Introduction to Incremental View Maintenance

Dan Olteanu (University of Zurich)

fdbresearch.github.io

Berkeley CS 294-248: Topics in Database Theory
https://berkeley-cs294-248.github.io/
September 12 & 14, 2023

Short Biography

Since 2020: Professor of Computer Science, University of Zurich

- Since 2017: Computer Scientist at RelationalAI
- 2007 2020: Professor of Computer Science, University of Oxford
- 2013 2014: Visiting professor at University of California, Berkeley

Taught CS 186 & worked for LogicBlox

Acknowledgments

Members of the DaST IVM team







Ahmet Kara

Milos Nikolic

Haozhe Zhang

Real-Time Analytics

- Datasets continuously evolve over time
 - E.g.: data streams from sensors, social networks, apps
- Real-time analytics over streaming data
 - Users want fresh up-to-date computation results, e.g., models















Setting & Objective of this Lecture

Incremental View Maintenance (IVM)

- Well-established and longstanding research problem
- Confusing naming: incremental vs decremental Alternative common naming: *Fully dynamic*

Setting & Objective of this Lecture

Incremental View Maintenance (IVM)

- Well-established and longstanding research problem
- Confusing naming: incremental vs decremental

Alternative common naming: Fully dynamic

Setting

- Fully dynamic algorithms (i.e., supports inserts and deletes)
- Single-tuple updates to relational databases
- Relational queries (non-recursive)

Setting & Objective of this Lecture

Incremental View Maintenance (IVM)

- Well-established and longstanding research problem
- Confusing naming: incremental vs decremental

Alternative common naming: Fully dynamic

Setting

- Fully dynamic algorithms (i.e., supports inserts and deletes)
- Single-tuple updates to relational databases
- Relational queries (non-recursive)

Objectives

- Main IVM approaches: First/higher order, adaptive
- Which queries can be maintained optimally?





Answer







Answer







update





Agenda

Part 1. Data model and query language

- Part 2. Main IVM techniques by example
 - First-order IVM using delta queries
 - Higher-order IVM using DBToaster and F-IVM
 - Adaptive IVM using IVM^e
- Part 3. Which queries can be maintained optimally?
 - Constant update time and enumeration delay
 - Beyond constant time

Part 1. Data Model & Query Language

Classical Representation of Relations and Updates

Insertions and deletions are represented as separate tables:



Classical Representation of Relations and Updates

Insertions and deletions are represented as separate tables:



Order of updates matters!

This affects the parallel and distributed execution of updates

Inserts and deletes are treated differently

They rely on different execution mechanisms

Uniform Representation of Relations and Updates

Relations and updates are represented as factors mapping tuples to multiplicities:

employee	age	\rightarrow	#		employee	age	\rightarrow	#		employee	age	\rightarrow	#
Elise Steve Joe	35 40 30	$\stackrel{\rightarrow}{\stackrel{\rightarrow}{\rightarrow}}$	2 1 2	+	Steve Steve Joe Mary	40 38 30 33	$\stackrel{\rightarrow}{\rightarrow}\stackrel{\rightarrow}{\rightarrow}\stackrel{\rightarrow}{\rightarrow}$	$^{-1}_{-2}$	=	Elise Steve Mary	35 38 33	$\stackrel{\rightarrow}{}$	2 1 2
fa	actor				u	pdat	e			updat	ed fa	acto	or

Uniform Representation of Relations and Updates

Relations and updates are represented as factors mapping tuples to multiplicities:

employee	age	\rightarrow	#		employee	age	\rightarrow	#		employee	age	\rightarrow	-
Elise Steve Joe	35 40 30	$\stackrel{\rightarrow}{}$	2 1 2	+	Steve Steve Joe Mary	40 38 30 33	$\stackrel{\rightarrow}{\rightarrow}\stackrel{\rightarrow}{\rightarrow}\stackrel{\rightarrow}{\rightarrow}$	$^{-1}_{-2}$	=	Elise Steve Mary	35 38 33	$\stackrel{\rightarrow}{\rightarrow}_{\rightarrow}$	212
fa	actor				u	pdat	e			updat	ed f	acto	or

- Order of updates does not matter: Addition is associative
- Database may be in inconsistent state: tuples with negative multiplicities
- Multiplicity:

negative = tuple delete; positive = tuple insert; 0 = tuple not in database

Uniform Representation of Relations and Updates

Relations and updates are represented as factors mapping tuples to multiplicities:

employee	age	\rightarrow	#		employee	age	\rightarrow	#		employee	age	\rightarrow	
Elise Steve Joe	35 40 30	$\stackrel{\rightarrow}{}$	2 1 2	+	Steve Steve Joe Mary	40 38 30 33	$\stackrel{\rightarrow}{\rightarrow}\stackrel{\rightarrow}{\rightarrow}\stackrel{\rightarrow}{\rightarrow}$	$^{-1}_{-2}$	=	Elise Steve Mary	35 38 33	$\stackrel{\rightarrow}{\rightarrow}_{\rightarrow}$	
fa	actor				u	pdat	e			updat	ed fa	acto	or

Order of updates does not matter: Addition is associative

- Database may be in inconsistent state: tuples with negative multiplicities
- Multiplicity:

negative = tuple delete; positive = tuple insert; 0 = tuple not in database

From \mathbb{Z} to arbitrary rings (r_1 and r_2 are elements from a ring):

А	В	\rightarrow	R(A, B)
a ₁ a ₂	$b_1 \\ b_1$	$\stackrel{\rightarrow}{\rightarrow}$	r ₁ r ₂

Rings

A ring $(D, +, \cdot, 0, 1)$ is a set D with two binary ops:

Additive commutativity Additive associativity Additive identity Additive inverse

Multiplicative associativity Multiplicative identity

Left and right distributivity

$$a + b = b + a$$

$$(a + b) + c = a + (b + c)$$

$$\mathbf{0} + a = a + \mathbf{0} = a$$

$$\exists -a \in \mathbf{D} : a + (-a) = (-a) + a = \mathbf{0}$$

$$(a \cdot b) \cdot c = a \cdot (b \cdot c)$$

$$a \cdot \mathbf{1} = \mathbf{1} \cdot a = a$$

$$a \cdot (b + c) = a \cdot b + a \cdot c \text{ and}$$

$$(a+b)\cdot c = a\cdot c + b\cdot c$$

The ring is commutative if multiplication is also commutative

- Examples: $\mathbb{Z}, \mathbb{Q}, \mathbb{R}, \mathbb{C}$
- More on semi-rings later in the course
- We use the ring $(\mathbb{Z}, +, \cdot, 0, 1)$



R(a,b)						
А	В	\rightarrow	#			
a_1 a_2	$b_1 \\ b_1$	$\stackrel{\rightarrow}{\rightarrow}$	r ₁ r ₂			

S(a,b)						
А	В	\rightarrow	#			
a 2	b_1	\rightarrow	s_1			
a 3	b_2	\rightarrow	s ₂			

T(b	T(b,c)						
В	С	\rightarrow	#				
$b_1 \\ b_2$	С ₁ С2	$\stackrel{\rightarrow}{\rightarrow}$	t_1 t_2				

 $\begin{array}{l} \textbf{Union/Sum} \ V(a,b) = \\ R(a,b) + S(a,b) \end{array}$

А	В	\rightarrow	#
а ₁ а ₂ а ₃	$b_1 \\ b_1 \\ b_2$	$\stackrel{\rightarrow}{\rightarrow}_{\rightarrow}$	$r_1 \\ r_2 + s_1 \\ s_2$

R(a,b)						
Α	В	\rightarrow	#			
а ₁ а ₂	$b_1 \\ b_1$	$\stackrel{\rightarrow}{\rightarrow}$	r ₁ r ₂			

S(a,b)						
А	В	\rightarrow	#			
a 2	b_1	\rightarrow	s 1			
a 3	b_2	\rightarrow	s 2			

T(b,c)					
В	С	\rightarrow	#		
b_1 b_2	С ₁ С2	$\stackrel{\rightarrow}{\rightarrow}$	t_1 t_2		

 $\begin{array}{l} \textbf{Union/Sum} \ V(a,b) = \\ R(a,b) + S(a,b) \end{array}$

А	В	\rightarrow	#
a1 a2 a3	$b_1 \\ b_1 \\ b_2$	$\stackrel{\rightarrow}{\rightarrow}_{\rightarrow}$	$r_1 + r_1 + s_1 \\ s_2$

Join/Product $W(a, b, c) = V(a, b) \cdot T(b, c)$

А	В	С	\rightarrow	#
a1 a2 a3	$b_1\ b_1\ b_2$	c ₁ c ₁ c ₂	$\stackrel{\rightarrow}{\rightarrow} \stackrel{\rightarrow}{\rightarrow}$	$(r_1 \cdot t_1 \\ (r_2 + s_1) \cdot t_1 \\ s_2 \cdot t_2$

R(a, b)	S(a,b)	T(b,c)
$ \begin{array}{c ccc} A & B & \rightarrow & \# \\ \hline a_1 & b_1 & \rightarrow & r_1 \\ a_2 & b_1 & \rightarrow & r_2 \end{array} $	$\begin{array}{cccc} A & B & \rightarrow & \# \\ \hline a_2 & b_1 & \rightarrow & s_1 \\ a_3 & b_2 & \rightarrow & s_2 \end{array}$	$\begin{array}{cccccccccccccccccccccccccccccccccccc$
$\begin{array}{l} \textbf{Union/Sum} \ V(a,b) = \\ R(a,b) + S(a,b) \end{array}$		$\label{eq:static} \begin{array}{l} \mbox{Join/Product} \\ W(a,b,c) = V(a,b) \cdot T(b,c) \end{array}$
$ \begin{array}{c ccc} A & B & \rightarrow & \# \\ \hline a_1 & b_1 & \rightarrow & r_1 \\ a_2 & b_1 & \rightarrow & r_2 + s_1 \\ a_3 & b_2 & \rightarrow & s_2 \end{array} $		$ \begin{array}{c ccccccccccccccccccccccccccccccccccc$

Project/Marginalization

 $U(b,c) = \sum_{a} W(a,b,c)$

 $\begin{array}{cccc} B & C & \rightarrow & \# \\ \hline b_1 & c_1 & \rightarrow & r_1 * t_1 + (r_2 + s_1) \cdot t_1 \\ b_2 & c_2 & \rightarrow & s_2 \cdot t_2 \end{array}$

Query Language Considered in This Lecture

Language of queries with operations sum, product, and variable marginalization over factors. Examples:

List the paths of length two in a graph with edge factor *E*:

$$Q(a,b,c) = E(a,b) \cdot E(b,c)$$
Language of queries with operations sum, product, and variable marginalization over factors. Examples:

List the paths of length two in a graph with edge factor *E*:

$$Q(a, b, c) = E(a, b) \cdot E(b, c)$$

List the source-target nodes of the 2-length paths (b ranges over the vertices in the graph):

$$Q(a,c) = \sum_{b} E(a,b) \cdot E(b,c)$$

Language of queries with operations sum, product, and variable marginalization over factors. Examples:

List the paths of length two in a graph with edge factor *E*:

$$Q(a, b, c) = E(a, b) \cdot E(b, c)$$

List the source-target nodes of the 2-length paths (b ranges over the vertices in the graph):

$$Q(a,c) = \sum_{b} E(a,b) \cdot E(b,c)$$

List the triangles:

$$Q(a, b, c) = E(a, b) \cdot E(b, c) \cdot E(c, a)$$

Language of queries with operations sum, product, and variable marginalization over factors. Examples:

List the paths of length two in a graph with edge factor *E*:

$$Q(a, b, c) = E(a, b) \cdot E(b, c)$$

List the source-target nodes of the 2-length paths (b ranges over the vertices in the graph):

$$Q(a,c) = \sum_{b} E(a,b) \cdot E(b,c)$$

List the triangles:

$$Q(a, b, c) = E(a, b) \cdot E(b, c) \cdot E(c, a)$$

• Count the triangles (*a*, *b*, *c* range over the set of vertices):

$$Q = \sum_{a,b,c} E(a,b) \cdot E(b,c) \cdot E(c,a)$$

Language of queries with operations sum, product, and variable marginalization over factors. Examples:

List the paths of length two in a graph with edge factor *E*:

$$Q(a, b, c) = E(a, b) \cdot E(b, c)$$

List the source-target nodes of the 2-length paths (b ranges over the vertices in the graph):

$$Q(a,c) = \sum_{b} E(a,b) \cdot E(b,c)$$

List the triangles:

$$Q(a, b, c) = E(a, b) \cdot E(b, c) \cdot E(c, a)$$

• Count the triangles (*a*, *b*, *c* range over the set of vertices):

$$Q = \sum_{a,b,c} E(a,b) \cdot E(b,c) \cdot E(c,a)$$

• Update a factor E with changes given by another factor δE :

$$E_{new}(a, b) = E(a, b) + \delta E(a, b)$$

Compute COUNT over the natural join of factors R, S, and T (using Yannakakis)

Factors map tuples to multiplicity 1



Join hypergraph

Compute COUNT over the natural join of factors R, S, and T (using Yannakakis)

Factors map tuples to multiplicity 1



Join hypergraph

Naïve: compute the join and then COUNT

$$\mathsf{Q} = \sum_{a,b,c,d,e} \mathsf{R}(\mathsf{a},\mathsf{b}) \cdot \mathsf{S}(\mathsf{a},\mathsf{c},\mathsf{e}) \cdot \mathsf{T}(\mathsf{c},\mathsf{d})$$

Takes $\mathcal{O}(N^3)$ time if each factor has size $\mathcal{O}(N)$!

There can be $\mathcal{O}(N^3)$ possible combinations of b, e, and d values

Compute COUNT over the natural join of factors R, S, and T (using Yannakakis)

Factors map tuples to multiplicity 1



Join hypergraph

Naïve: compute the join and then COUNT

$$\mathsf{Q} = \sum_{a,b,c,d,e} \mathsf{R}(\mathsf{a},\mathsf{b}) \cdot \mathsf{S}(\mathsf{a},\mathsf{c},\mathsf{e}) \cdot \mathsf{T}(\mathsf{c},\mathsf{d})$$

Takes $\mathcal{O}(N^3)$ time if each factor has size $\mathcal{O}(N)$!

There can be $\mathcal{O}(N^3)$ possible combinations of b, e, and d values

Can we compute COUNT in $\mathcal{O}(N)$ time?

