HyperBlocker: Accelerating Rule-Based Blocking in Entity Resolution Using GPUs

Research track Information Integration and Data Quality III

Xiaoke Zhu
(Beihang University → Shenzhen Institute of Computing Sciences)

Co-authors: Min Xie (Shenzhen Institute of Computing Sciences),
Ting Deng (Beihang University), Qi Zhang (Meta Platforms)

What is Entity Resolution?

Problem of identifying and linking/grouping different manifestations of the same real world object.

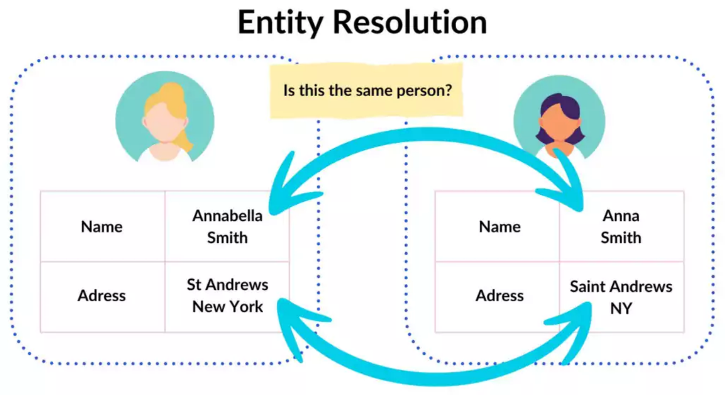

What is Entity Resolution?

Problem of identifying and linking/grouping different manifestations of the same real world object.

What Caused the ER Problem?

- Name/Attribute ambiguity
- Errors due to data entry
- Data integration
- ...

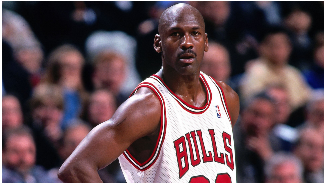

Figure 1: Basketball Player

Figure 2: Professor at UC, Berkeley

What is Entity Resolution?

Problem of identifying and linking/grouping different manifestations of the same real world object.

What Caused the ER Problem?

- Name/Attribute ambiguity
- Errors due to data entry
- Data integration
- ...

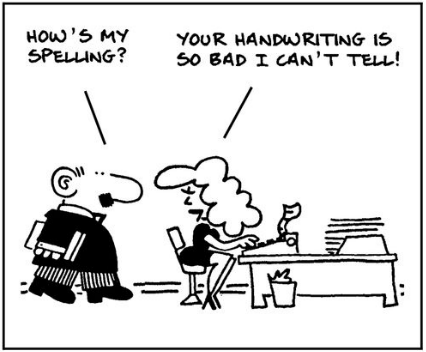

What is Entity Resolution?

Problem of identifying and linking/grouping different manifestations of the same real world object.

What Caused the ER Problem?

- Name/Attribute ambiguity
- Errors due to data entry
- Data integration
- ...

BigData ER Challenges

ER suffers from a quadratic time complexity, $O(n^2)$

1000 business listings each from 1,000 different cities across the world

- 1 trillion comparisons
- 11 days (if each comparison is 1 us)

Candidate selection step

- Blocking (e.g., hashing/clustering): groups similar entities into blocks.
- Filtering (e.g., similarity join): removes dissimilar entity pairs while potentially keeping some non-matching pairs.

