Location-based Matching in Publish/Subscribe Revisited

Mohammad Sadoghi and Hans-Arno Jacobsen

University of Toronto

December 2012

Computational Advertising (A Billion-dollar Industry)

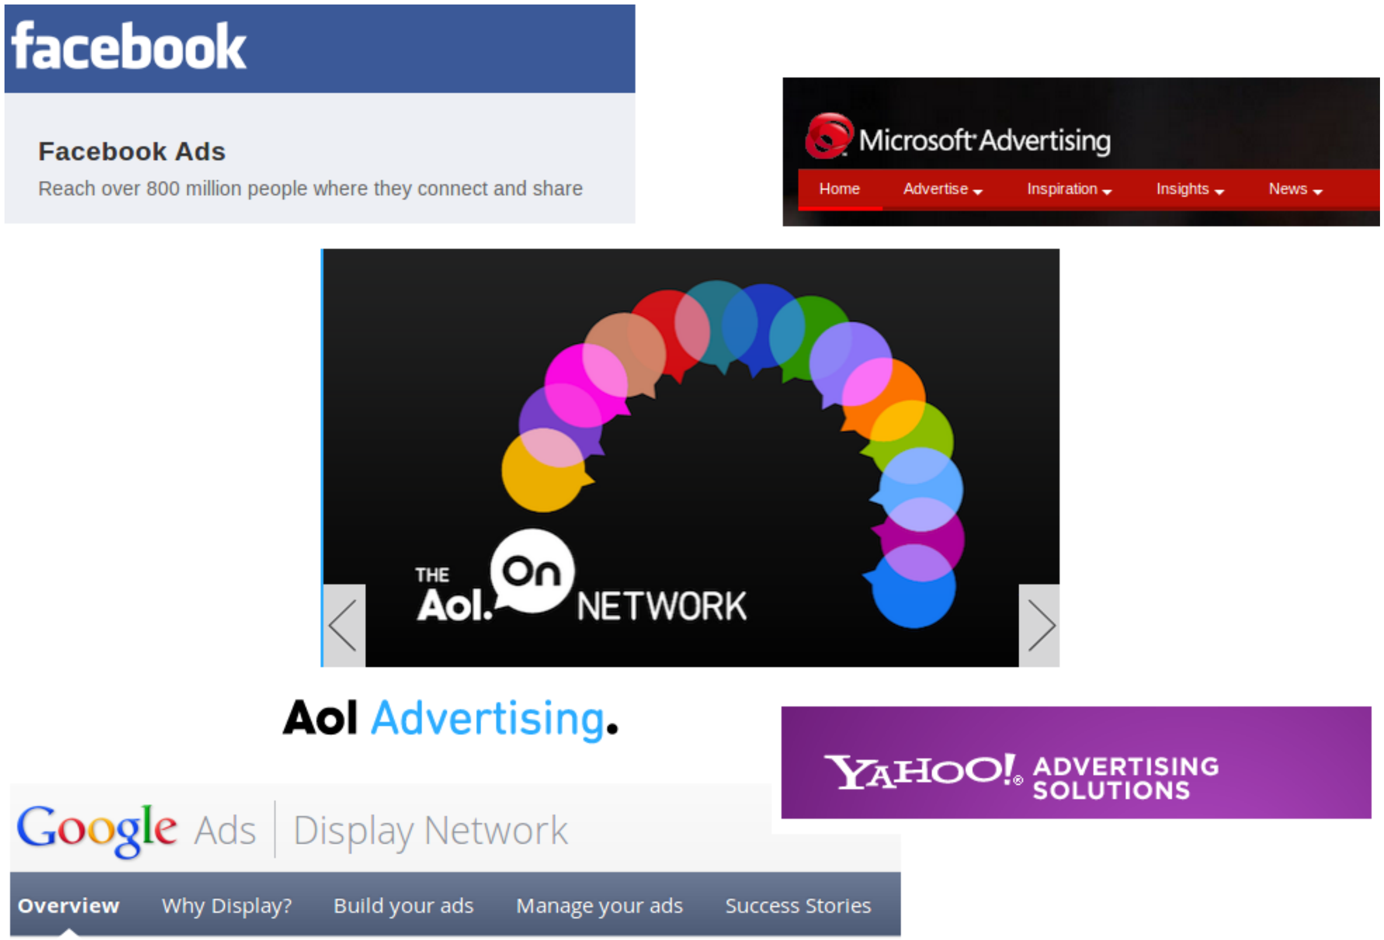

UNIVERSITY OF TORONTO

MIDDLEWARE SYSTEMS RESEARCH GROUP

Computational Advertising (A Billion-dollar Industry)