Idea: Push COUNT past joins and marginalize variables



Idea: Push COUNT past joins and marginalize variables



 $\mathcal{O}(N)$

Idea: Push COUNT past joins and marginalize variables



Idea: Push COUNT past joins and marginalize variables

 $\begin{array}{cccc} V_1 & A & & & V_1 \\ \Rightarrow & & & \\ & & & & \\ & & & \\ & & & \\ & & & \\ & & & & \\ & & & \\ & & & \\ & & & \\ & & & \\ & & & \\ & & & \\ & & & \\ &$ R $V_1(a) = \sum_{i} R(a, b)$ $\mathcal{O}(N)$ $\mathsf{V}_2(\mathsf{a},\mathsf{c}) = \sum_\mathsf{e} \mathsf{S}(\mathsf{a},\mathsf{c},\mathsf{e})$ $\mathcal{O}(N)$ $V_3(c) = \sum_d T(c,d)$ $\mathcal{O}(N)$

Idea: Push COUNT past joins and marginalize variables

 $\begin{array}{cccc} V_1 & A & & V_1 & V_2 & V_1 & A \\ \Rightarrow & & & \\ & & & & \\ & & &$ R $V_1(a) = \sum_{i} R(a, b)$ $\mathcal{O}(N)$ $V_2(a,c) = \sum_a S(a,c,e)$ $\mathcal{O}(N)$ $V_3(c) = \sum_d T(c,d)$ $\mathcal{O}(N)$ $\mathsf{V}_4(\mathsf{a}) = \sum_c \left(\mathsf{V}_2(\mathsf{a},\mathsf{c})\cdot\mathsf{V}_3(\mathsf{c})\right)$ $\mathcal{O}(N)$

Idea: Push COUNT past joins and marginalize variables

 $\Rightarrow \begin{array}{c} V_{1} \\ \Rightarrow \\ T \end{array} \begin{array}{c} V_{1} \\ \Rightarrow \\ T \end{array} \begin{array}{c} V_{2} \\ \Rightarrow \\ V_{3} \\ \end{array} \begin{array}{c} V_{2} \\ \Rightarrow \\ V_{3} \\ \end{array} \begin{array}{c} V_{2} \\ \Rightarrow \\ V_{3} \\ \end{array} \begin{array}{c} V_{2} \\ \Rightarrow \\ V_{4} \end{array} \begin{array}{c} \Rightarrow \\ V_{4} \\ \end{array} \begin{array}{c} V_{2} \\ \Rightarrow \\ V_{4} \end{array} \begin{array}{c} V_{1} \\ \Rightarrow \\ V_{4} \end{array} \begin{array}{c} V_{2} \\ \Rightarrow \\ V_{4} \end{array} \begin{array}{c} V_{1} \\ \Rightarrow \\ V_{4} \end{array} \begin{array}{c} V_{2} \\ \Rightarrow \\ V_{4} \end{array} \begin{array}{c} V_{1} \\ \Rightarrow \\ V_{4} \end{array} \begin{array}{c} V_{2} \\ \Rightarrow \\ V_{4} \end{array} \begin{array}{c} V_{1} \\ \Rightarrow \\ V_{4} \end{array} \begin{array}{c} V_{2} \\ \end{array} \begin{array}{c} V_{2} \\ \Rightarrow \\ V_{4} \end{array} \begin{array}{c} V_{2} \\ \end{array} \end{array}$ R $V_1(a) = \sum_{i} R(a, b)$ $\mathcal{O}(N)$ $V_2(a,c) = \sum_e S(a,c,e)$ $\mathcal{O}(N)$ $V_3(c) = \sum_d T(c,d)$ $\mathcal{O}(N)$ $\mathsf{V_4}(\mathsf{a}) = \sum_c \left(\mathsf{V_2}(\mathsf{a},\mathsf{c})\cdot\mathsf{V_3}(\mathsf{c})\right)$ $\mathcal{O}(N)$ $\mathsf{Q} = \sum_{\mathsf{a}} \left(\mathsf{V}_1(\mathsf{a}) \cdot \mathsf{V}_4(\mathsf{a})\right)$ $\mathcal{O}(N)$

Factor Payloads Enable Different Aggregates (1/2)

Task: Compute $SUM(C \cdot D)$ over the join of R, S, and T

Factor Payloads Enable Different Aggregates (1/2)

Task: Compute $SUM(C \cdot D)$ over the join of R, S, and T

■ Use the same query:

$$\mathsf{Q} = \sum_{a,b,c,d,e} \mathsf{R}(\mathsf{a},\mathsf{b}) \cdot \mathsf{S}(\mathsf{a},\mathsf{c},\mathsf{e}) \cdot \mathsf{T}(\mathsf{c},\mathsf{d})$$

Factor Payloads Enable Different Aggregates (1/2)

Task: Compute $SUM(C \cdot D)$ over the join of R, S, and T

■ Use the same query:

$$\mathsf{Q} = \sum_{a,b,c,d,e} \mathsf{R}(\mathsf{a},\mathsf{b}) \cdot \mathsf{S}(\mathsf{a},\mathsf{c},\mathsf{e}) \cdot \mathsf{T}(\mathsf{c},\mathsf{d})$$

Use factors with different payloads, e.g.,:

- R(a, b) = 1 if $(a, b) \in R$ and 0 otherwise
- S(a, c, e) = c if $(a, c, e) \in S$ and 0 otherwise
- ▶ T(c,d) = d if $(c,d) \in T$ and 0 otherwise

■ For every tuple (a, b, c, d, e), we sum up the products R(a, b) · S(a, c, e) · T(c, d)

Factor Payloads Enable Different Aggregates (2/2)

Factors R, S, and T with payloads from a ring (D, +, *, 0, 1)

 \blacktriangleright Each tuple has the payload $\mathbf{1} \in \mathsf{D}$

• Lift factors Λ_A , Λ_B , Λ_C , Λ_D , Λ_E map the domain of a variable to **D**

 $\begin{array}{ll} \textbf{COUNT} & \mbox{all lift factors map to } 1 \in \mathbb{Z} \\ \textbf{SUM(C*D)} & \mbox{} \Lambda_C[c] = c \in \mathbb{R} \mbox{ and } \Lambda_D[d] = d \in \mathbb{R}; \mbox{ all others map to } 1 \in \mathbb{R} \end{array}$

Factor Payloads Enable Different Aggregates (2/2)

Factors R, S, and T with payloads from a ring (D, +, *, 0, 1)

 \blacktriangleright Each tuple has the payload $\mathbf{1} \in \mathbf{D}$

• Lift factors Λ_A , Λ_B , Λ_C , Λ_D , Λ_E map the domain of a variable to **D**

COUNTall lift factors map to $1 \in \mathbb{Z}$ SUM(C*D) $\Lambda_C[c] = c \in \mathbb{R}$ and $\Lambda_D[d] = d \in \mathbb{R}$; all others map to $1 \in \mathbb{R}$

Lift values of a variable just before its marginalization

$$\begin{array}{c} \mathsf{R} \\ \mathsf{R} \\ \mathsf{B} \\ \mathsf{C} \\ \mathsf$$

 $V_{1}(a) = \sum_{b} (R(a, b) \cdot \Lambda_{B}(b)) \qquad V_{2}(a, c) = \sum_{e} (S(a, c, e) \cdot \Lambda_{E}(e))$ $V_{3}(c) = \sum_{d} (T(c, d) \cdot \Lambda_{D}(d)) \qquad V_{4}(a) = \sum_{c} (V_{2}(a, c) \cdot V_{3}(c) \cdot \Lambda_{C}(c))$ $Q = \sum_{c} (V_{1}(a) \cdot V_{4}(a) \cdot \Lambda_{A}(a))$

R		S		Т	
ΑB	#	ВC	#	СA	#
$a_1 \ b_1$	2	$b_1 c_1$	2	$c_1 a_1$	1
$a_2 b_1$	3	$b_1 c_2$	1	<i>c</i> ₂ <i>a</i> ₁	3
				C ₂ a ₂	3

R	S	Т	$R \cdot S \cdot T$
A B #	B C #	<i>C A</i> #	A B C #
<i>a</i> ₁ <i>b</i> ₁ 2	$b_1 c_1 = 2$	<i>c</i> ₁ <i>a</i> ₁ 1	$a_1 \ b_1 \ c_1 \ \Big \ 2 \cdot 2 \cdot 1 = 4$
$a_2 b_1 = 3$	$b_1 c_2 \mid 1$	$c_2 a_1 = 3$	
		$c_2 a_2 = 3$	

R		S		Т			R·	$S \cdot T$
ΑΒ	#	ВC	#	СA	#	Α	ВC	#
$a_1 b_1$	2	$b_1 c_1$	2	c ₁ a ₁	1	a_1	$b_1 c_1$	$2 \cdot 2 \cdot 1 = 4$
$a_2 b_1$	3	$b_1 c_2$	1	<i>c</i> ₂ <i>a</i> ₁	3	a_1	$b_1 c_2$	$2\cdot 1\cdot 3=6$
				c ₂ a ₂	3	a_2	$b_1 c_2$	$3\cdot 1\cdot 3=9$

Triangle Count Query:
$$Q = \sum_{a,b,c} R(a,b) \cdot S(b,c) \cdot T(c,a)$$

R	5	Т	R ·	$S \cdot T$
A B #	B C #	<i>C A</i> #	АВС	#
a ₁ b ₁ 2	$b_1 c_1 = 2$	$c_1 a_1 1$	$a_1 \ b_1 \ c_1$	$2 \cdot 2 \cdot 1 = 4$
$a_2 b_1 = 3$	$b_1 c_2 = 1$	<i>c</i> ₂ <i>a</i> ₁ 3	$a_1 b_1 c_2$	$2 \cdot 1 \cdot 3 = 6$
		c a 3	a> b1 C>	$3 \cdot 1 \cdot 3 = 9$



Database of three factors R, S, T

Triangle Count Query:
$$Q = \sum_{a,b,c} R(a,b) \cdot S(b,c) \cdot T(c,a)$$

 A single-tuple update is a factor mapping a tuple to a non-zero value (positive for insertions, negative for deletions)

R	S	Т	$R \cdot S \cdot T$
A B #	B C #	<i>C A</i> #	A B C #
$\begin{array}{c c} a_1 & b_1 & 2 \\ a_2 & b_1 & 3 \end{array}$	$\begin{array}{c cccc} b_1 & c_1 & 2 \\ b_1 & c_2 & 1 \end{array}$	$ \begin{array}{c cccc} c_1 & a_1 & 1 \\ c_2 & a_1 & 3 \\ c_2 & a_2 & 3 \end{array} $	$\begin{array}{c ccccccccccccccccccccccccccccccccccc$

$\delta R = \{(a_2, b_1) \mapsto -2\}$				
A B	#			
$a_2 b_1$	-2			



Database of three factors R, S, T

Triangle Count Query:
$$Q = \sum_{a,b,c} R(a,b) \cdot S(b,c) \cdot T(c,a)$$

 A single-tuple update is a factor mapping a tuple to a non-zero value (positive for insertions, negative for deletions)

R	5	Т	$R \cdot S \cdot T$
A B #	B C #	<i>C A</i> #	A B C #
$\begin{array}{c c} a_1 & b_1 & 2 \\ a_2 & b_1 & 3 \end{array}$	$\begin{array}{c ccc} b_1 & c_1 & 2 \\ b_1 & c_2 & 1 \end{array}$	$ \begin{array}{c ccc} c_1 & a_1 & 1 \\ c_2 & a_1 & 3 \\ c_2 & a_2 & 3 \end{array} $	$\begin{array}{c c c c c c c c c c c c c c c c c c c $

$\delta R = \{(a_2, b_1) \mapsto -2\}$				
A B	#			
$a_2 b_1$	-2			



Database of three factors R, S, T

Triangle Count Query:
$$Q = \sum_{a,b,c} R(a,b) \cdot S(b,c) \cdot T(c,a)$$

 A single-tuple update is a factor mapping a tuple to a non-zero value (positive for insertions, negative for deletions)

R	5	Т	$R \cdot S \cdot T$
A B #	B C #	<i>C A</i> #	A B C #
$ \begin{array}{c ccccccccccccccccccccccccccccccccccc$	$\begin{array}{c cccc} b_1 & c_1 & 2 \\ b_1 & c_2 & 1 \end{array}$	$ \begin{array}{c cccc} c_1 & a_1 & 1 \\ c_2 & a_1 & 3 \\ c_2 & a_2 & 3 \end{array} $	$\begin{array}{c ccccccccccccccccccccccccccccccccccc$

$\delta R = \{(a_2, b_1) \mapsto -2\}$				
AB	#			
$a_2 b_1$	-2			



Database of three factors R, S, T

Triangle Count Query:
$$Q = \sum_{a,b,c} R(a,b) \cdot S(b,c) \cdot T(c,a)$$

 A single-tuple update is a factor mapping a tuple to a non-zero value (positive for insertions, negative for deletions)

A B #	B C #	C A #	A B C #
$ \begin{array}{c ccccccccccccccccccccccccccccccccccc$	$\begin{array}{c ccc} b_1 & c_1 & 2 \\ b_1 & c_2 & 1 \end{array}$	$ \begin{array}{cccc} c_1 & a_1 & 1 \\ c_2 & a_1 & 3 \\ c_2 & a_2 & 3 \end{array} $	$\begin{array}{c c c c c c c c c c c c c c c c c c c $

$\delta R = \{(a_2, b_1) \mapsto -2\}$				
ΑB	#			
$a_2 b_1$	-2			



• Database of three factors R, S, T

Triangle Count Query:
$$Q = \sum_{a,b,c} R(a,b) \cdot S(b,c) \cdot T(c,a)$$

 A single-tuple update is a factor mapping a tuple to a non-zero value (positive for insertions, negative for deletions)

R	<i>S</i>	T	$R \cdot S \cdot T$
A B #	B C #	<i>C A #</i>	A B C #
<i>a</i> ₁ <i>b</i> ₁ 2	$b_1 c_1 2$	$c_1 a_1 \mid 1$	$a_1 b_1 c_1 2 \cdot 2 \cdot 1 = 4$
$-a_2 b_1 - 3$	$b_1 c_2 = 1$	$c_2 a_1 = 3$	$a_1 \ b_1 \ c_2 2 \cdot 1 \cdot 3 = 6$
$a_2 b_1 = 1$. <u> </u>	<i>c</i> ₂ <i>a</i> ₂ 3	$\begin{array}{c c c c c c c c c c c c c c c c c c c $
		<u>,</u>	$a_2 \ b_1 \ c_2 \ \ 1 \cdot 1 \cdot 3 = 3$
Ť			\downarrow
$\delta R = \{(a_2, b_1) \mapsto -2\}$			Q
A B	#		Ø #
a ₂ b ₁	-2		() 4+6+9=19

Database of three factors R, S, T

Triangle Count Query:
$$Q = \sum_{a,b,c} R(a,b) \cdot S(b,c) \cdot T(c,a)$$

 A single-tuple update is a factor mapping a tuple to a non-zero value (positive for insertions, negative for deletions)

R	5	<i>T</i>	$R \cdot S \cdot T$
A B #	<i>B C \ #</i>	<i>C A</i> #	A B C #
<i>a</i> ₁ <i>b</i> ₁ 2	$b_1 c_1 = 2$	$c_1 a_1 \mid 1$	$a_1 b_1 c_1 2 \cdot 2 \cdot 1 = 4$
$-a_2 b_1 = 3$	$b_1 c_2 = 1$	$c_2 a_1 = 3$	$a_1 b_1 c_2 2 \cdot 1 \cdot 3 = 6$
$a_2 b_1 1$	<u>.</u>	<i>c</i> ₂ <i>a</i> ₂ 3	$-\frac{a_2}{b_1} \frac{b_1}{c_2} \frac{3 \cdot 1 \cdot 3}{3 \cdot 1 \cdot 3} = 9$
<u> </u>			$a_2 \ b_1 \ c_2 \ \ 1 \cdot 1 \cdot 3 = 3$
1			\downarrow
$\delta R = \{(a_2, b_1) \mapsto -2\}$			Q
A B	#		0 #

Part 2. Main IVM techniques: First-Order IVM

Incremental Computation for Queries

Input: Query Q, database D, updates δD to D

Task: Maintain the query result Q(D) under changes δD

$$Q(D + \delta D) = Q(D) + \delta Q(D, \delta D)$$

IVM is faster than re-computation when:

- The "merge" operation of the two query results is fast
 Upserts into hashmaps, appends to lists
- **Delta query** δQ is faster

Lower computational complexity, less data input and output

Incremental Computation for Queries

Input: Query Q, database D, updates δD to D

Task: Maintain the query result Q(D) under changes δD

$$Q(D + \delta D) = Q(D) + \delta Q(D, \delta D)$$

IVM is faster than re-computation when:

- The "merge" operation of the two query results is fast
 Upserts into hashmaps, appends to lists
- Delta query δQ is faster

Lower computational complexity, less data input and output

Question: How can we express the delta queries?