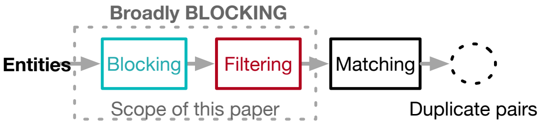

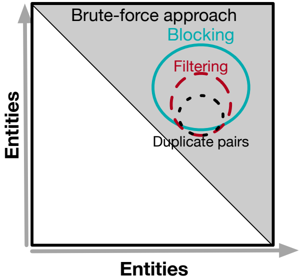

Rule-based BLOCKING

eid	pno	pname	price	sname	description	color	saddress
e_1	t_1	Apple Mac Air	$909	Comp. World	Apple MacBook Air (13-inch, 8GB RAM, 256GB SSD)	Ashen	9 Barton Grove, McCulloughmouth
e_2	t_2	ThinkPad	-	Smith's Tech	ThinkPad E15, 15.6-inch full HD IPS display, Intel Core i5-1235U processor, (16GB) RAM | 512GB PCIe SSD)	Gray	Seg Plaza, Hua qiang North Road
e_2	t_3	ThinkPad	$849	Smith's Tech	Lenovo E15 Business ThinkPad, 15.6-inch full HD IPS display, 12 generation Intel Core i5, 16GB RAM, 512GB SSD	Gray	Seg Plaza, Hua qiang North Road
e_1	t_4	MacBook Air	$909	Comp. World	Apple 2022 MacBook Air M2 chip 13-inch, 8 GB RAM, 256 GB SSD storage gray	Gray	-
e_1	t_5	MacBook Air	$909	Comp. World	-	Gray	Barton Grove, McCulloughmouth

- $\varphi_1 : t.\text{color} = s.\text{color} \wedge t.\text{price} = s.\text{price} \wedge t.\text{sname} = s.\text{sname} \wedge t.\text{pname} \approx_{\text{ED}} s.\text{pname} \rightarrow t.\text{eid} = s.\text{eid}$

- $\varphi_2 : t.\text{sname} = s.\text{sname} \wedge t.\text{description} \approx_{\text{JD}} s.\text{description} \rightarrow t.\text{eid} = s.\text{eid}$

- $\varphi_3 : t.\text{saddress} \approx_{\text{ED}} s.\text{saddress} \wedge t.\text{description} \approx_{\text{JD}} s.\text{description} \rightarrow t.\text{eid} = s.\text{eid}$

Rule-based BLOCKING

eid	pno	pname	price	sname	description	color	saddress
e_1	t_1	Apple Mac Air	\$909	Comp. World	Apple MacBook Air (13-inch, 8GB RAM, 256GB SSD)	Ashen	9 Barton Grove, McCulloughmouth
e_2	t_2	ThinkPad	-	Smith's Tech	ThinkPad E15, 15.6-inch full HD IPS display, Intel Core i5-1235U processor, (16GB) RAM | 512GB PCIe SSD)	Gray	Seg Plaza, Hua qiang North Road
e_2	t_3	ThinkPad	\$849	Smith's Tech	Lenovo E15 Business ThinkPad, 15.6-inch full HD IPS display, 12 generation Intel Core i5, 16GB RAM, 512GB SSD	Gray	Seg Plaza, Hua qiang North Road
e_1	t_4	MacBook Air	\$909	Comp. World	Apple 2022 MacBook Air M2 chip 13-inch, 8 GB RAM, 256 GB SSD storage gray	Gray	-
e_1	t_5	MacBook Air	\$909	Comp. World	-	Gray	Barton Grove, McCulloughmouth

- $\varphi_1 : t.\text{color} = s.\text{color} \wedge t.\text{price} = s.\text{price} \wedge t.\text{sname} = s.\text{sname} \wedge t.\text{pname} \approx_{\text{ED}} s.\text{pname} \rightarrow t.\text{eid} = s.\text{eid}$
  - t_4 is a potential match for t_5
- $\varphi_2 : t.\text{sname} = s.\text{sname} \wedge t.\text{description} \approx_{\text{JD}} s.\text{description} \rightarrow t.\text{eid} = s.\text{eid}$
  - t_1 is a potential match for t_4
- $\varphi_3 : t.\text{saddress} \approx_{\text{ED}} s.\text{saddress} \wedge t.\text{description} \approx_{\text{JD}} s.\text{description} \rightarrow t.\text{eid} = s.\text{eid}$
  - t_2 is a potential match for t_3

Entity Resolution on 10+ Million Tuples

Dataset: TFACC: 10 million tuples, each described by 16 attributes

Enviroments:

A cluster of 30 HPC servers, powered with 2.40GHz Intel Xeon Gold CPU

1 × Nvidia Tesla V100 GPU

Summary:

Solution	Execution Time	Characteristics
Dedoop	>3 hours	Distributed implementation (30-nodes)
DisDedup	>3 hours	Distributed implementation (30-nodes)
SparkER	>3 hours	Distributed implementation (30-nodes)
Faiss + Ditto	Out of Memory	Simlarity Join + Matcher (GPU)
DeepBlocker + Ditto	Out of Memory	Learned Blocker + Matcher (GPU, SOTA)

Existing methods cannot process 10M+ tuples, even on a 30-node cluster

Our Contributions

- The data/rule-aware execution plan for BLOCKING
  - Tree representation of execution plan
  - Cost models for User Dedefie Function(UDFs)
- The GPU hardware-aware parallelism for BLOCKING
  - Task stealing + parallel sliding window
- Multi-GPUs collaboration

VS Learned Blocker
(SOTA in Accuracy)

Accuracy
≈

VS Learned Blocker
(SOTA in Accuracy)

Space
≈ 75%

VS 30-Node Distributed Blocker
& GPU-Accelerated Blocker
(SOTA in Efficiency)

Speed
80% - 9.1X

Observation 1

Not all predicates entail equivalent satisfaction or computational costs

eid	pno	pname	price	sname	description	color	saddress
e_1	t_1	Apple Mac Air	$909	Comp. World	Apple MacBook Air (13-inch, 8GB RAM, 256GB SSD)	Ashen	9 Barton Grove, McCulloughmouth
e_2	t_2	ThinkPad	-	Smith's Tech	ThinkPad E15, 15.6-inch full HD IPS display, Intel Core i5-1235U processor, (16GB) RAM | 512GB PCIe SSD)	Gray	Seg Plaza, Hua qiang North Road
e_2	t_3	ThinkPad	$849	Smith's Tech	Lenovo E15 Business ThinkPad, 15.6-inch full HD IPS display, 12 generation Intel Core i5, 16GB RAM, 512GB SSD	Gray	Seg Plaza, Hua qiang North Road
e_1	t_4	MacBook Air	$909	Comp. World	Apple 2022 MacBook Air M2 chip 13-inch, 8 GB RAM, 256 GB SSD storage gray	Gray	-
e_1	t_5	MacBook Air	$909	Comp. World	-	Gray	Barton Grove, McCulloughmouth

More discriminative in the "description" attribute

"description": Long-text field containing discriminative features for entity resolution

Observation 1

Not all predicates entail equivalent satisfaction or computational costs

eid	pno	pname	price	sname	description	color	saddress
e_1	t_1	Apple Mac Air	$909	Comp. World	Apple MacBook Air (13-inch, 8GB RAM, 256GB SSD)	Ashen	9 Barton Grove, McCulloughmouth
e_2	t_2	ThinkPad	-	Smith's Tech	ThinkPad E15, 15.6-inch full HD IPS display, Intel Core i5-1235U processor, (16GB) RAM	512GB PCIe SSD)	Gray	Seg Plaza, Hua qiang North Road
e_2	t_3	ThinkPad	$849	Smith's Tech	Lenovo E15 Business ThinkPad, 15.6-inch full HD IPS display, 12 generation Intel Core i5, 16GB RAM, 512GB SSD	Gray	Seg Plaza, Hua qiang North Road
e_1	t_4	MacBook Air	$909	Comp. World	Apple 2022 MacBook Air M2 chip 13-inch, 8 GB RAM, 256 GB SSD storage gray	Gray	-
e_1	t_5	MacBook Air	$909	Comp. World	-	Gray	Barton Grove, McCulloughmouth

Higher computational efficiency on the "color" attribute.