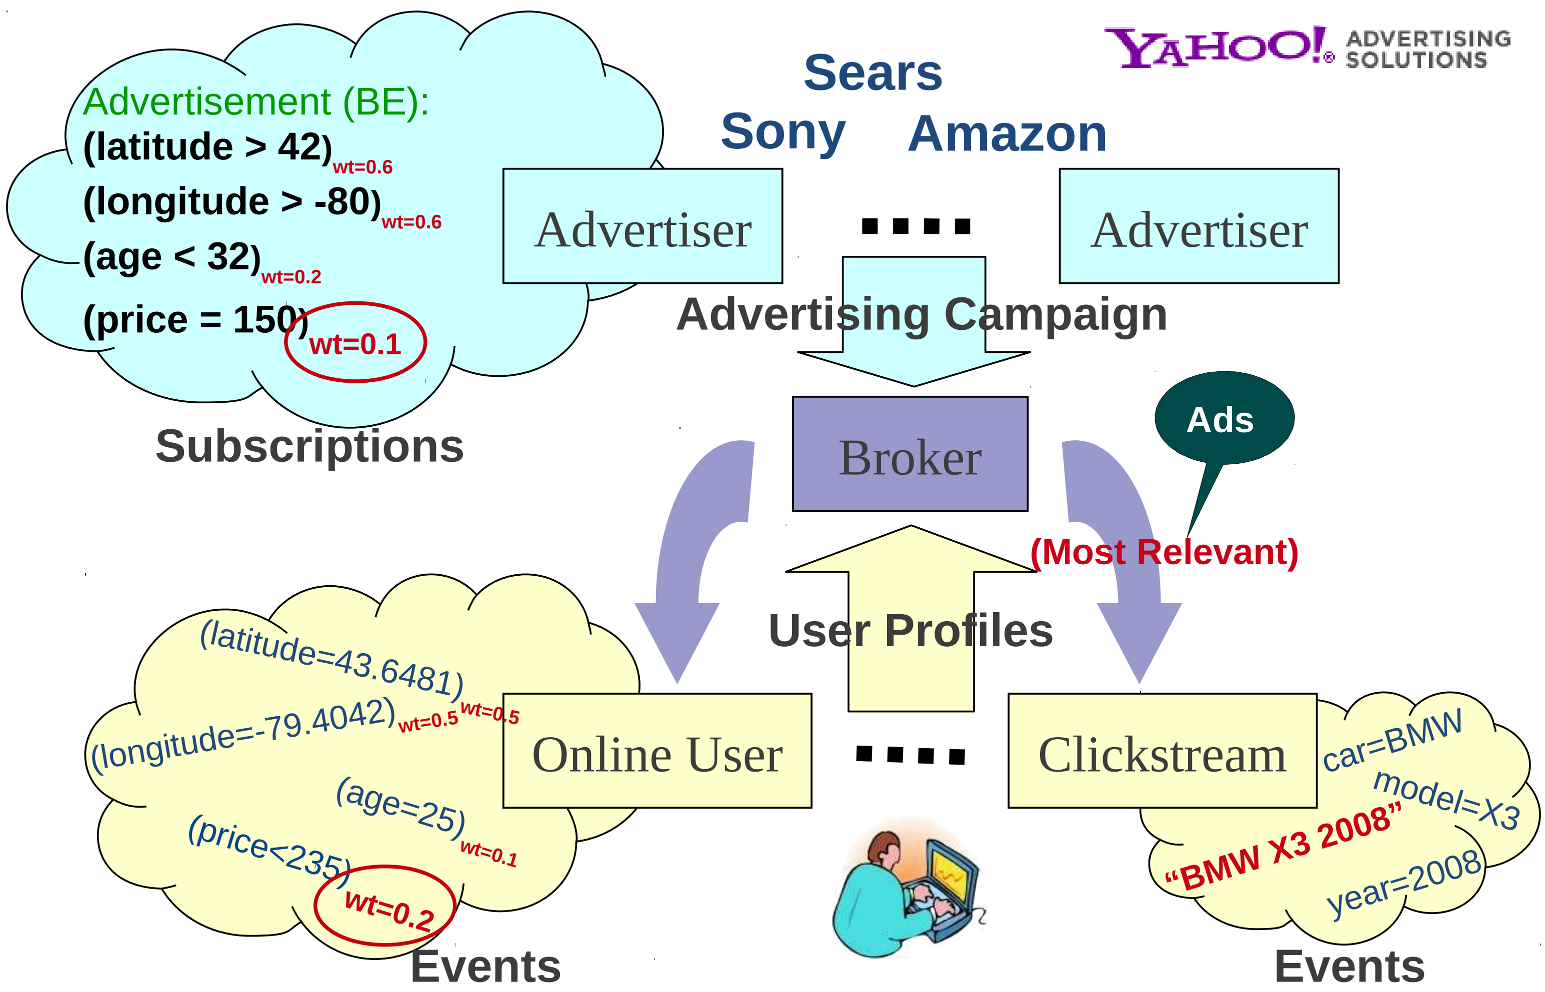

UNIVERSITY OF TORONTO

MIDDLEWARE SYSTEMS RESEARCH GROUP

Application Scenarios

1. Computational advertising (targeted advertising)
2. Computational finance (algorithmic trading)
3. Intrusion detection (deep packet inspection)
4. Real-time data analysis (data analytics)
5. Emerging mobile applications in co-spaces (location-based services)

Problem Statement

To continuously evaluate a set of predefined patterns/specifications (subscriptions) over incoming event stream.

Challenges Derived from Application Scenarios

Key matching problem challenges addressed in this work

1. Retrieve only the most relevant subscriptions for given a event.
2. Handle subscriptions with expressive operators (over discrete/continuous domains) that impose conditions only on a few dimensions, resulting in a high degree of overlap among subscriptions.