Query Language Closed under Taking Deltas

Given:

Query Q

• Update δT for factor T in Q

The delta query is defined on the structure of Q:

 $\delta(\mathsf{R} + \mathsf{S}) = \delta\mathsf{R} + \delta\mathsf{S}$

$$\delta(\mathsf{R} \cdot \mathsf{S}) = (\delta\mathsf{R} \cdot \mathsf{S}) + (\mathsf{R} \cdot \delta\mathsf{S}) + (\delta\mathsf{R} \cdot \delta\mathsf{S})$$

$$\delta(\sum_{a} \mathsf{R}) = \sum_{a} \delta \mathsf{R}$$

 $\delta \mathsf{R} = \begin{cases} \delta \mathsf{T} & \text{ if } \mathsf{R} = \mathsf{T} \\ 0 & \text{ otherwise } // \text{ Factor 0 maps empty tuple to value 0} \end{cases}$

Example: Delta for Simple Join Query

Given: Factors R and S of size N

Query:

$$Q(a,b) = R(a,b) \cdot S(b)$$

Evaluation time: $\mathcal{O}(N)$

Example: Delta for Simple Join Query

Given: Factors R and S of size N

Query:

$$Q(a,b) = R(a,b) \cdot S(b)$$

Evaluation time: $\mathcal{O}(N)$

Delta queries:

$$\delta Q(a,b) = \delta (R(a,b) \cdot S(b)) = \delta R(a,b) \cdot S(b)$$

$$\delta Q(a,b) = \delta (R(a,b) \cdot S(b)) = R(a,b) \cdot \delta S(b)$$

Single-tuple update: $\mathcal{O}(1)$ for $\delta R(a, b)$ and $\mathcal{O}(N)$ for $\delta S(b)$ Enumeration delay: $\mathcal{O}(1)$

Example: Delta for Simple Join Query

Given: Factors R and S of size N

Query:

$$Q(a,b) = R(a,b) \cdot S(b)$$

Evaluation time: $\mathcal{O}(N)$

Delta queries:

$$\delta Q(a,b) = \delta (R(a,b) \cdot S(b)) = \delta R(a,b) \cdot S(b)$$

$$\delta Q(a,b) = \delta (R(a,b) \cdot S(b)) = R(a,b) \cdot \delta S(b)$$

Single-tuple update: $\mathcal{O}(1)$ for $\delta R(a, b)$ and $\mathcal{O}(N)$ for $\delta S(b)$ Enumeration delay: $\mathcal{O}(1)$

Optimal: $\mathcal{O}(1)$ single-tuple update time and enumeration delay

Example: Delta for Simple Project-Join Query

Given: Factors R and S of size N

Query:

$$Q(a) = \sum_{b} R(a, b) \cdot S(b)$$

Evaluation time: $\mathcal{O}(N)$
Example: Delta for Simple Project-Join Query

Given: Factors R and S of size N

Query:

$$Q(a) = \sum_{b} R(a, b) \cdot S(b)$$

Evaluation time: $\mathcal{O}(N)$

Delta queries:

$$\delta Q(a) = \sum_{b} \delta R(a, b) \cdot S(b)$$

 $\delta Q(a) = \sum_{b} R(a, b) \cdot \delta S(b)$

Single-tuple update time: $\mathcal{O}(1)$ for δR and $\mathcal{O}(N)$ for δS Enumeration delay: $\mathcal{O}(1)$

Example: Delta for Simple Project-Join Query

Given: Factors R and S of size N

Query:

$$Q(a) = \sum_{b} R(a, b) \cdot S(b)$$

Evaluation time: $\mathcal{O}(N)$

Delta queries:

$$\delta Q(a) = \sum_{b} \delta R(a, b) \cdot S(b)$$

 $\delta Q(a) = \sum_{b} R(a, b) \cdot \delta S(b)$

Single-tuple update time: $\mathcal{O}(1)$ for δR and $\mathcal{O}(N)$ for δS Enumeration delay: $\mathcal{O}(1)$

Optimal: $\mathcal{O}(\sqrt{N})$ single-tuple update time and enumeration delay

Example: Delta for Triangle Join Query (1/3)

Given: Graph with edge factor E, update δE

Query:

$$Q(a, b, c) = E(a, b) \cdot E(b, c) \cdot E(c, a)$$

Evaluation time: $\mathcal{O}(|E|^{1.5})$

(clarified later in the course)

Example: Delta for Triangle Join Query (1/3)

Given: Graph with edge factor E, update δE

Query:

$$Q(a, b, c) = E(a, b) \cdot E(b, c) \cdot E(c, a)$$

Evaluation time: $\mathcal{O}(|E|^{1.5})$ (clarified later in the course) Delta query:

$$\begin{split} \delta Q(a, b, c) &= \delta \left(E(a, b) \cdot E(b, c) \cdot E(c, a) \right) \\ &= \delta E(a, b) \cdot E(b, c) \cdot E(c, a) + E(a, b) \cdot \delta \left(E(b, c) \cdot E(c, a) \right) \\ &+ \delta E(a, b) \cdot \delta \left(E(b, c) \cdot E(c, a) \right) \end{split}$$

We next expand:

$$\begin{split} \delta\left(E(b,c)\cdot E(c,a)\right) &= \delta E(b,c)\cdot E(c,a) + E(b,c)\cdot \delta E(c,a) \\ &+ \delta E(b,c)\cdot \delta E(c,a) \end{split}$$

Example: Delta for Triangle Join Query (2/3)

$$\begin{split} \delta Q(a, b, c) &= \delta E(a, b) \cdot E(b, c) \cdot E(c, a) + E(a, b) \cdot \delta E(b, c) \cdot E(c, a) \\ &+ E(a, b) \cdot E(b, c) \cdot \delta E(c, a) + E(a, b) \cdot \delta E(b, c) \cdot \delta E(c, a) \\ &+ \delta E(a, b) \cdot \delta E(b, c) \cdot E(c, a) + \delta E(a, b) \cdot E(b, c) \cdot \delta E(c, a) \\ &+ \delta E(a, b) \cdot \delta E(b, c) \cdot \delta E(c, a) \end{split}$$

Example: Delta for Triangle Join Query (2/3)

$$\begin{split} \delta Q(a, b, c) &= \delta E(a, b) \cdot E(b, c) \cdot E(c, a) + E(a, b) \cdot \delta E(b, c) \cdot E(c, a) \\ &+ E(a, b) \cdot E(b, c) \cdot \delta E(c, a) + E(a, b) \cdot \delta E(b, c) \cdot \delta E(c, a) \\ &+ \delta E(a, b) \cdot \delta E(b, c) \cdot E(c, a) + \delta E(a, b) \cdot E(b, c) \cdot \delta E(c, a) \\ &+ \delta E(a, b) \cdot \delta E(b, c) \cdot \delta E(c, a) \end{split}$$

Can we simplify this query?

Example: Delta for Triangle Join Query (2/3)

$$\begin{split} \delta Q(a, b, c) &= \delta E(a, b) \cdot E(b, c) \cdot E(c, a) + E(a, b) \cdot \delta E(b, c) \cdot E(c, a) \\ &+ E(a, b) \cdot E(b, c) \cdot \delta E(c, a) + E(a, b) \cdot \delta E(b, c) \cdot \delta E(c, a) \\ &+ \delta E(a, b) \cdot \delta E(b, c) \cdot E(c, a) + \delta E(a, b) \cdot E(b, c) \cdot \delta E(c, a) \\ &+ \delta E(a, b) \cdot \delta E(b, c) \cdot \delta E(c, a) \end{split}$$

Can we simplify this query?

Yes! Idea: Q is symmetric, so

$$\delta E(a,b) \cdot E(b,c) \cdot E(c,a) = E(a,b) \cdot \delta E(b,c) \cdot E(c,a)$$

= $E(a,b) \cdot E(b,c) \cdot \delta E(c,a)$

Example: Delta for Triangle Join Query (3/3)

$$\begin{split} \delta Q(a, b, c) &= \delta E(a, b) \cdot E(b, c) \cdot E(c, a) + E(a, b) \cdot \delta E(b, c) \cdot E(c, a) \\ &+ E(a, b) \cdot E(b, c) \cdot \delta E(c, a) + E(a, b) \cdot \delta E(b, c) \cdot \delta E(c, a) \\ &+ \delta E(a, b) \cdot \delta E(b, c) \cdot E(c, a) + \delta E(a, b) \cdot E(b, c) \cdot \delta E(c, a) \\ &+ \delta E(a, b) \cdot \delta E(b, c) \cdot \delta E(c, a) \end{split}$$

Let 3 be a factor mapping the empty tuple to multiplicity 3. Then,

$$\begin{split} \delta Q(a, b, c) &= 3 \cdot \delta E(a, b) \cdot E(b, c) \cdot E(c, a) \\ &+ 3 \cdot E(a, b) \cdot \delta E(b, c) \cdot \delta E(c, a) \\ &+ \delta E(a, b) \cdot \delta E(b, c) \cdot \delta E(c, a) \end{split}$$

Example: Delta for Triangle Join Query (3/3)

$$\begin{split} \delta Q(a, b, c) &= \delta E(a, b) \cdot E(b, c) \cdot E(c, a) + E(a, b) \cdot \delta E(b, c) \cdot E(c, a) \\ &+ E(a, b) \cdot E(b, c) \cdot \delta E(c, a) + E(a, b) \cdot \delta E(b, c) \cdot \delta E(c, a) \\ &+ \delta E(a, b) \cdot \delta E(b, c) \cdot E(c, a) + \delta E(a, b) \cdot E(b, c) \cdot \delta E(c, a) \\ &+ \delta E(a, b) \cdot \delta E(b, c) \cdot \delta E(c, a) \end{split}$$

Let 3 be a factor mapping the empty tuple to multiplicity 3. Then,

$$\begin{split} \delta Q(a, b, c) &= 3 \cdot \delta E(a, b) \cdot E(b, c) \cdot E(c, a) \\ &+ 3 \cdot E(a, b) \cdot \delta E(b, c) \cdot \delta E(c, a) \\ &+ \delta E(a, b) \cdot \delta E(b, c) \cdot \delta E(c, a) \end{split}$$

Single-tuple update time: $\mathcal{O}(|E|)$ Enumeration delay: $\mathcal{O}(1)$

Example: Delta for Triangle Join Query (3/3)

$$\begin{split} \delta Q(a, b, c) &= \delta E(a, b) \cdot E(b, c) \cdot E(c, a) + E(a, b) \cdot \delta E(b, c) \cdot E(c, a) \\ &+ E(a, b) \cdot E(b, c) \cdot \delta E(c, a) + E(a, b) \cdot \delta E(b, c) \cdot \delta E(c, a) \\ &+ \delta E(a, b) \cdot \delta E(b, c) \cdot E(c, a) + \delta E(a, b) \cdot E(b, c) \cdot \delta E(c, a) \\ &+ \delta E(a, b) \cdot \delta E(b, c) \cdot \delta E(c, a) \end{split}$$

Let 3 be a factor mapping the empty tuple to multiplicity 3. Then,

$$\delta Q(a, b, c) = 3 \cdot \delta E(a, b) \cdot E(b, c) \cdot E(c, a)$$

+ $3 \cdot E(a, b) \cdot \delta E(b, c) \cdot \delta E(c, a)$
+ $\delta E(a, b) \cdot \delta E(b, c) \cdot \delta E(c, a)$

Single-tuple update time: $\mathcal{O}(|E|)$ Enumeration delay: $\mathcal{O}(1)$ Optimal: $\mathcal{O}(\sqrt{|E|})$ single-tuple update time and $\mathcal{O}(1)$ delay

Example: Delta for Triangle Count Query

Graph with edge factor E, update δE

Query:

$$Q = \sum_{a,b,c} E(a,b) \cdot E(b,c) \cdot E(c,a)$$

Evaluation time: $\mathcal{O}(|E|^{1.5})$. Strassen-like algorithm: $\mathcal{O}(|E|^{1.41})$

Example: Delta for Triangle Count Query

Graph with edge factor E, update δE

Query:

$$Q = \sum_{a,b,c} E(a,b) \cdot E(b,c) \cdot E(c,a)$$

Evaluation time: $\mathcal{O}(|E|^{1.5})$. Strassen-like algorithm: $\mathcal{O}(|E|^{1.41})$

Delta query:

$$\begin{split} \delta Q &= \delta \left(\sum_{a,b,c} E(a,b) \cdot E(b,c) \cdot E(c,a) \right) \\ &= \sum_{a,b,c} \delta \left(E(a,b) \cdot E(b,c) \cdot E(c,a) \right) = \dots \text{ (as for the triangle join)} \end{split}$$

Single-tuple update time: $\mathcal{O}(|E|)$ Enumeration delay: $\mathcal{O}(1)$

Example: Delta for Triangle Count Query

Graph with edge factor E, update δE

Query:

$$Q = \sum_{a,b,c} E(a,b) \cdot E(b,c) \cdot E(c,a)$$

Evaluation time: $\mathcal{O}(|E|^{1.5})$. Strassen-like algorithm: $\mathcal{O}(|E|^{1.41})$

Delta query:

$$\begin{split} \delta Q &= \delta \left(\sum_{a,b,c} E(a,b) \cdot E(b,c) \cdot E(c,a) \right) \\ &= \sum_{a,b,c} \delta \left(E(a,b) \cdot E(b,c) \cdot E(c,a) \right) = \dots \text{ (as for the triangle join)} \end{split}$$

Single-tuple update time: $\mathcal{O}(|E|)$ Enumeration delay: $\mathcal{O}(1)$ Optimal: $\mathcal{O}(\sqrt{|E|})$ single-tuple update time and $\mathcal{O}(1)$ delay

Example: Delta for Acyclic Query

Given: Factors R, S, T of size N Query: $Q(a, b, c) = \sum_{d,e} R(a, b) \cdot S(a, c, e) \cdot T(a, c, d)$ Evaluation time: $O(N^2)$

Example: Delta for Acyclic Query

Given: Factors R, S, T of size N

• Query:
$$Q(a, b, c) = \sum_{d,e} R(a, b) \cdot S(a, c, e) \cdot T(a, c, d)$$

Evaluation time: $\mathcal{O}(N^2)$

Delta query for update δR (similar for δS and δT):

$$\delta Q(a, b, c) = \delta \left(\sum_{d, e} R(a, b) \cdot S(a, c, e) \cdot T(a, c, d) \right)$$
$$= \sum_{d, e} \delta \left(R(a, b) \cdot S(a, c, e) \cdot T(a, c, d) \right)$$
$$= \delta R(a, b) \cdot \left(\sum_{e} S(a, c, e) \right) \cdot \left(\sum_{d} T(a, c, d) \right)$$

Single-tuple update time: $\mathcal{O}(N)$. Enumeration delay: $\mathcal{O}(1)$

Example: Delta for Acyclic Query

Given: Factors R, S, T of size N

• Query:
$$Q(a, b, c) = \sum_{d,e} R(a, b) \cdot S(a, c, e) \cdot T(a, c, d)$$

Evaluation time: $\mathcal{O}(N^2)$

Delta query for update δR (similar for δS and δT):

$$\delta Q(a, b, c) = \delta \left(\sum_{d, e} R(a, b) \cdot S(a, c, e) \cdot T(a, c, d) \right)$$
$$= \sum_{d, e} \delta \left(R(a, b) \cdot S(a, c, e) \cdot T(a, c, d) \right)$$
$$= \delta R(a, b) \cdot \left(\sum_{e} S(a, c, e) \right) \cdot \left(\sum_{d} T(a, c, d) \right)$$

Single-tuple update time: O(N). Enumeration delay: O(1)Optimal: O(1) single-tuple update time and delay

Beyond Delta Queries

Can we achieve the aforementioned optimal update times by using additional storage?

Part 2. Main IVM techniques: Higher-Order IVM

Using Materialized Views To Lower Update Time

We reconsider the following query with factors *R* and *S* of size *N* Query:

$$Q(a,b) = R(a,b) \cdot S(b)$$

Using Materialized Views To Lower Update Time

We reconsider the following query with factors *R* and *S* of size *N* Query:

$$Q(a,b) = R(a,b) \cdot S(b)$$

• Delta query for update $\delta S(b_0)$:

$$\delta Q(a, b_0) = R(a, b_0) \cdot \delta S(b_0)$$

Single-tuple update: $\mathcal{O}(N)$, Enumeration delay: $\mathcal{O}(1)$

Using Materialized Views To Lower Update Time

We reconsider the following query with factors *R* and *S* of size *N* Query:

$$Q(a,b) = R(a,b) \cdot S(b)$$

• Delta query for update $\delta S(b_0)$:

$$\delta Q(a, b_0) = R(a, b_0) \cdot \delta S(b_0)$$

Single-tuple update: $\mathcal{O}(N)$, Enumeration delay: $\mathcal{O}(1)$

Using the views

$$V_R(b) = \sum_a R(a, b)$$
 and $V_{RS}(b) = V_R(b) \cdot S(b)$

we achieve optimality: Single-tuple update: $\mathcal{O}(1)$, Enumeration delay: $\mathcal{O}(1)$

Tree of Materialized Views

We organize the input factors and the extra views in a tree





Computation time: $\mathcal{O}(N)$

View tree

$$V'_{RS}$$

$$|$$

$$V_{RS}(b)$$

$$/$$

$$V_{R}(b)$$

$$S(b)$$

$$|$$

$$R(a, b)$$

Single-tuple update $\delta R(a_0, b_0)$



$$V'_{RS}$$

$$|$$

$$V_{RS}(b)$$

$$V_{R}(b) \quad S(b)$$

$$|$$

$$\delta R(a_{0}, b_{0})$$

Single-tuple update $\delta R(a_0, b_0)$





View tree $\delta V_{RS}'$ | $\delta V_{RS}(b_0)$ $\langle \cdot \rangle$ $\delta V_R(b_0) \quad S(b)$ | $\delta R(a_0, b_0)$ Single-tuple update $\delta R(a_0, b_0)$ $\delta V'_{RS} = \sum_{b_0} \delta V_{RS}(b_0) = V_{RS}(b_0)$ $\delta V_{RS}(b_0) = \delta V_R(b_0) \cdot S(b_0)$ $\delta V_R(b_0) = \sum_{a_0} \delta R(a_0, b_0) = \delta R(a_0, b_0)$

View tree $\delta V'_{RS}$ $\delta V_{RS}(b_0)$ $\delta V_R(b_0) = S(b)$ $\delta R(a_0, b_0)$ Single-tuple update $\delta R(a_0, b_0)$ $\delta V'_{RS} = \sum_{b_0} \delta V_{RS}(b_0) = V_{RS}(b_0)$ $\delta V_{RS}(b_0) = \delta V_R(b_0) \cdot S(b_0)$ $\delta V_R(b_0) = \sum_{a_0} \delta R(a_0, b_0) = \delta R(a_0, b_0)$

Update time: O(1)











Enumeration from View Tree

Enumeration

View tree

$$V'_{RS}$$

$$|$$

$$V_{RS}(b)$$

$$V_{R}(b) \quad S(b)$$

$$|$$

$$R(a, b)$$

Enumeration from View Tree

Enumeration

View tree

$$V_{RS}' | V_{RS}(b) \\ \swarrow \\ V_{RS}(b) \\ \swarrow \\ V_{R}(b) \\ S(b) \\ | \\ R(a, b)$$

Enumeration from View Tree

Enumeration

Is V'_{RS} empty? If yes, stop!