"color": The computation of short-text similarity necessitates only few operations

Solution: Query optimization for User Dedefie Function

Quantifies the likelihood that predicate p is satisfied on attribute a of dataset D

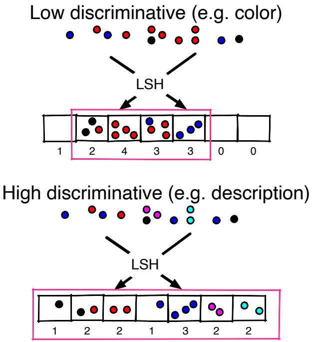

Quantifies the efficiency of computing p on attribute a of dataset D

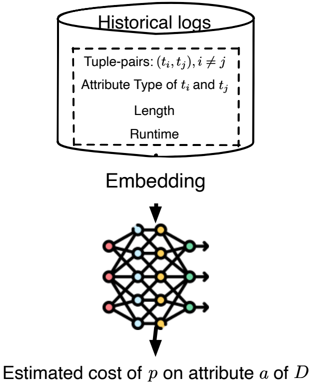

A cost model balancing computational cost and selectivity of predicates was proposed for execution plan optimization

Solution: Query optimization for User Dedefie Function

Quantifies the likelihood that predicate p is satisfied when applied to attribute a

- Evaluating selectivity of p on attribute a of D

$$\mathrm{sp}(p,D) = \mathrm{Norm}\left(\sqrt{\frac{1}{k}\sum_{i}^{k}\left(b_i - \frac{|D|}{k}\right)^2}\right)$$

Quantifies the efficiency of computing p on attribute a

- Cost of predicate p on attribute a of D

$$\mathrm{cost}_a(p,D) = \sum_{(t_1,t_2)\in D\times D} T_p(t_1,t_2)$$

- Estimated cost

$$\hat{\mathrm{cost}}_a(p,D) = \mathrm{Norm}\left(\sum_{(t_1,t_2)\in D\times D} \mathcal{N}(p,t_1,t_2)\right)$$

cost-effectiveness mode

$$\frac{1 - sp(p,D)}{\hat{\mathrm{cost}}_a(p,D)}$$

Observation 2 & Solution

Redundant Computations

$\varphi_1 : t.\text{color} = s.\text{color} \wedge t.\text{price} = s.\text{price} \wedge t.\text{sname} = s.\text{sname} \wedge t.\text{pname} \approx_{\text{ED}} s.\text{pname} \rightarrow t.\text{eid} = s.\text{eid}$

$\varphi_2 : t.\text{sname} = s.\text{sname} \wedge t.\text{description} \approx_{\text{JD}} s.\text{description} \rightarrow t.\text{eid} = s.\text{eid}$

$\varphi_3 : t.\text{saddress} \approx_{\text{ED}} s.\text{saddress} \wedge t.\text{description} \approx_{\text{JD}} s.\text{description} \rightarrow t.\text{eid} = s.\text{eid}$

Solution: Execution Tree Representation

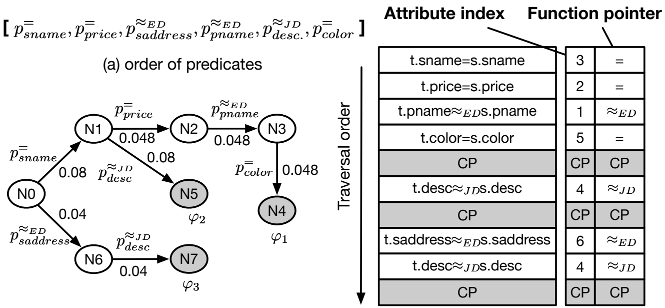

(b) Execution Tree (c) Sequential execution path

Observation 3

GPU thread workload imbalance occurs during rule deduction on *D*

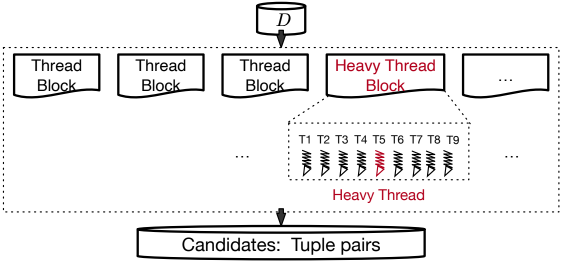

Unique Challenges for Rule-Based Blocking

- Early termination for a thread (first rule suffice) vs. full traversal for another thread (no early witness).
- Even for the same predicate, the evaluation time on different tuples is different (e.g., long vs. short text)

