San Jose, California

1980

The format model: A theory of database organization

Richard Hull, Chee K. Yap

1990



The Information Manifold

Thomas Kirk
AT&T Bell Laboratories
tk@research.att.com

Alon Y. Levy
AT&T Bell Laboratories
levy@research.att.com

Yehoshua Sagiv
Hebrew University
sagiv@cs.huji.ac.il

Divesh Srivastava
AT&T Bell Laboratories
divesh@research.att.com

2000



Data exchange: semantics and query answering ☆

Ronald Fagin^a ✉, Phokion G. Kolaitis^{b 1} ✉, Renée J. Miller^{c 1} ✉, Lucian Popa^a 👤 ✉

2010



Sharing Distributed Knowledge on the Web

Serge Abiteboul

Collège de France, INRIA Saclay & ENS Cachan, France

2020



Lead-In To Great Thoughts Friday

1970



A Relational Model of Data for Large Shared Data Banks

E. F. Codd
IBM Research Laboratory, San Jose, California

1980

The format model: A theory of database organization

Richard Hull, Chee K. Yap

1990



The Information Manifold

Thomas Kirk
AT&T Bell Laboratories
tk@research.att.com

Alon Y. Levy
AT&T Bell Laboratories
levy@research.att.com

Yehoshua Sagiv
Hebrew University
sagiv@cs.huji.ac.il

Divesh Srivastava
AT&T Bell Laboratories
divesh@research.att.com

2000



Data exchange: semantics and query answering ☆

Ronald Fagin^a ✉, Phokion G. Kolaitis^{b 1} ✉, Renée J. Miller^{c 1} ✉, Lucian Popa^a 👤 ✉

2010



Sharing Distributed Knowledge on the Web

Serge Abiteboul

Collège de France, INRIA Saclay & ENS Cachan, France

2020





When you need more complicated views

R



S



Query to support specified as

$$Q = \exists \mathbf{x} \mathbf{y} \ R(\mathbf{x}, \mathbf{y}) \wedge S(\mathbf{x}, \mathbf{y}) \wedge S(\mathbf{y}, \mathbf{x})$$

When you need more complicated views

R



S



Query to support specified as

$$Q = \exists \mathbf{x} \mathbf{y} \ R(\mathbf{x}, \mathbf{y}) \wedge S(\mathbf{x}, \mathbf{y}) \wedge S(\mathbf{y}, \mathbf{x})$$

Minimal information supporting view at the Simons source:

$$S(\mathbf{x}, \mathbf{y}) \vee S(\mathbf{y}, \mathbf{x})$$