■ Iterate over (distinct) b's in V_{RS}

 V_{RS}' $V_{RS}(b)$ $V_{R}(b) \qquad (b)$

R(a, b)

View tree
Enumeration from View Tree

View tree

 V'_{RS}

 $V_{RS}(b)$

 $V_R(b) = S(b)$

R(a, b)

Enumeration

- Is V'_{RS} empty? If yes, stop!
- Iterate over (distinct) b's in V_{RS}
- For each *b*, iterate over (distinct) *a*'s in *R*(*a*, *b*)

Requires an index $b \mapsto a$ over R

Enumeration from View Tree

View tree

 V'_{RS}

 $V_{RS}(b)$

 $V_R(b) = S(b)$

R(a, b)

Enumeration

- Is V'_{RS} empty? If yes, stop!
 - Iterate over (distinct) b's in V_{RS}
 - For each *b*, iterate over (distinct) *a*'s in *R*(*a*, *b*)

Requires an index $b \mapsto a$ over R

Output (*b*, *a*)

Enumeration from View Tree

Enumeration

- Is V'_{RS} empty? If yes, stop!
- Iterate over (distinct) b's in V_{RS}
- For each *b*, iterate over (distinct) *a*'s in *R*(*a*, *b*)

Requires an index $b \mapsto a$ over R

Output (b, a)

Enumeration delay: $\mathcal{O}(1)$

View tree

 V'_{RS} | $V_{RS}(b)$ / $V_{R}(b)$ S(b) | R(a, b)

How to Generalize This Example?

We need to construct view trees for the query

- Similar to query plans for static query evaluation
- Goal: Minimize the propagation time for each update

We next sketch two main approaches

DBToaster https://dbtoaster.github.io

F-IVM https://github.com/fdbresearch/FIVM

How to Generalize This Example?

We need to construct view trees for the query

- Similar to query plans for static query evaluation
- Goal: Minimize the propagation time for each update

We next sketch two main approaches

- DBToaster https://dbtoaster.github.io
 - 1. One view tree per updatable factor, which is child of root
 - 2. Materialized view over each connected component of the remaining factors
 - 3. This view exports the join and head variables
 - 4. This view is treated like a query (Go to 1)

F-IVM https://github.com/fdbresearch/FIVM

How to Generalize This Example?

We need to construct view trees for the query

- Similar to query plans for static query evaluation
- Goal: Minimize the propagation time for each update

We next sketch two main approaches

- DBToaster https://dbtoaster.github.io
 - 1. One view tree per updatable factor, which is child of root
 - 2. Materialized view over each connected component of the remaining factors
 - 3. This view exports the join and head variables
 - 4. This view is treated like a query (Go to 1)

F-IVM https://github.com/fdbresearch/FIVM

- 1. One view tree for all updatable factors
- 2. Structure of view tree follows a partial order of query variables
- 3. Query result possibly distributed over several views (previous example)

Extra Slides: DBToaster Example

Higher-Order IVM using DBToaster (1/10)

Consider again the factors R, S, T of size N and the query

$$Q(a, b, c) = \sum_{d, e} R(a, b) \cdot S(a, c, e) \cdot T(a, c, d)$$

Higher-Order IVM using DBToaster (1/10)

Consider again the factors R, S, T of size N and the query

$$Q(a, b, c) = \sum_{d, e} R(a, b) \cdot S(a, c, e) \cdot T(a, c, d)$$

DBToaster constructs a view tree for each factor:

For update δR to R

For update δS to S For upd

For update δT to T

Higher-Order IVM using DBToaster (2/10)

$$Q(a, b, c) = \sum_{d, e} R(a, b) \cdot S(a, c, e) \cdot T(a, c, d)$$

For update δR to R:

DBToaster creates the materialized view V_{ST}

$$V_{ST}(a,c) = \sum_{d,e} S(a,c,e) \cdot T(a,c,d)$$

$$= \sum_{e} S(a,c,e) \cdot \sum_{d} T(a,c,d)$$

$$= V_{S}(a,c) \cdot V_{T}(a,c)$$

Computation time: $\mathcal{O}(N)$

$$Q(a,b,c)$$

$$R(a,b) \quad V_{ST}(a,c)$$

$$V_{S}(a,c) \quad V_{T}(a,c)$$

 $S(a, c, e) \quad T(a, c, d)$

Higher-Order IVM using DBToaster (2/10)

$$Q(a, b, c) = \sum_{d, e} R(a, b) \cdot S(a, c, e) \cdot T(a, c, d)$$

For update δR to R:

DBToaster creates the materialized view V_{ST}

$$V_{ST}(a, c) = \sum_{d, e} S(a, c, e) \cdot T(a, c, d)$$

$$= \sum_{e} S(a, c, e) \cdot \sum_{d} T(a, c, d)$$

$$= V_{S}(a, c) \cdot V_{T}(a, c)$$

$$Computation time: \mathcal{O}(N)$$
The delta query for update $\delta R(a_{0}, b_{0})$ using V_{ST} is then:
$$S(a, c, e) T(a, c, d)$$

 $\delta Q(a, b, c) = \delta R(a_0, b_0) \cdot V_{ST}(a_0, c)$

Single-tuple update time: $\mathcal{O}(N)$

Higher-Order IVM using DBToaster (3/10)

$$Q(a, b, c) = \sum_{d, e} R(a, b) \cdot S(a, c, e) \cdot T(a, c, d)$$

The materialized view $V_{\mbox{\scriptsize ST}}$ needs to be updated, too

Delta query for update $\delta S(a_0, c_0, e_0)$:

$$\delta V_{S}(a_{0}, c_{0}) = \sum_{e_{0}} \delta S(a_{0}, c_{0}, e_{0}) = \delta S(a_{0}, c_{0}, e_{0})$$

$$\delta V_{ST}(a_{0}, c_{0}) = \delta V_{S}(a_{0}, c_{0}) \cdot V_{T}(a_{0}, c_{0})$$

$$Q(a, b, c)$$

$$(a, b) \delta V_{ST}(a_0, c_0)$$

$$\delta V_{S}(a_0, c_0) V_{T}(a, c)$$

$$|$$

$$\delta S(a_0, c_0, c_0) T(a, c, d)$$

Higher-Order IVM using DBToaster (3/10)

$$Q(a, b, c) = \sum_{d, e} R(a, b) \cdot S(a, c, e) \cdot T(a, c, d)$$

The materialized view $V_{\it ST}$ needs to be updated, too

Delta query for update $\delta S(a_0, c_0, e_0)$:

$$\delta V_S(a_0, c_0) = \sum_{e_0} \delta S(a_0, c_0, e_0) = \delta S(a_0, c_0, e_0)$$

 $\delta V_{ST}(a_0, c_0) = \delta V_S(a_0, c_0) \cdot V_T(a_0, c_0)$

Single-tuple update time: $\mathcal{O}(1)$

$$Q(a, b, c)$$

$$(a, b) \quad \delta V_{ST}(a_0, c_0)$$

$$\delta V_S(a_0, c_0) \quad V_T(a, c)$$

$$| \qquad |$$

$$\delta S(a_0, c_0, e_0)T(a, c, d)$$

Higher-Order IVM using DBToaster (4/10)

$$Q(a, b, c) = \sum_{d, e} R(a, b) \cdot S(a, c, e) \cdot T(a, c, d)$$

The materialized view V_{ST} needs to be updated, too Delta query for update $\delta T(a_0, c_0, d_0)$:

$$\delta V_T(a_0, c_0) = \sum_{d_0} \delta T(a_0, c_0, d_0) = \delta T(a_0, c_0, d_0)$$

 $\delta V_{ST}(a_0, c_0) = V_S(a_0, c_0) \cdot \delta V_T(a_0, c_0)$

$$Q(a, b, c)$$

$$R(a, b) \ \delta V_{ST}(a_0, c_0)$$

$$V_S(a, c) \ \delta V_T(a_0, c_0)$$

$$| \qquad |$$

$$S(a, c, e) \delta T(a_0, c_0, d_0)$$

Higher-Order IVM using DBToaster (4/10)

$$Q(a, b, c) = \sum_{d, e} R(a, b) \cdot S(a, c, e) \cdot T(a, c, d)$$

The materialized view V_{ST} needs to be updated, too Delta query for update $\delta T(a_0, c_0, d_0)$:

$$\begin{split} \delta V_T(a_0,c_0) &= \sum_{d_0} \delta T(a_0,c_0,d_0) = \delta T(a_0,c_0,d_0) \\ \delta V_{ST}(a_0,c_0) &= V_S(a_0,c_0) \cdot \delta V_T(a_0,c_0) \end{split}$$

Single-tuple update time: $\mathcal{O}(1)$

$$Q(a, b, c)$$

$$(a, b) \quad \delta V_{ST}(a_0, c_0)$$

$$V_S(a, c) \quad \delta V_T(a_0, c_0)$$

$$| \qquad |$$

$$S(a, c, e) \delta T(a_0, c_0, d_0)$$

Higher-Order IVM using DBToaster (5/10)

$$Q(a, b, c) = \sum_{d, e} R(a, b) \cdot S(a, c, e) \cdot T(a, c, d)$$

For update δS to S:

DBToaster creates the materialized view V_{RT}

$$V_{RT}(a, b, c) = \sum_{d} R(a, b) \cdot T(a, c, d)$$

= $R(a, b) \cdot \sum_{d} T(a, c, d)$
= $R(a, b) \cdot V_{T}(a, c)$
Computation time: $\mathcal{O}(N^{2})$
$$Q(a, b, c)$$

 $S(a, c, e) V_{RT}(a, b, c)$
 $R(a, b) V_{T}(a, c)$

T(a, c, d)

Higher-Order IVM using DBToaster (5/10)

$$Q(a, b, c) = \sum_{d, e} R(a, b) \cdot S(a, c, e) \cdot T(a, c, d)$$

For update δS to S:

DBToaster creates the materialized view V_{RT}

$$V_{RT}(a, b, c) = \sum_{d} R(a, b) \cdot T(a, c, d)$$

= $R(a, b) \cdot \sum_{d} T(a, c, d)$
= $R(a, b) \cdot V_{T}(a, c)$
Computation time: $\mathcal{O}(N^{2})$
ta query for update $\delta S(a_{0}, c_{0}, e_{0})$ using V_{RT} is:

T(a, c, d)

The delta query for update
$$\delta S(a_0, c_0, e_0)$$
 using V_{RT} is:

$$\begin{split} \delta Q(a_0, b, c_0) &= \sum_{e_0} \delta S(a_0, c_0, e_0) \cdot V_{RT}(a_0, b, c_0) \\ &= \delta S(a_0, c_0, e_0) \cdot V_{RT}(a_0, b, c_0) \\ \text{Single-tuple update time: } \mathcal{O}(N) \end{split}$$

Higher-Order IVM using DBToaster (6/10)

$$Q(a, b, c) = \sum_{d, e} R(a, b) \cdot S(a, c, e) \cdot T(a, c, d)$$

The materialized view V_{RT} needs to be updated, too Delta query for update $\delta R(a_0, b_0)$:

$$\delta V_{RT}(a_0, b_0, c) = \delta R(a_0, b_0) \cdot V_T(a_0, c)$$

Single-tuple update time: $\mathcal{O}(N)$

Q(a, b, c) $S(a, c, e) \quad \delta V_{RT}(a_0, b_0, c)$ $\delta R(a_0, b_0) \quad V_T(a, c)$ | T(a, c, d)

Higher-Order IVM using DBToaster (7/10)

$$Q(a, b, c) = \sum_{d, e} R(a, b) \cdot S(a, c, e) \cdot T(a, c, d)$$

The materialized view V_{RT} needs to be updated, too

Delta queries for update $\delta T(a_0, c_0, d_0)$:

$$\delta V_T(a_0, c_0) = \sum_{d_0} \delta T(a_0, c_0, d_0) = \delta T(a_0, c_0, d_0)$$

$$\delta V_{RT}(a_0,b,c_0)=R(a_0,b)\cdot \delta V_T(a_0,c_0)$$

Single-tuple update time: $\mathcal{O}(1)$ and $\mathcal{O}(N)$, respectively

$$Q(a, b, c)$$

$$S(a, c, e) \quad \delta V_{RT}(a_0, b, c_0)$$

$$R(a, b) \quad \delta V_T(a_0, c_0)$$

$$|$$

$$\delta T(a_0, c_0, d_0)$$

Higher-Order IVM using DBToaster (8/10)

$$Q(a, b, c) = \sum_{d, e} R(a, b) \cdot S(a, c, e) \cdot T(a, c, d)$$

For update δT to T (analogous to update to S):

DBToaster creates the materialized view V_{RS}

$$V_{RS}(a, b, c) = \sum_{e} R(a, b) \cdot S(a, c, e)$$

= $R(a, b) \cdot \sum_{e} S(a, c, e)$
= $R(a, b) \cdot V_{S}(a, c)$
Computation time: $\mathcal{O}(N^{2})$
$$Q(a, b, c)$$

 $T(a, c, d) V_{RS}(a, b, c)$
 $R(a, b) V_{S}(a, c)$

S(a, c, e)

Higher-Order IVM using DBToaster (8/10)

$$Q(a, b, c) = \sum_{d, e} R(a, b) \cdot S(a, c, e) \cdot T(a, c, d)$$

For update δT to T (analogous to update to S):

DBToaster creates the materialized view V_{RS}

$$V_{RS}(a, b, c) = \sum_{e} R(a, b) \cdot S(a, c, e)$$

= $R(a, b) \cdot \sum_{e} S(a, c, e)$
= $R(a, b) \cdot V_{S}(a, c)$
Computation time: $\mathcal{O}(N^{2})$
Welts group for undate $\delta T(a, c, d)$ using V_{CS} :
$$R(a, b) = V_{S}(a, c)$$

$$R(a, b) = V_{S}(a, c)$$

S(a, c, e)

The delta query for update $\delta T(a_0, c_0, d_0)$ using V_{RS} :

$$\begin{split} \delta Q(a_0, b, c_0) &= \sum_{d_0} \delta T(a_0, c_0, d_0) \cdot V_{RS}(a_0, b, c_0) \\ &= \delta T(a_0, c_0, d_0) \cdot V_{RS}(a_0, b, c_0) \\ \text{Single-tuple update time: } \mathcal{O}(N) \end{split}$$

Higher-Order IVM using DBToaster (9/10)

$$Q(a, b, c) = \sum_{d, e} R(a, b) \cdot S(a, c, e) \cdot T(a, c, d)$$

The materialized view V_{RS} needs to be updated, too Delta query for update $\delta R(a_0, b_0)$:

$$\delta V_{RS}(a_0, b_0, c) = \delta R(a_0, b_0) \cdot V_S(a_0, c)$$

Single-tuple update time: $\mathcal{O}(N)$

Q(a, b, c) $T(a, c, d) \quad \delta V_{RS}(a_0, b_0, c)$ $\delta R(a_0, b_0) \quad V_S(a, c)$ | S(a, c, e)

Higher-Order IVM using DBToaster (10/10)

$$Q(a, b, c) = \sum_{d, e} R(a, b) \cdot S(a, c, e) \cdot T(a, c, d)$$

The materialized view V_{RS} needs to be updated, too

Delta queries for update $\delta S(a_0, c_0, e_0)$: $\delta V_S(a_0, c_0) = \sum_{e_0} \delta S(a_0, c_0, e_0) = \delta S(a_0, c_0, e_0)$ $T(a, c, d) \quad \delta V_{RS}(a_0, b, c_0)$ $R(a, b) \quad \delta V_S(a_0, c_0)$ $R(a, b) \quad \delta V_S(a_0, c_0)$

Q(a, b, c)

Single-tuple update time: O(1) and O(N), respectively

DBToaster's Update Time Can Be Sub-Optimal

Two sources of sub-optimality (among others)

- 1. Too many views and view trees to maintain
- 2. The entire (super-linear) query result available at root views

DBToaster's Update Time Can Be Sub-Optimal

Two sources of sub-optimality (among others)

- 1. Too many views and view trees to maintain
- 2. The entire (super-linear) query result available at root views

How to overcome them?