Solutions

Parallel Sliding Window

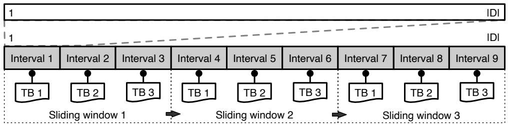

Task Stealing

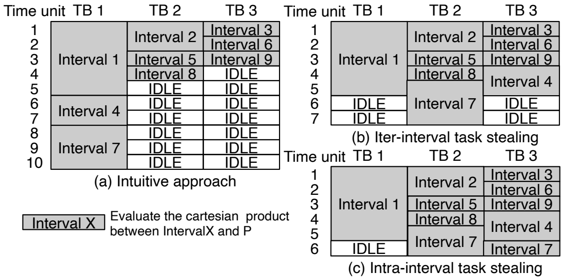

Experimental

Dataset

Dataset	Domain	#Tuples	Max #Pairs	#GT Pairs	#Attrs	#Rules	#Partitions
Fodors-Zagat	restaurant	866	1.8×10^4	112	6	1	1
DBLP-ACM	citation	4591	6.0×10^6	2294	4	10	8
DBLP-Scholar	citation	66881	1.7×10^8	5348	4	10	8
IMDB	movie	1.5M	8.1×10^{10}	0.2M+	6	10	128
Songs	music	0.5M	2.7×10^{11}	1.2M	8	10	128
NCV	vote	2M	1.0×10^{12}	0.5M+	5	10	512
TPCH	synthetic	4M	1.6×10^{13}	#	8	30	512
TFACC	traffic	10M	1.0×10^{14}	#	16	50	1024
$\text{TFACC}_{\text{large}}$	traffic	36M	1.3×10^{15}	#	16	50	1024

Baselines

- Distributed on 30 nodes: SparkER[EDBT'19], DisDedup[VLDB'16], Dedoop[VLDB'12]
- GPU-based: DeepBlocker[VLDB'21], GPUDet[HPIFuture'13]
- Matcher: Ditto[VLDB'20]
- Variants: $\text{HyperBlocker}_{\text{noEPG}}$: without EPG, $\text{HyperBlocker}_{\text{HO}}$: disables all hardware optimizations, $\text{HyperBlocker}_{\text{Ditto}}$: HyperBlocker as blocker and Ditto as the matcher