- 1. One view tree for all updatable factors
 - \Rightarrow Share views across updates
- 2. Keep the query result distributed over several views
 - \Rightarrow Factorized representation of the query result

Both features are supported by F-IVM

End of Extra Slides

Higher-Order IVM using F-IVM

F-IVM uses a partial order on the query variables:

The variables of a factor are on the same path

Challenge: Find variable order that guarantees asymptotically minimal update time (among all variable orders)

 Hard problem already for static query evaluation, more in the FAQ lecture

Higher-Order IVM using F-IVM

F-IVM uses a partial order on the query variables:

The variables of a factor are on the same path

Challenge: Find variable order that guarantees asymptotically minimal update time (among all variable orders)

- Hard problem already for static query evaluation, more in the FAQ lecture
- F-IVM creates a view tree for a given variable order:
 - Each variable induces extra views
 - The views join the child views & marginalize the variables
 - The input factors are the leaves

Higher-Order IVM using F-IVM: Variable Order

Consider again the factors R, S, T of size N and the query

$$Q(a, b, c) = \sum_{d, e} R(a, b) \cdot S(a, c, e) \cdot T(a, c, d)$$

A possible variable order:



• A is root as it dominates all other variables

- Dominate: It appears in all factors of the dominated variables
- B is independent of the other variables given A

D and E are under C as they are dominated by C

D and E are independent given C

Higher-Order IVM using F-IVM: View Tree

View tree construction:

Place factors at leaves



Higher-Order IVM using F-IVM: View Tree

View tree construction:

- Place factors at leaves
- Create parent view to join children $V'_C(a,c) = V_D(a,c) \cdot V_E(a,c)$

$$V_A'(a) = V_B(a) \cdot V_C(a)$$



Higher-Order IVM using F-IVM: View Tree

View tree construction:

- Place factors at leaves
- Create parent view to join children $V'_C(a,c) = V_D(a,c) \cdot V_E(a,c)$

$$V_A'(a) = V_B(a) \cdot V_C(a)$$

 Aggregate away variables not needed for further joins

$$V_A = \sum_a V'_A(a)$$

$$V_B(a) = \sum_b R(a, b)$$

 $V_C(a) = \sum_c V'_C(a, c)$
 $V_D(a, c) = \sum_d T(a, c, d)$
 $V_F(a, c) = \sum S(a, c, e)$





Single-tuple update $\delta R(a_0, b_0)$ to R:










Single-tuple update $\delta S(a_0, c_0, e_0)$ to *S*:

View tree V_A $V_A'(a)$ $V_B(a) = V_C(a)$ $R(a,b) \quad V'_C(a,c)$ $V_D(a,c) \quad V_E(a,c)$ T(a, c, d) S(a, c, e)

Single-tuple update $\delta S(a_0, c_0, e_0)$ to *S*:

View tree V_A $V_A'(a)$ $V_B(a) = V_C(a)$ $R(a,b) \quad V'_C(a,c)$ $V_D(a,c) \quad V_E(a,c)$ $T(a, c, d) \delta S(a_0, c_0, e_0)$

Single-tuple update $\delta S(a_0, c_0, e_0)$ to S:

$$\delta V_{E}(a_{0},c_{0}) = \sum_{e_{0}} \delta S(a_{0},c_{0},e_{0}) = \delta S(a_{0},c_{0},e_{0})$$



Single-tuple update $\delta S(a_0, c_0, e_0)$ **to** *S*:

$$\delta V_E(a_0, c_0) = \sum_{e_0} \delta S(a_0, c_0, e_0) = \delta S(a_0, c_0, e_0)$$

$$\delta V_C'(\mathsf{a}_0, \mathsf{c}_0) = \delta V_E(\mathsf{a}_0, \mathsf{c}_0) \cdot V_D(\mathsf{a}_0, \mathsf{c}_0)$$



Single-tuple update $\delta S(a_0, c_0, e_0)$ to S: $\delta V_E(a_0, c_0) = \sum_{e_0} \delta S(a_0, c_0, e_0) = \delta S(a_0, c_0, e_0)$ $\delta V'_C(a_0, c_0) = \delta V_E(a_0, c_0) \cdot V_D(a_0, c_0)$ $\delta V_C(a_0) = \sum_{c_0} \delta V'_C(a_0, c_0) = \delta V'_C(a_0, c_0)$

$$V_{A} = \begin{bmatrix} V_{A} \\ V_{A}(a) \\ V_{B}(a) & \delta V_{C}(a_{0}) \\ V_{B}(a, b) & \delta V_{C}'(a_{0}, c_{0}) \\ V_{D}(a, c) & \delta V_{E}(a_{0}, c_{0}) \\ V_{D}(a, c, c) & \delta V_{E}(a_{0}, c_{0}) \\ V_{D}(a, c, c_{0}) & \delta S(a_{0}, c_{0}, e_{0}) \end{bmatrix}$$

Single-tuple update $\delta S(a_0, c_0, e_0)$ **to** *S*: $\delta V_E(a_0, c_0) = \sum_{a_0} \delta S(a_0, c_0, e_0) = \delta S(a_0, c_0, e_0)$ $\delta V'_{C}(a_{0}, c_{0}) = \delta V_{F}(a_{0}, c_{0}) \cdot V_{D}(a_{0}, c_{0})$ $\delta V_C(a_0) = \sum_{\alpha} \delta V'_C(a_0, c_0) = \delta V'_C(a_0, c_0)$ $\delta V'_{A}(a_0) = \delta V_{C}(a_0) \cdot V_{B}(a_0)$

View tree

 V_A $\delta V'_{\Lambda}(a_0)$ $V_B(a) = \delta V_C(a_0)$ $R(a,b) \delta V'_{c}(a_0,c_0)$ $V_D(a,c) \delta V_E(a_0,c_0)$ $T(a, c, d) \delta S(a_0, c_0, e_0)$

Single-tuple update $\delta S(a_0, c_0, e_0)$ **to** *S*: $\delta V_E(a_0,c_0) = \sum_{a} \delta S(a_0,c_0,e_0) = \delta S(a_0,c_0,e_0)$ $\delta V'_{C}(a_{0}, c_{0}) = \delta V_{F}(a_{0}, c_{0}) \cdot V_{D}(a_{0}, c_{0})$ $\delta V_C(a_0) = \sum_{c} \delta V'_C(a_0, c_0) = \delta V'_C(a_0, c_0)$ $\delta V'_{A}(a_0) = \delta V_{C}(a_0) \cdot V_{B}(a_0)$ $\delta V_A = \sum_{a_0} \delta V'_A(a_0) = \delta V'_A(a_0)$

View tree

 δV_A $\delta V'_{\rm A}(a_0)$ $/ \setminus$ $V_B(a) \quad \delta V_C(a_0)$ $R(a,b) \delta V'_{c}(a_0,c_0)$ $V_D(a,c) \delta V_E(a_0,c_0)$ $T(a, c, d) \delta S(a_0, c_0, e_0)$

View tree δV_A

 $\delta V'_{\rm A}(a_0)$

 $/ \setminus$

Single-tuple update $\delta S(a_0, c_0, e_0)$ **to** *S*: $\delta V_E(a_0,c_0) = \sum_{a_0} \delta S(a_0,c_0,e_0) = \delta S(a_0,c_0,e_0)$ $\delta V'_{C}(a_{0}, c_{0}) = \delta V_{E}(a_{0}, c_{0}) \cdot V_{D}(a_{0}, c_{0})$ $\delta V_C(\mathbf{a}_0) = \sum_{\mathbf{a}} \delta V'_C(\mathbf{a}_0, \mathbf{c}_0) = \delta V'_C(\mathbf{a}_0, \mathbf{c}_0)$ $V_B(a) = \delta V_C(a_0)$ $\delta V'_{A}(a_0) = \delta V_{C}(a_0) \cdot V_{B}(a_0)$ $R(a,b) \delta V'_{c}(a_0,c_0)$ $\delta V_A = \sum_{a_0} \delta V'_A(a_0) = \delta V'_A(a_0)$ $V_D(a,c) \delta V_E(a_0,c_0)$ For each updated view/factor A: $T(a, c, d) \delta S(a_0, c_0, e_0)$ $A := A + \delta A$

Each view update takes $\mathcal{O}(1)$ time

Single-tuple update $\delta T(a_0, c_0, d_0)$ to T:

View tree V_A $V_A'(a)$ $V_B(a)$ $V_C(a)$ $R(a,b) = V'_C(a,c)$ $V_D(a,c) = V_E(a,c)$ T(a, c, d) S(a, c, e)

Single-tuple update $\delta T(a_0, c_0, d_0)$ to T:

View tree V_A $V_A'(a)$ $V_B(a)$ $V_C(a)$ $R(a,b) \quad V'_C(a,c)$ $V_D(a,c) = V_E(a,c)$ $\delta T(a_0, c_0, d_0)S(a, c, e)$

Single-tuple update $\delta T(a_0, c_0, d_0)$ to T:

$$\delta V_D(\mathsf{a}_0, \mathsf{c}_0) = \sum_{\mathsf{d}_0} \delta T(\mathsf{a}_0, \mathsf{c}_0, \mathsf{d}_0) = \delta T(\mathsf{a}_0, \mathsf{c}_0, \mathsf{d}_0)$$



Single-tuple update $\delta T(a_0, c_0, d_0)$ to T:

$$\delta V_D(a_0, c_0) = \sum_{d_0} \delta T(a_0, c_0, d_0) = \delta T(a_0, c_0, d_0)$$

View tree

 $\delta V_C'(a_0,c_0) = V_E(a_0,c_0) \cdot \delta V_D(a_0,c_0)$



Single-tuple update $\delta T(a_0, c_0, d_0)$ to T: $\delta V_D(a_0, c_0) = \sum_{d_0} \delta T(a_0, c_0, d_0) = \delta T(a_0, c_0, d_0)$ $\delta V'_C(a_0, c_0) = V_E(a_0, c_0) \cdot \delta V_D(a_0, c_0)$ $\delta V_C(a_0) = \sum_{c_0} \delta V'_C(a_0, c_0) = \delta V'_C(a_0, c_0)$

View tree

 V_A $V'_{\Lambda}(a)$ $V_B(a) \qquad \delta V_C(a_0)$ $R(a, b) \quad \delta V'_C(a_0, c_0)$ $\delta V_D(a_0, c_0) V_E(a, c)$ $\delta T(a_0, c_0, d_0)S(a, c, e)$

Single-tuple update $\delta T(a_0, c_0, d_0)$ to T: $\delta V_D(a_0, c_0) = \sum_{d_0} \delta T(a_0, c_0, d_0) = \delta T(a_0, c_0, d_0)$ $\delta V'_{C}(a_{0}, c_{0}) = V_{E}(a_{0}, c_{0}) \cdot \delta V_{D}(a_{0}, c_{0})$ $\delta V_C(\mathbf{a}_0) = \sum_{c_0} \delta V'_C(\mathbf{a}_0, c_0) = \delta V'_C(\mathbf{a}_0, c_0)$ $\delta V'_{A}(a_0) = \delta V_{B}(a_0) \cdot \delta V_{C}(a_0)$ $\delta V_D(a_0, c_0) V_E(a, c)$

 $\delta T(a_0, c_0, d_0)S(a, c, e)$

View tree V_A

 $\delta V_A'(a_0)$ $V_B(a) \qquad \delta V_C(a_0)$

 $R(a, b) \quad \delta V'_C(a_0, c_0)$

Single-tuple update $\delta T(a_0, c_0, d_0)$ to T: $\delta V_D(a_0, c_0) = \sum_{d_0} \delta T(a_0, c_0, d_0) = \delta T(a_0, c_0, d_0)$ $\delta V'_{C}(a_{0}, c_{0}) = V_{E}(a_{0}, c_{0}) \cdot \delta V_{D}(a_{0}, c_{0})$ $\delta V_C(\mathbf{a}_0) = \sum_{c_0} \delta V'_C(\mathbf{a}_0, \mathbf{c}_0) = \delta V'_C(\mathbf{a}_0, \mathbf{c}_0)$ $\delta V'_{A}(a_0) = \delta V_{B}(a_0) \cdot \delta V_{C}(a_0)$ $\delta V_A = \sum \delta V'_A(a_0) = \delta V'_A(a_0)$

View tree

 δV_A $\delta V'_A(a_0)$ $V_B(a) \quad \delta V_C(a_0)$ $R(a,b) \delta V'_{c}(a_0,c_0)$ $\delta V_D(a_0, c_0) V_E(a, c)$ $\delta T(a_0, c_0, d_0)S(a, c, e)$

Single-tuple update $\delta T(a_0, c_0, d_0)$ to T: $\delta V_D(a_0, c_0) = \sum_{d_0} \delta T(a_0, c_0, d_0) = \delta T(a_0, c_0, d_0)$ View tree $\delta V'_{C}(a_{0}, c_{0}) = V_{E}(a_{0}, c_{0}) \cdot \delta V_{D}(a_{0}, c_{0})$ δV_A $\delta V_{\mathcal{C}}(\mathsf{a}_0) = \sum_{c_0} \delta V_{\mathcal{C}}'(\mathsf{a}_0, c_0) = \delta V_{\mathcal{C}}'(\mathsf{a}_0, c_0)$ $\delta V'_{\rm A}(a_0)$ $V_B(a) \qquad \delta V_C(a_0)$ $\delta V'_{A}(a_0) = \delta V_{B}(a_0) \cdot \delta V_{C}(a_0)$ $R(a, b) \quad \delta V'_C(a_0, c_0)$ $\delta V_A = \sum \delta V'_A(a_0) = \delta V'_A(a_0)$ $\delta V_D(a_0, c_0) \quad V_E(a, c)$ For each updated view/factor A: $\delta T(a_0, c_0, d_0)S(a, c, e)$ $A := A + \delta A$

Each view update takes $\mathcal{O}(1)$ time

Higher-Order IVM using F-IVM: Enumeration



Enumeration for Q(a, b, c) with constant delay

- Top-down in the view tree
- Views calibrated for variables underneath
- Guaranteed to get matching tuples in views below

Higher-Order IVM using F-IVM: Enumeration

View tree

 V_A $V'_{\rm A}(a)$ $V_B(a) = V_C(a)$ $R(a,b) = V'_{c}(a,c)$ $V_D(a,c) \quad V_E(a,c)$ T(a, c, d) S(a, c, e) Enumeration for Q(a, b, c) with constant delay

- Top-down in the view tree
- Views calibrated for variables underneath
- Guaranteed to get matching tuples in views below
- Is V_A empty? If yes, stop.
- Iterate over a's in $V'_A(a)$
- For each a, iterate over b's in R(a, b) Requires index a → b in R
- For each a, b, iterate over c's in V'_C(a, c) Requires index a → c in V'_C

Output (a, b, c)

DBToaster & F-IVM are not optimal

for a variety of queries,

including the triangle count

Extra Slides: DBToaster and F-IVM for Triangle Count

F-IVM for the Triangle Count Query (1/2)

$$Q = \sum_{a,b,c} R(a,b) \cdot S(b,c) \cdot T(c,a)$$

A single-tuple update δR to R takes time $\mathcal{O}(1)$:

F-IVM creates the materialized view V_{ST}

$$V_{ST}(a,b) = \sum_{c} S(b,c) \cdot T(c,a)$$

Computation time: $\mathcal{O}(N^2)$

 $\begin{array}{c}
Q \\
R(a,b) V_{ST}(a,b) \\
S(b,c) T(c,a)
\end{array}$

F-IVM for the Triangle Count Query (1/2)

$$Q = \sum_{a,b,c} R(a,b) \cdot S(b,c) \cdot T(c,a)$$

A single-tuple update δR to R takes time $\mathcal{O}(1)$:

F-IVM creates the materialized view V_{ST}

$$V_{ST}(a,b) = \sum_{c} S(b,c) \cdot T(c,a)$$

The delta query for update δR using V_{ST} is:

$$\delta Q = \sum_{a_0,b_0} \delta R(a_0,b_0) \cdot V_{ST}(a_0,b_0)$$

$$= \frac{\delta R(a_0, b_0) \cdot V_{ST}(a_0, b_0)}{\delta V_{ST}(a_0, b_0)}$$



F-IVM for the Triangle Count Query (2/2)

$$Q = \sum_{a,b,c} R(a,b) \cdot S(b,c) \cdot T(c,a)$$

Q and the view V_{ST} need $\mathcal{O}(N)$ for single-tuple updates:

Delta queries for updates δS and δT :

$$\delta V_{ST}(a, b_0) = \delta S(b_0, c_0) \cdot T(c_0, a)$$

 $\delta Q = \delta V_{ST}(a, b_0) \cdot R(a, b_0)$

$$R(a,b) \qquad \delta V_{ST}(a,b_0) \\ \delta S(b_0,c_0) \qquad T(c,a)$$

F-IVM for the Triangle Count Query (2/2)

$$Q = \sum_{a,b,c} R(a,b) \cdot S(b,c) \cdot T(c,a)$$

Q and the view V_{ST} need $\mathcal{O}(N)$ for single-tuple updates: Delta queries for updates δS and δT :

$$\delta V_{ST}(a, b_0) = \delta S(b_0, c_0) \cdot T(c_0, a)$$

 $\delta Q = \delta V_{ST}(a, b_0) \cdot R(a, b_0)$
 $\delta V_{ST}(a_0, b) = \delta T(c_0, a_0) \cdot S(b, c_0)$
 $\delta Q = \delta V_{ST}(a_0, b) \cdot R(a_0, b)$

$$R(a,b) \quad \delta V_{ST}(a_0,b)$$

$$S(b,c) \quad \delta T(c_0,a_0)$$

DBToaster for the Triangle Count Query

$$Q = \sum_{a,b,c} R(a,b) \cdot S(b,c) \cdot T(c,a)$$

DBToaster also creates the materialized views V_{RT} and V_{RS} to support updates to S and T:





DBToaster for the Triangle Count Query

$$Q = \sum_{a,b,c} R(a,b) \cdot S(b,c) \cdot T(c,a)$$

DBToaster also creates the materialized views V_{RT} and V_{RS} to support updates to S and T:



They again require $\mathcal{O}(N)$ update time

DBToaster for the Triangle Count Query

$$Q = \sum_{a,b,c} R(a,b) \cdot S(b,c) \cdot T(c,a)$$

DBToaster also creates the materialized views V_{RT} and V_{RS} to support updates to S and T:



They again require $\mathcal{O}(N)$ update time

Single-tuple updates to each of the three factors require $\mathcal{O}(N)$ time

$$Q = \sum_{a,b,c} R(a,b) \cdot S(b,c) \cdot T(c,a)$$

First-order IVM: $\mathcal{O}(N)$ time for update to any factor

$$Q = \sum_{a,b,c} R(a,b) \cdot S(b,c) \cdot T(c,a)$$

- First-order IVM: $\mathcal{O}(N)$ time for update to any factor
- Higher-order IVM using DBToaster
 - creates three view trees
 - $\mathcal{O}(N)$ time for update to any factor

$$Q = \sum_{a,b,c} R(a,b) \cdot S(b,c) \cdot T(c,a)$$

- First-order IVM: O(N) time for update to any factor
- Higher-order IVM using DBToaster
 - creates three view trees
 - $\mathcal{O}(N)$ time for update to any factor
- Higher-order IVM using F-IVM
 - creates a single view tree
 - $\mathcal{O}(1)$ time for update to one factor
 - $\mathcal{O}(N)$ time for update to the other two factors

$$Q = \sum_{a,b,c} R(a,b) \cdot S(b,c) \cdot T(c,a)$$

- First-order IVM: O(N) time for update to any factor
- Higher-order IVM using DBToaster
 - creates three view trees
 - $\mathcal{O}(N)$ time for update to any factor
- Higher-order IVM using F-IVM
 - creates a single view tree
 - $\mathcal{O}(1)$ time for update to one factor
 - $\mathcal{O}(N)$ time for update to the other two factors

Can we achieve the optimal update time of $\mathcal{O}(N^{1/2})$?

$$Q = \sum_{a,b,c} R(a,b) \cdot S(b,c) \cdot T(c,a)$$

- First-order IVM: O(N) time for update to any factor
- Higher-order IVM using DBToaster
 - creates three view trees
 - $\mathcal{O}(N)$ time for update to any factor
- Higher-order IVM using F-IVM
 - creates a single view tree
 - $\mathcal{O}(1)$ time for update to one factor
 - $\mathcal{O}(N)$ time for update to the other two factors

Can we achieve the optimal update time of $\mathcal{O}(N^{1/2})$? Yes! Adaptive IVM

Extra Slides End

Part 2. Main IVM techniques: Adaptive IVM
The Triangle Count Query

The triangle count query Q returns the number of tuples in the join of R, S, and T:

$$Q = \sum_{a,b,c} R(a,b) \cdot S(b,c) \cdot T(c,a)$$



Problem: Maintain Q under single-tuple updates to R, S, and T

The Triangle Count Query

The triangle count query Q returns the number of tuples in the join of R, S, and T:

$$Q = \sum_{a,b,c} R(a,b) \cdot S(b,c) \cdot T(c,a)$$



Problem: Maintain Q under single-tuple updates to R, S, and T

Is there a **fully dynamic algorithm** that can maintain the **exact triangle count** in **worst-case optimal** time?

"Recompute from scratch"



"Recompute from scratch" $\delta R = \{(\alpha, \beta) \mapsto m\}$



"Recompute from scratch" $\delta R = \{(\alpha, \beta) \mapsto m\}$



"Recompute from scratch" $\delta R = \{(\alpha, \beta) \mapsto m\}$



- N is the database size
- Update time: O(N^{1.5}) using worst-case optimal join algorithms [Algorithmica 1997, SIGMOD R. 2013, ICDT 2014]
 Slightly better using Strassen-like matrix multiplication
- Space: $\mathcal{O}(N)$ to store input factors















• Update time: $\mathcal{O}(N)$ to intersect C-values from S and T

Space: $\mathcal{O}(N)$ to store input factors















• Time for updates to R: $\mathcal{O}(1)$ to look up in V_{ST}

Maintain V_{ST} under updates











 $V_{ST}(\beta, a) = V_{ST}(\beta, a) + \delta V_{ST}(\beta, a)$



 $V_{ST}(\beta, a) = V_{ST}(\beta, a) + \delta V_{ST}(\beta, a)$

Time for updates to S and T: O(N) to maintain V_{ST}
Space: O(N²) to store input factors and V_{ST}

Lower Bound for Maintaining the Triangle Count

The Boolean Triangle Detection Problem

Boolean Triangle Detection Query

$$Q_b = \bigvee_{a,b,c} R(a,b) \wedge S(b,c) \wedge T(c,a)$$

The Boolean Triangle Detection Problem

Boolean Triangle Detection Query

$$Q_b = \bigvee_{a,b,c} R(a,b) \wedge S(b,c) \wedge T(c,a)$$

Let **D** be the database instance and N the number of tuples in **D**. For any $\gamma > 0$, there is no algorithm that incrementally maintains Q_b with

update time	and	enumeration delay	
$\mathcal{O}(N^{rac{1}{2}-\gamma})$		$\mathcal{O}(N^{1-\gamma})$	

unless the Online Vector-Matrix-Vector Multiplication (OuMv) Conjecture fails.

Online Vector-Matrix-Vector Multiplication

The OuMv problem:

- Input: An n × n Boolean matrix M and n pairs (u₁, v₁), ..., (u_n, v_n) of Boolean column-vectors of size n arriving one after the other.
- Goal: After seeing each pair $(\mathbf{u}_r, \mathbf{v}_r)$, output $\mathbf{u}_r^\top \mathbf{M} \mathbf{v}_r$

Online Vector-Matrix-Vector Multiplication

The OuMv problem:

- Input: An n × n Boolean matrix M and n pairs (u₁, v₁), ..., (u_n, v_n) of Boolean column-vectors of size n arriving one after the other.
- **Goal:** After seeing each pair $(\mathbf{u}_r, \mathbf{v}_r)$, output $\mathbf{u}_r^\top \mathbf{M} \mathbf{v}_r$

The OuMv Conjecture

[STOC 2015]

For any $\gamma > 0$, there is no algorithm that solves OuMv in time $\mathcal{O}(n^{3-\gamma})$.

Proof Idea

Assume there is an algorithm \mathcal{A} that can maintain Triangle Detection Query Q_b with

amortized update time and enumeration delay $\mathcal{O}(N^{\frac{1}{2}-\gamma})$ $\mathcal{O}(N^{1-\gamma})$ for some $\gamma > 0$.

■ We design an algorithm *B* that uses the oracle *A* to solve OuMv in subcubic time in *n*. ⇒ Contradicts the OuMv Conjecture!

Proof Idea

 Assume there is an algorithm A that can maintain Triangle Detection Query Q_b with

amortized update time and enumeration delay $\mathcal{O}(N^{\frac{1}{2}-\gamma})$ $\mathcal{O}(N^{1-\gamma})$ for some $\gamma > 0$.

■ We design an algorithm B that uses the oracle A to solve OuMv in subcubic time in n. ⇒ Contradicts the OuMv Conjecture!

Algorithm \mathcal{B}

- Factor S encodes the matrix M: S(i,j) = M[i,j]
- In each round $r \in [n]$:
 - Factor R encodes the vector \mathbf{u}_r : $R(\mathbf{a}, i) = \mathbf{u}_r[i]$, for constant \mathbf{a}
 - Factor T encodes the vector \mathbf{v}_r : $T(j, \mathbf{a}) = \mathbf{v}_r[j]$, for constant \mathbf{a}
 - ► Then $\mathbf{u}_r^\top \mathbf{M} \mathbf{v}_r = Q_b$
 - Check whether $Q_b = 1$ using algorithm A.

Example Encoding for u, M, and v

$\mathbf{u}^ op$	М	v	u [⊤] Mv
0 1 0	0 1 0 1 1 0 1 0 1	1 0 0	1
R	S	Т	Q_b
A B val	B C val	C A val	Ø val
a 2 1	2 1 1	1 a 1	() 1
	3 1 1		
	1 2 1		
	2 2 1		

1

3 3

.

Proof Sketch: Algorithm \mathcal{B}

(1) For $i, j \in [n]$: S(i, j) = M[i, j]

 $(\leq n^2 \text{ insertions})$

Proof Sketch: Algorithm \mathcal{B}

(1) For $i, j \in [n]$: S(i, j) = M[i, j]

 $(\leq n^2 \text{ insertions})$

(2) In each round $r \in [n]$:

• Delete all tuples in R and T

 $(\leq 2n \text{ deletions})$
Proof Sketch: Algorithm \mathcal{B}

(1) For $i, j \in [n]$: S(i, j) = M[i, j] ($\leq n$

 $(\leq n^2 \text{ insertions})$

(2) In each round $r \in [n]$:

• Delete all tuples in R and T

 $(\leq 2n \text{ deletions})$

▶ Insert into R and T:
For
$$i, j \in [n]$$
: $R(\mathbf{a}, i) = \mathbf{u}_r[i]$ and $T(j, \mathbf{a}) = \mathbf{v}_r[j]$ (≤ 2n insertions)

Proof Sketch: Algorithm \mathcal{B}

(1) For $i, j \in [n]$: S(i, j) = M[i, j] ($\leq n^2$ insertions)

(2) In each round $r \in [n]$:

Delete all tuples in R and T

 $(\leq 2n \text{ deletions})$

Insert into R and T:
For
$$i, j \in [n]$$
: $R(a, i) = \mathbf{u}_r[i]$ and $T(j, a) = \mathbf{v}_r[j]$ ($\leq 2n$ insertions)

• Check $Q_b = 1$: This holds if and only if $\mathbf{u}_r^T \mathbf{M} \mathbf{v}_r = 1$

 $\mathbf{u}_{r}^{T}\mathbf{M}\mathbf{v}_{r} = 1 \Leftrightarrow \exists i, j \in [n]: \mathbf{u}_{r}[i] = 1, \mathbf{M}[i, j] = 1, \mathbf{v}_{r}[j] = 1$

Proof Sketch: Algorithm \mathcal{B}

(1) For $i, j \in [n]$: S(i, j) = M[i, j] ($\leq n^2$ insertions)

(2) In each round $r \in [n]$:

Delete all tuples in R and T

 $(\leq 2n \text{ deletions})$

► Insert into R and T: For $i, j \in [n]$: $R(\mathbf{a}, i) = \mathbf{u}_r[i]$ and $T(j, \mathbf{a}) = \mathbf{v}_r[j]$ ($\leq 2n$ insertions)

• Check $Q_b = 1$: This holds if and only if $\mathbf{u}_r^T \mathbf{M} \mathbf{v}_r = 1$

 $\mathbf{u}_r^T \mathbf{M} \mathbf{v}_r = 1 \Leftrightarrow \exists i, j \in [n] : \mathbf{u}_r[i] = 1, \mathbf{M}[i, j] = 1, \mathbf{v}_r[j] = 1$

 \mathcal{B} constructs a database of size $N = \mathcal{O}(n^2)$.