HyperBlocker is Highly Efficient and Scalable

Achieves speeds 6.818× faster than baselines, with comparable accuracy

Method	DBLP-ACM		IMDB		Songs		NCV	
	F1-score	Time (s)	F1-score	Time (s)	F1-score	Time (s)	F1-score	Time (s)
SparkER	0.77 (-0.17)	11.0 (13.8×)	0.31 (-0.65)	242.9 (6.8×)	0.08 (-0.72)	203.4 (15.2×)	0.26 (-0.66)	229.3 (49.8×)
GPUDet	0.92 (-0.02)	20.1 (25.1×)	0.94 (-0.02)	323.8 (9.1×)	0.80 (+0)	404.8 (30.2×)	0.90 (-0.02)	1252.6 (272.3×)
DeepBlocker	0.98 (+0.04)	8.3 (10.4×)	/	>3h	/	>3h	/	>3h
$\text{HyperBlocker}_{\text{noEPG}}$	0.94 (+0)	9.9 (12.4×)	/	>3h	0.80 (+0)	1904.1 (142×)	0.92 (+0)	2408.6 (523.6×)
$\text{HyperBlocker}_{\text{noHO}}$	0.94 (+0)	9.5 (11.9×)	0.96 (+0)	472.6 (13.2×)	0.80 (+0)	45.0 (3.4×)	0.92 (+0)	35.9 (7.8×)
HyperBlocker	0.94	0.8	0.96	35.7	0.80	13.4	0.92	4.6
Dedoop	0.90 (-0.08)	59.4 (9.4×)	0.67 (-0.29)	534.0 (15.0×)	0.80 (-0.08)	7643.4 (6.5×)	/	>3h
DisDedup	0.45 (-0.53)	94.0 (14.9×)	0.67 (-0.29)	644.0 (18.0×)	0.06 (-0.82)	917.0 (0.8×)	/	>3h
$\text{Ditto}_{\text{top2}}$	0.98 (+0)	9.0 (1.4×)	0.79 (-0.17)	6741.2 (188.8×)	0.88 (+0)	2308.6 (2.0×)	0.97 (+0.03)	381.8 (2.1×)
$\text{DeepBlocker}_{\text{Ditto}}$	0.99 (+0.01)	12.4 (2.0×)	/	>3h	/	>3h	/	>3h
$\text{HyperBlocker}_{\text{Ditto}}$	0.98	6.3	*0.96	*35.7	0.88	1179.0	0.94	180.6

HyperBlocker can scale to large dataset: it process 36M tuples within 3h.

HyperBlocker Reduces Memory Costs

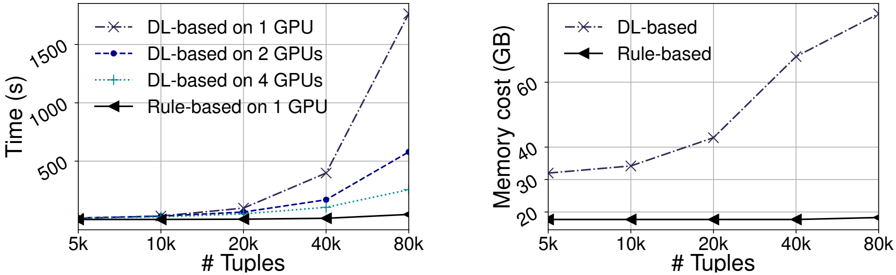

(a) Execution time

(b) Memory cost on a single GPU

Figure 3: DL-based blocking vs. rule-based blocking

HyperBlocker Takeaways

- Improves Performance
- Scale up to 36M tuples
- Benefit beyound Blockfing in Entity Resolution
- Source code: https://github.com/hsiaoko/HyperBlocker

Thanks

Backup Slides

HyperBlocker is Scalable

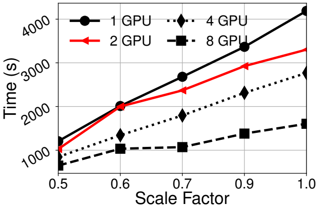

(a) TFACC: Varying #GPUs

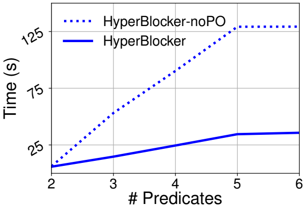

(b) IMDB: Varying $|\varphi|$

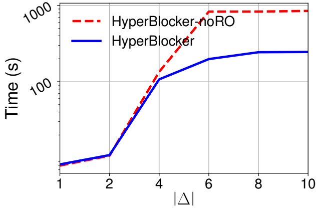

(c) IMDB: Varying $|\Delta|$

EPG VS Existing Query Optimizer

Table 1: Time under different execution plans

Method	Songs	DBLP-ACM	NCV	IMDB
PostgreSQL	40.6 (3.0×)	0.7 (2.3×)	13.5 (2.2×)	13036.0 (449.5×)
TupleX[SIGMOD'21]	14.2 (1.04×)	0.4 (1.3×)	12.4 (2.0×)	36.1 (1.2×)
EPG	13.6	0.3	6.2	29.0