Recall $\mathcal A$ needs $\mathcal O((n^2)^{rac{1}{2}-\gamma})$ update time and $\mathcal O((n^2)^{1-\gamma})$ delay

(1) For $i, j \in [n]$: S(i, j) = M[i, j]

(2) In each round $r \in [n]$:

- ▶ Delete all tuples in *R* and *T*
- ▶ Insert into R and T: For $i, j \in [n]$: $R(a, i) = \mathbf{u}_r[i]$ and $T(j, a) = \mathbf{v}_r[j]$

Recall \mathcal{A} needs $\mathcal{O}((n^2)^{rac{1}{2}-\gamma})$ update time and $\mathcal{O}((n^2)^{1-\gamma})$ delay

(1) For $i, j \in [n]$: S(i, j) = M[i, j]



(2) In each round $r \in [n]$:

- Delete all tuples in R and T
- ▶ Insert into R and T: For $i, j \in [n]$: $R(a, i) = \mathbf{u}_r[i]$ and $T(j, a) = \mathbf{v}_r[j]$

Recall \mathcal{A} needs $\mathcal{O}((n^2)^{rac{1}{2}-\gamma})$ update time and $\mathcal{O}((n^2)^{1-\gamma})$ delay

(1) For $i, j \in [n]$: S(i, j) = M[i, j]



(2) In each round $r \in [n]$:

- Delete all tuples in R and T
- ▶ Insert into R and T: For $i, j \in [n]$: $R(a, i) = \mathbf{u}_r[i]$ and $T(j, a) = \mathbf{v}_r[j]$

$$\mathcal{O}(\underbrace{4n}_{\#\text{updates}} \cdot \underbrace{(n^2)^{\frac{1}{2}-\gamma}}_{\text{update time}}) = \mathcal{O}(n^{2-2\gamma})$$

Recall \mathcal{A} needs $\mathcal{O}((n^2)^{rac{1}{2}-\gamma})$ update time and $\mathcal{O}((n^2)^{1-\gamma})$ delay

(1) For $i, j \in [n]$: S(i, j) = M[i, j]



(2) In each round $r \in [n]$:

- Delete all tuples in R and T
- ▶ Insert into R and T: For $i, j \in [n]$: $R(a, i) = \mathbf{u}_r[i]$ and $T(j, a) = \mathbf{v}_r[j]$

$$\mathcal{O}(\underbrace{4n}_{\#\text{updates}} \cdot \underbrace{(n^2)^{\frac{1}{2}-\gamma}}_{\text{update time}}) = \mathcal{O}(n^{2-2\gamma})$$

$$\mathcal{O}(\underbrace{(n^2)^{1-\gamma}}_{\text{delay}}) = \mathcal{O}(n^{2-2\gamma})$$

Recall \mathcal{A} needs $\mathcal{O}((n^2)^{rac{1}{2}-\gamma})$ update time and $\mathcal{O}((n^2)^{1-\gamma})$ delay

(1) For $i, j \in [n]$: S(i, j) = M[i, j]



(2) In each round $r \in [n]$:

- Delete all tuples in R and T
- ▶ Insert into R and T: For $i, j \in [n]$: $R(a, i) = \mathbf{u}_r[i]$ and $T(j, a) = \mathbf{v}_r[j]$

$$\mathcal{O}(\underbrace{4n}_{\#\text{updates}} \cdot \underbrace{(n^2)^{\frac{1}{2}-\gamma}}_{\text{update time}}) = \mathcal{O}(n^{2-2\gamma})$$

• Check $Q_b = 1$: This holds if and only if $\mathbf{u}_r^T \mathbf{M} \mathbf{v}_r = 1$

$$\mathcal{O}(\underbrace{(n^2)^{1-\gamma}}_{\text{delay}}) = \mathcal{O}(n^{2-2\gamma})$$

For *n* rounds: $\mathcal{O}(n(n^{2-2\gamma} + n^{2-2\gamma})) = \mathcal{O}(n^{3-2\gamma})$

Recall \mathcal{A} needs $\mathcal{O}((n^2)^{rac{1}{2}-\gamma})$ update time and $\mathcal{O}((n^2)^{1-\gamma})$ delay

(1) For $i, j \in [n]$: S(i, j) = M[i, j]



(2) In each round $r \in [n]$:

- Delete all tuples in R and T
- ▶ Insert into R and T: For $i, j \in [n]$: $R(a, i) = \mathbf{u}_r[i]$ and $T(j, a) = \mathbf{v}_r[j]$

$$\mathcal{O}(\underbrace{4n}_{\#\text{updates}} \cdot \underbrace{(n^2)^{\frac{1}{2}-\gamma}}_{\text{update time}}) = \mathcal{O}(n^{2-2\gamma})$$

• Check $Q_b = 1$: This holds if and only if $\mathbf{u}_r^T \mathbf{M} \mathbf{v}_r = 1$

$$\mathcal{O}(\underbrace{(n^2)^{1-\gamma}}_{\text{delay}}) = \mathcal{O}(n^{2-2\gamma})$$

For *n* rounds: $\mathcal{O}(n(n^{2-2\gamma}+n^{2-2\gamma})) = \mathcal{O}(n^{3-2\gamma})$

Overall time: $\mathcal{O}(n^{3-2\gamma} + n^{3-2\gamma}) = \mathcal{O}(n^{3-2\gamma}) \Rightarrow$ Contradicts OuMv Conjecture!

Complexity bounds for the maintenance of the triangle count

Known Upper Bound	
Update Time:	0(N)
Space:	0(N)



Update time: not $\mathcal{O}(N^{\frac{1}{2}-\gamma})$ for any $\gamma > 0$

under the OuMv Conjecture

Complexity bounds for the maintenance of the triangle count

Known Upper Bound		
Update Time:	<i>O</i> (<i>N</i>)	
Space:	$\mathcal{O}(N)$	

Can the triangle count be maintained with sublinear update time?



Complexity bounds for the maintenance of the triangle count

Known Upper Bound	
Update Time:	$\mathcal{O}(N)$
Space:	$\mathcal{O}(N)$

Can the triangle count be maintained with sublinear update time? Yes: IVM^{ε} Amortized update time: $\mathcal{O}(N^{\frac{1}{2}})$ This is worst-case optimal

Known Lower Bound

Update time: not $\mathcal{O}(N^{\frac{1}{2}-\gamma})$ for any $\gamma > 0$

under the OuMv Conjecture

IVM^ε Exhibits a Time-Space Tradeoff

Given $\varepsilon \in [0,1], \, \mathsf{IVM}^{\varepsilon}$ maintains the triangle count with

- $\mathcal{O}(N^{\max\{\varepsilon,1-\varepsilon\}})$ amortized update time
- $\mathcal{O}(N^{1+\min\{\varepsilon,1-\varepsilon\}})$ space
- $\mathcal{O}(N^{\frac{3}{2}})$ preprocessing time
- $\mathcal{O}(1)$ answer time.



(Linear space possible with a slightly more involved argument)

Inside IVM $^{\varepsilon}$

Main Techniques used in IVM $^{\varepsilon}$

- Compute the delta like in first-order IVM
- Materialize views like in higher-order IVM
- New ingredient: Use adaptive processing based on data skew
 Treat *heavy* values differently from *light* values

Heavy/Light Partitioning of Factors

Partition R based on A into

- a light part $R_L = \{t \in R \mid |\sigma_{A=t.A}| < N^{\varepsilon}\},\$
- a heavy part $R_H = R \setminus R_L!$



Derived Bounds

from light part: for all A-values $a, |\sigma_{A=a}R_L| < N^{\epsilon}$.

from heavy part:

for all A-values a, $|\sigma_{A=a}R_H| \ge N^{\epsilon}$ and $|\pi_A R_H| \cdot N^{\epsilon} \le N$

 $\implies |\pi_A R_H| \le N^{1-\varepsilon}$

Heavy/Light Partitioning of Factors

Partition R based on A into

- a light part $R_L = \{t \in R \mid |\sigma_{A=t,A}| < N^{\varepsilon}\},\$
- a heavy part $R_H = R \setminus R_L!$



Derived Bounds

Heavy/Light Partitioning of Factors

Likewise, partition

- $S = S_L \cup S_H$ based on B, and
- $T = T_L \cup T_H$ based on C!

Q is the sum of skew-aware queries

$$Q = \sum_{a,b,c} R_U(a,b) \cdot S_V(b,c) \cdot T_W(c,a), \text{ for } U,V,W \in \{L,H\}.$$

Given an update $\delta R_* = \{(\alpha, \beta) \mapsto m\}$, compute the delta for each of the following skew-aware queries using a different strategy:

$$Q_{*LL} = \sum_{a,b,c} R_*(a,b) \cdot S_L(b,c) \cdot T_L(c,a)$$

$$Q_{*HH} = \sum_{a,b,c} R_*(a,b) \cdot S_H(b,c) \cdot T_H(c,a)$$

$$Q_{*LH} = \sum_{a,b,c} R_*(a,b) \cdot S_L(b,c) \cdot T_H(c,a)$$

$$Q_{*HL} = \sum_{a,b,c} R_*(a,b) \cdot S_H(b,c) \cdot T_L(c,a)$$

Given an update $\delta R_* = \{(\alpha, \beta) \mapsto m\}$, compute the delta for each of the following skew-aware queries using a different strategy:

$$\delta Q_{*LL} = \delta R_*(\alpha, \beta) \cdot \sum_c S_L(\beta, c) \cdot T_L(c, \alpha)$$

$$\delta Q_{*HH} = \delta R_*(\alpha,\beta) \cdot \sum_c S_H(\beta,c) \cdot T_H(c,\alpha)$$

$$\delta Q_{*LH} = \delta R_*(\alpha, \beta) \cdot \sum_c S_L(\beta, c) \cdot T_H(c, \alpha)$$

$$\delta Q_{*HL} = \delta R_*(\alpha,\beta) \cdot \sum_c S_H(\beta,c) \cdot T_L(c,\alpha)$$

$$\delta Q_{*LL} = \delta R_*(\alpha,\beta) \cdot \sum_c S_L(\beta,c) \cdot T_L(c,\alpha)$$



$$\delta Q_{*LL} = \delta R_*(\alpha, \beta) \cdot \sum_c S_L(\beta, c) \cdot T_L(c, \alpha)$$





Update time: $\mathcal{O}(N^{\varepsilon})$ to intersect the lists of *C*-values from S_L and T_L







Update time: $\mathcal{O}(N^{1-\varepsilon})$ to intersect the lists of C-values from S_H and T_H









Update time: $\mathcal{O}(N^{\min\{\varepsilon,1-\varepsilon\}})$ to intersect the lists of C-values from S_L and T_H

$$\delta Q_{*HL} = \delta R_*(\alpha,\beta) \cdot \sum_c S_H(\beta,c) \cdot T_L(c,\alpha)$$



 $V_{ST}(b,a) = \sum S_H(b,c) \cdot T_L(c,a)$

$$\delta Q_{*HL} = \delta R_*(\alpha,\beta) \cdot \sum_c S_H(\beta,c) \cdot T_L(c,\alpha)$$



$$V_{ST}(b,a) = \sum_{c} S_{H}(b,c) \cdot T_{L}(c,a)$$

 $T_L(c, \alpha)$

$$\delta Q_{*HL} = \delta R_*(\alpha,\beta) \cdot \sum_c S_H(\beta,c) \cdot T_L(c,\alpha)$$



 $V_{ST}(b, a) = \sum S_H(b, c) \cdot T_L(c, a)$





 $V_{ST}(b, a) = \sum S_H(b, c) \cdot T_L(c, a)$

Update time: O(1) to look up in V_{ST} , assuming V_{ST} is already materialized

Summary of Adaptive Maintenance Strategies

Maintenance for an update $\delta R_* = \{(\alpha, \beta) \mapsto m\}$:

Skew-aware View Evaluation from left to right Time $\sum R_*(a,b) \cdot S_L(b,c) \cdot T_L(c,a) \quad \delta R_*(\alpha,\beta) \cdot \sum S_L(\beta,c) \cdot T_L(c,\alpha)$ $\mathcal{O}(N^{\varepsilon})$ a.b.c $\sum R_*(a,b) \cdot S_H(b,c) \cdot T_H(c,a) \quad \frac{\delta R_*(\alpha,\beta)}{\delta R_*(\alpha,\beta)} \cdot \sum T_H(c,\alpha) \cdot S_H(\beta,c)$ $\mathcal{O}(N^{1-\varepsilon})$ a,b,c $\delta R_*(\alpha,\beta) \cdot \sum S_L(\beta,c) \cdot T_H(c,\alpha)$ $\mathcal{O}(N^{\varepsilon})$ $\sum R_*(a,b) \cdot S_L(b,c) \cdot T_H(c,a)$ or a,b,c $\mathcal{O}(N^{1-\varepsilon})$ $\delta R_*(\alpha,\beta) \cdot \sum T_H(c,\alpha) \cdot S_L(\beta,c)$ $\sum R_*(a,b) \cdot S_H(b,c) \cdot T_L(c,a) = \frac{\delta R_*(\alpha,\beta)}{\delta R_*(\alpha,\beta)} \cdot V_{ST}(\beta,\alpha)$ $\mathcal{O}(1)$ a.b.c

Overall update time: $\mathcal{O}(N^{\max(\varepsilon,1-\varepsilon)})$
Auxiliary Materialized Views

$$V_{RS}(a,c) = \sum_{b} R_{H}(a,b) \cdot S_{L}(b,c)$$

$$V_{ST}(b,a) = \sum_{c} S_{H}(b,c) \cdot T_{L}(c,a)$$

$$V_{TR}(a,c) = \sum_{a} T_H(c,a) \cdot R_L(a,b)$$

Maintain $V_{ST}(b, a) = \sum_{c} S_{H}(b, c) \cdot T_{L}(c, a)$ under update $\delta S_{H} = \{(\beta, \gamma) \mapsto m\}$



Maintain $V_{ST}(b, a) = \sum_{c} S_{H}(b, c) \cdot T_{L}(c, a)$ under update $\delta S_{H} = \{(\beta, \gamma) \mapsto m\}$



Maintain $V_{ST}(b, a) = \sum_{c} S_{H}(b, c) \cdot T_{L}(c, a)$ under update $\delta S_{H} = \{(\beta, \gamma) \mapsto m\}$



Update time: $\mathcal{O}(N^{\varepsilon})$ to iterate over *a*-values paired with γ from T_L

 $\text{Maintain } V_{ST}(b,a) = \sum_{c} S_{H}(b,c) \cdot T_{L}(c,a) \text{ under update } \delta T_{L} = \{(\gamma,\alpha) \mapsto m\}$



Maintain $V_{ST}(b, a) = \sum_{c} S_{H}(b, c) \cdot T_{L}(c, a)$ under update $\delta T_{L} = \{(\gamma, \alpha) \mapsto m\}$



Maintain $V_{ST}(b, a) = \sum_{c} S_{H}(b, c) \cdot T_{L}(c, a)$ under update $\delta T_{L} = \{(\gamma, \alpha) \mapsto m\}$



Update time: $\mathcal{O}(N^{1-\varepsilon})$ to iterate over *b*-values paired with γ from S_H

Maintenance of Auxiliary Views: Summary

$$V_{RS}(a,c) = \sum_{b} R_{H}(a,b) \cdot S_{L}(b,c)$$

$$V_{ST}(b,a) = \sum_{c} S_H(b,c) \cdot T_L(c,a)$$

$$V_{TR}(a,c) = \sum_{a} T_H(c,a) \cdot R_L(a,b)$$

Maintenance Complexity

Time: $\mathcal{O}(N^{\max\{\varepsilon,1-\varepsilon\}})$

Updates can change frequencies of values & heavy/light threshold

Rebalancing Partitions

Updates can change the frequencies of values in the factor parts



Minor Rebalancing

• Transfer $\mathcal{O}(N^{\varepsilon})$ tuples from one to the other part of the same factor

Time complexity:
$$\mathcal{O}(N^{\varepsilon + \max{\{\varepsilon, 1-\varepsilon\}}})$$

Rebalancing Partitions

Updates can change the heavy-light threshold!



Major Rebalancing

Recompute partitions and views from scratch

Amortization of Rebalancing Times

Both forms of rebalancing require superlinear time

Amortization of Rebalancing Times

- Both forms of rebalancing require superlinear time
- The rebalancing times amortize over sequences of updates
 - Amortized minor rebalancing time: $\mathcal{O}(N^{\max{\{\varepsilon, 1-\varepsilon\}}})$

• Amortized major rebalancing time: $\mathcal{O}(N^{\max{\{\varepsilon,1-\varepsilon\}}})$

$$\begin{array}{c} & & & \\ & & \\ & & \\ & & \\ & & \\ & & \\ & & \\ & & \\$$

Amortization of Rebalancing Times

- Both forms of rebalancing require superlinear time
- The rebalancing times amortize over sequences of updates
 - Amortized minor rebalancing time: *O*(*N*^{max {ε,1-ε}})

• Amortized major rebalancing time: $\mathcal{O}(N^{\max{\{\varepsilon, 1-\varepsilon\}}})$

• Overall amortized rebalancing time: $\mathcal{O}(N^{\max{\{\varepsilon, 1-\varepsilon\}}})$

 $\Omega(N)$

References i

[Algorithmica 1997] Noga Alon, Raphael Yuster, and Uri Zwick. *Finding and counting given length cycles.*

[SODA 2002] Ziv Bar-Yossef, Ravi Kumar, and D Sivakumar. Reductions in Streaming Algorithms, with an Application to Counting Triangles in Graphs.

[COCOON 2005] Hossein Jowhari and Mohammad Ghodsi. *New Streaming Algorithms for Counting Triangles in Graphs.*

[PODS 2006] Luciana S Buriol, Gereon Frahling, Stefano Leonardi, Alberto Marchetti-Spaccamela, and Christian Sohler. *Counting Triangles in Data Streams.*

[Found. & Trends DB 2012] Rada Chirkova and Jun Yang. Materialized Views.

[SIGMOD R. 2013] Hung Q Ngo, Christopher Ré, and Atri Rudra. Skew strikes back: new developments in the theory of join algorithms.

References ii

[ICDT 2014] Todd L. Veldhuizen. Leapfrog Triejoin: A Simple, Worst-Case Optimal Join Algorithm.

[VLDB J. 2014] Christoph Koch, Yanif Ahmad, Oliver Kennedy, Milos Nikolic, Andres Nötzli, Daniel Lupei, and Amir Shaikhha. *DBToaster: Higher-Order Delta Processing for Dynamic, Frequently Fresh Views.*

[FOCS 2015] Talya Eden, Amit Levi, Dana Ron, and C. Seshadhri. Approximately Counting Triangles in Sublinear Time.

[STOC 2015] Monika Henzinger, Sebastian Krinninger, Danupon Nanongkai, and Thatchaphol Saranurak. Unifying and Strengthening Hardness for Dynamic Problems via the Online Matrix-Vector Multiplication Conjecture.

[PODS 2016] Andrew McGregor, Sofya Vorotnikova, and Hoa T Vu. Better Algorithms for Counting Triangles in Data Streams.

[PODS 2017] Christoph Berkholz, Jens Keppeler, and Nicole Schweikardt. *Answering Conjunctive Queries Under Updates.*

References iii

[SIGMOD 2017] Muhammad Idris, Martín Ugarte, and Stijn Vansummeren. The Dynamic Yannakakis Algorithm: Compact and Efficient Query Processing Under Updates.

[Theor. Comput. Sci. 2017] Graham Cormode and Hossein Jowhari. A Second Look at Counting Triangles in Graph Streams (Corrected).

[Found. & Trends DB 2018] Paraschos Koutris, Semih Salihoglu, and Dan Suciu. *Algorithmic Aspects of Parallel Data Processing*.

[ICDT 2018] Christoph Berkholz, Jens Keppeler, and Nicole Schweikardt. Answering UCQs Under Updates and in the Presence of Integrity Constraints.

[ICM 2018] Virginia Vassilevska Williams. On Some Fine-Grained Questions in Algorithms and Complexity.

[SIGMOD 2018] Nikolic, Milos, and Dan Olteanu. Incremental view maintenance with triple lock factorization benefits.

References iv

[ICDT 2019] Ahmet Kara, Milos Nikolic, Hung Q. Ngo, Dan Olteanu, and Haozhe Zhang. *Counting Triangles under Updates in Worst-Case Optimal Time.*

[APOCS 2021] Laxman Dhulipala, Quanquan C. Liu, Julian Shun, Shangdi Yu. *Parallel Batch-Dynamic k-Clique Counting.*

[ICDT 2021] Shangqi Lu, Yufei Tao. Towards Optimal Dynamic Indexes for Approximate (and Exact) Triangle Counting.

Part 3. IVM Optimality

IVM Optimality

Two query classes with known worst-case optimal IVM algorithms

- Q-hierarchical queries
 - The only queries that admit constant update time and delay
- δ_1 -hierarchical queries
 - They admit $\mathcal{O}(\sqrt{N})$ update time and delay

Query examples shown in Part 2

Part 3. IVM Optimality: Constant Update Time & Delay

Hierarchical Queries

A query is hierarchical if for any two variables, their sets of atoms in the query are either disjoint or one is contained in the other [VLDB 2004]

hierarchical

$$Q(b,d) = \sum_{a,c,e,f} R(a,b,d) \cdot S(a,b) \cdot T(a,c,e) \cdot U(a,c,f)$$



Hierarchical Queries

A query is hierarchical if for any two variables, their sets of atoms in the query are either disjoint or one is contained in the other [VLDB 2004]

hierarchical

$$Q(b,d) = \sum_{a,c,e,f} R(a,b,d) \cdot S(a,b) \cdot T(a,c,e) \cdot U(a,c,f)$$



not hierarchical

 $Q(a,b) = R(a) \cdot \frac{S}{S}(a,b) \cdot T(b)$



*Q***-Hierarchical Queries**

A query is *q*-hierarchical if it is hierarchical and the free variables dominate the bound variables [PODS 2017]

$$q\text{-hierarchical}$$

$$Q(a, b, c) = \sum_{d, e, f} R(a, b, d) \cdot \frac{S}{S}(a, b) \cdot \frac{S}{T}(a, c, e) \cdot U(a, c, f)$$



Q-Hierarchical Queries

A query is *q*-hierarchical if it is hierarchical and the free variables dominate the bound variables [PODS 2017]

$$q\text{-hierarchical}$$
$$Q(a, b, c) = \sum_{d, e, f} R(a, b, d) \cdot \frac{S(a, b)}{T(a, c, e)} \cdot U(a, c, f)$$



hierarchical but not q-hierarchical $Q(a) = \sum_{b} S(a, b) \cdot T(b)$



Dichotomy for *Q*-Hierarchical Queries

Let Q be any conjunctive query without self-joins and D a database.

- If Q is q-hierarchical, then the query answer admits O(1) single-tuple updates and enumeration delay.
- If Q is not q-hierarchical, then there is no algorithm with
 O(|D|^{1/2-γ}) update time and enumeration delay for any γ > 0,
 unless the OMv conjecture fails.

[PODS 2017]

Queries under Functional Dependencies

Rewriting queries under functional dependencies [ICDE 2009]

- Given: Query Q and set Σ of functional dependencies
- Replace the set of variables of each atom in Q by its closure under Σ called Σ-reduct

Under $\Sigma = \{x \rightarrow y, y \rightarrow z\}$, the closure of $\{x\}$ is $\{x, y, z\}$

 If the Σ-reduct is *q*-hierarchical, then *Q* admits constant update time and enumeration delay [VLDB J 2023]

Maintenance of Q-Hierarchical Queries

How to achieve constant update time and enumeration delay? Recipe: [PODS 2017]

- Construct a factorized representation of the query answer [ICDT 2012]
- Such factorizations admit constant-delay enumeration
- Apply updates directly on the factorization

- F-IVM system [https://github.com/fdbresearch/FIVM] [SIGMOD 2018]
 - Factorize the query answer as a tree of views
 - Materialize the views to speed up updates and enumeration

Example: Query Rewriting

$$Q(w, x, y, z) = R(w, x) \cdot S(x, y) \cdot T(y, z)$$

Assume the functional dependencies: $X \rightarrow Y$ and $Y \rightarrow Z$

Q is not q-hierarchical, but its rewriting under FDs is:

$$Q'(w, x, y, z) = R'(w, x, y, z) \cdot S'(x, y, z) \cdot T'(y, z)$$

Example: Variable Order

$$Q'(w, x, y, z) = R'(w, x, y, z) \cdot S'(x, y, z) \cdot T'(y, z)$$
Top-down construction of variable order for Q' :
$$Z \text{ and } Y \text{ are first as they dominate } X \text{ and } W$$

$$Then X, which dominates W$$

$$Finally W$$
We use this variable order also for Q

Example: View Tree



View tree construction:

- Place factors at leaves
- Create parent view to join children

$$V'_Z(z, y) = T(y, z) \cdot V_X(y)$$
$$V'_X(y, x) = S(x, y) \cdot V_W(x)$$

 Aggregate away variables not needed for further joins

$$V_Z() = \sum_{z} V_Y(z)$$
$$V_Y(z) = \sum_{y} V'_Z(y, z)$$
$$V_X(y) = \sum_{x} V'_X(x, y)$$
$$V_W(x) = \sum_{w} R'(x, w)$$








































Example: Enumeration of Query Answers



Enumeration for Q(z, y, x, w) with constant delay

- Top-down in the view tree
- Views calibrated for variables underneath
- Guaranteed to get matching tuples in views below

Example: Enumeration of Query Answers



Part 3. IVM Optimality: $\mathcal{O}(\sqrt{N})$ Update Time & Delay

$$Q(a) = \sum_{b} R(a, b) \cdot S(b)$$





 \log_N preprocessing time

Known approach: Eager update, quick enumeration

- Preprocessing: Materialize the result.
- Upon update: Maintain the materialized result.
- Enumeration: Enumerate from materialized result.



log_N preprocessing time

Known approach: Lazy update, heavy enumeration

- Preprocessing: Eliminate dangling tuples
- Upon update: Update only input factors
- Enumeration: Eliminate dangling tuples and enumerate from R



Yet, there is an algorithm that admits sub-linear update time and sub-linear enumeration delay





Weak Pareto optimality

Factor Partitioning

 $Q(a) = \sum_{b} R(a, b) \cdot S(b)$

Partition R based on the values b into

- a light part $R^{L} = \{(a, b) \in R \mid |\sigma_{B=b}R| < N^{\varepsilon}\}$
- a heavy part $R^H = R R^L$



Factor Partitioning

 $Q(a) = \sum_{b} R(a, b) \cdot S(b)$

Partition R based on the values b into

- a light part $R^L = \{(a, b) \in R \mid |\sigma_{B=b}R| < N^{\varepsilon}\}$
- a heavy part $R^H = R R^L$



$$egin{aligned} Q(a) &= Q_L(a) + Q_H(a) \ Q_L(a) &= \sum_b R^L(a,b) \cdot S(b) \ Q_H(a) &= \sum_b R^H(a,b) \cdot S(b) \end{aligned}$$

Light Case

 $Q_L(a) = \sum_b R^L(a, b) \cdot S(b)$

Materialize the result

Light Case

 $Q_L(a) = \sum_b R^L(a, b) \cdot S(b)$ Materialize the result $Q_L(a) = \sum_b R^L(a, b) \cdot S(b)$ ai aj S(b) $R^{L}(a, b)$ b_1' $a_1 b_1$ b'_m a_n b_n

Preprocessing in the Light Case

 $Q_L(a) = \sum_b R^L(a, b) \cdot S(b)$



• Q_L can be computed in time $\mathcal{O}(N)$

Enumeration in the Light Case



Q_L allows constant-time lookups and constant-delay enumeration









• Updates to R^L : $\mathcal{O}(1)$



• Updates to R^L : $\mathcal{O}(1)$



• Updates to R^L : $\mathcal{O}(1)$
Updates in the Light Case



Updates in the Light Case



- Updates to R^L: O(1)
- Updates to S: O(N^ε)

Heavy Case

$$Q_{H}(a) = \sum_{b} R^{H}(a,b) \cdot S(b)$$

Materialize the b values in the join result

Heavy Case





Preprocessing in the Heavy Case





■ V_{RS} can be computed in time $\mathcal{O}(N^{1-\varepsilon})$ and has at most $N^{1-\varepsilon}$ values

Enumeration in the Heavy Case

$$Q_H(a) = \sum_b R^H(a,b) \cdot S(b)$$



• V_{RS} contains at most $N^{1-\varepsilon}$ values b

For each value *b* in V_{RS} , the values *a* in R^H paired with *b* admit constant enumeration delay

Enumeration of Distinct Tuples from Union

• $V_{RS}(b)$ contains at most $N^{1-\varepsilon}$ values

- For each value *b* in *V_{RS}*, the values *a* in *R^H* paired with *b* admit constant enumeration delay
- Yet: For two distinct b_1 and b_2 , the sets of values a in $R^H(a, b_1)$ and $R^H(a, b_2)$ may not be disjoint

 \implies Enumerating all the values *a* in $R^H(a, b_1)$ and $R^H(a, b_2)$ can lead to duplicates

Enumeration of Distinct Tuples from Union

• $V_{RS}(b)$ contains at most $N^{1-\varepsilon}$ values

- For each value *b* in *V_{RS}*, the values *a* in *R^H* paired with *b* admit constant enumeration delay
- Yet: For two distinct b_1 and b_2 , the sets of values a in $R^H(a, b_1)$ and $R^H(a, b_2)$ may not be disjoint

 \implies Enumerating all the values *a* in $R^H(a, b_1)$ and $R^H(a, b_2)$ can lead to duplicates

Union Algorithm [CSL 2011] • The distinct values *a* can be enumerated with $\mathcal{O}(N^{1-\varepsilon})$ delay

Enumeration of the distinct tuples in the union of two sets

- Both sets allow lookup time ℓ and enumeration delay d
- > The union of the sets can be enumerated with $\mathcal{O}(\ell+d)$ delay

$$S_1$$
 S_2 $S_1 \cup S_2$
 $a_3 a_4 a_1 a_2$ EOF $a_5 a_6 a_2 a_4$ EOF

Enumeration of the distinct tuples in the union of two sets

- Both sets allow lookup time ℓ and enumeration delay d
- \Rightarrow The union of the sets can be enumerated with $\mathcal{O}(\ell+d)$ delay

Enumeration of the distinct tuples in the union of two sets

Both sets allow lookup time ℓ and enumeration delay d

Enumeration of the distinct tuples in the union of two sets

Both sets allow lookup time ℓ and enumeration delay d

Enumeration of the distinct tuples in the union of two sets

Both sets allow lookup time ℓ and enumeration delay d

Enumeration of the distinct tuples in the union of two sets

Both sets allow lookup time ℓ and enumeration delay d

Enumeration of the distinct tuples in the union of two sets

Both sets allow lookup time ℓ and enumeration delay d

Enumeration of the distinct tuples in the union of two sets

Both sets allow lookup time ℓ and enumeration delay d

Enumeration of the distinct tuples in the union of two sets

- **Both sets allow lookup time** ℓ and enumeration delay *d*
- > The union of the sets can be enumerated with $\mathcal{O}(\ell+d)$ delay

Generalization: Enumeration from the union of n sets

- Each set allows lookup time ℓ and enumeration delay d
- The union of the sets can be enumerated with $\mathcal{O}(n(\ell + d))$ delay



















- Updates to R^H : $\mathcal{O}(1)$
- Updates to S: O(1)

Summing Up

 $Q(a) = R(a, b) \cdot S(b)$

Preprocessing Time

light case	heavy case	overall
$\mathcal{O}(N)$	$\mathcal{O}(N^{1-arepsilon})$	$\mathcal{O}(N)$

Enumeration Delay

Update Time

light case	heavy case	overall
$\mathcal{O}(N^{\varepsilon})$	$\mathcal{O}(1)$	$\mathcal{O}(N^{\varepsilon})$

Are there more queries with the same weak Pareto optimality as our previous example?

δ_1 -Hierarchical Queries

- For any bound variable X and any atom α of X, there is at most one other atom β so that all free variables dominated by X are covered by α and β together
- The query is hierarchical and not *q*-hierarchical



δ_1 -Hierarchical Queries

- For any bound variable X and any atom α of X, there is at most one other atom β so that all free variables dominated by X are covered by α and β together
- The query is hierarchical and not *q*-hierarchical





Optimality for δ_1 -Hierarchical Queries

• For any δ_1 -hierarchical query, there is no algorithm that admits preprocessing time update time enumeration delay arbitrary $\mathcal{O}(N^{1/2-\gamma}) = \mathcal{O}(N^{1/2-\gamma})$ for any $\gamma > 0$, unless the OMv Conjecture (*) fails

(*) Online Matrix-Vector Multiplication cannot be solved in sub-cubic time log_Mupdate time



Optimality for δ_1 -Hierarchical Queries



(*) Online Matrix-Vector Multiplication cannot be solved in sub-cubic time



Optimality for δ_1 -Hierarchical Queries



log_N preprocessing time

References i

[VLDB 2004] Nilesh N. Dalvi, Dan Suciu. *Efficient Query Evaluation on Probabilistic Databases.*

[ICDE 2009] Dan Olteanu, Jiewen Huang, Christoph Koch. SPROUT: Lazy vs. Eager Query Plans for Tuple-Independent Probabilistic Databases.

[CSL 2011] Arnaud Durand, Yann Strozecki. *Enumeration Complexity* of Logical Query Problems with Second-order Variables. CSL 2011

[ICDT 2012] Dan Olteanu, Jakub Zavodny. Factorised representations of query results: size bounds and readability.

[PODS 2017] Christoph Berkholz, Jens Keppeler, Nicole Schweikardt. Answering Conjunctive Queries under Updates.

References ii

[SIGMOD 2018] Milos Nikolic, Dan Olteanu. Incremental View Maintenance with Triple Lock Factorization Benefits.

[ICDT 2023] Ahmet Kara, Milos Nikolic, Dan Olteanu, Haozhe Zhang. Conjunctive Queries with Free Access Patterns Under Updates.

[VLDBJ 2023] Ahmet Kara, Milos Nikolic, Dan Olteanu, Haozhe Zhang. *F-IVM: Analytics over Relational Databases under Updates.* (To appear)

[LMCS 2023] Ahmet Kara, Milos Nikolic, Dan Olteanu, Haozhe Zhang. Trade-offs in Static and Dynamic Evaluation of Hierarchical Queries.

Thank You!