3. Scale to large collections of subscriptions with thousands of dimensions.
4. Sustain high matching rates of events in presence of frequent insertions and deletions of subscriptions.
5. Adapt to skewed workload distributions (self-adjusting mechanism), i.e., avoid structure deterioration.

BE-Tree Family Core Design (Two-phase Space-cutting)

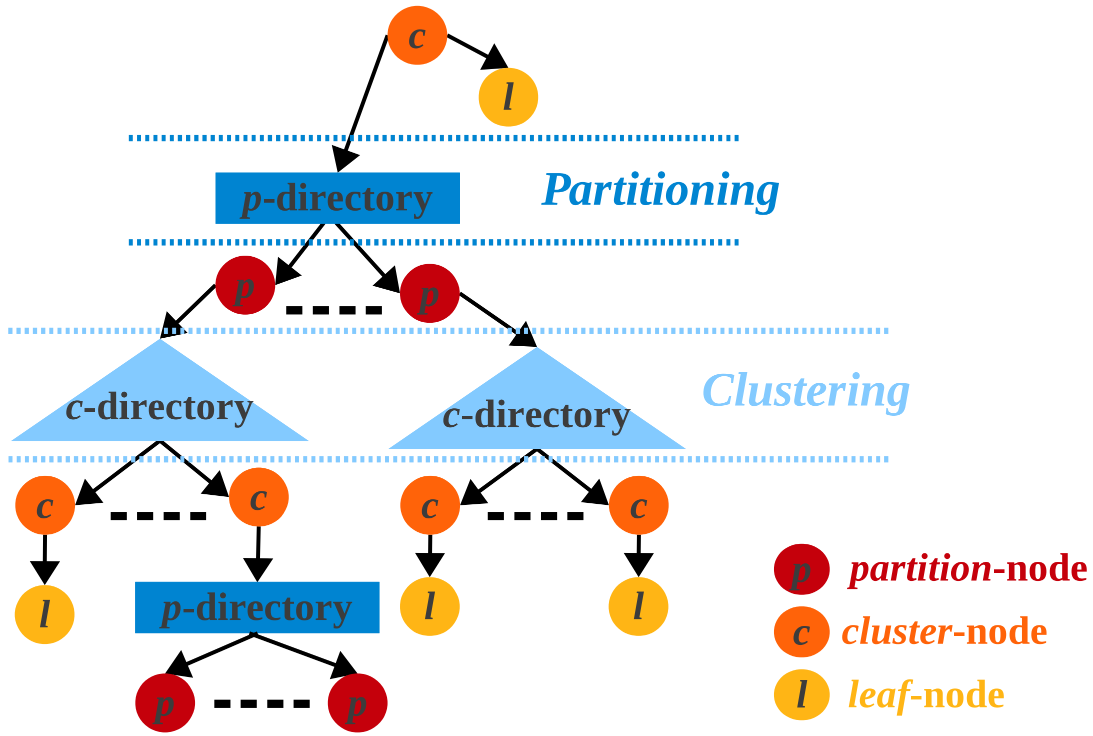

The two-phase space-cutting technique consists of

1. **space partitioning:** a global structuring to determine the best splitting dimension
2. **space clustering:** a local structuring for each partition to determine the best grouping of expressions with respect to the expressions' range of values of the chosen partition

Intuition Behind the Two-phase Space-cutting Technique

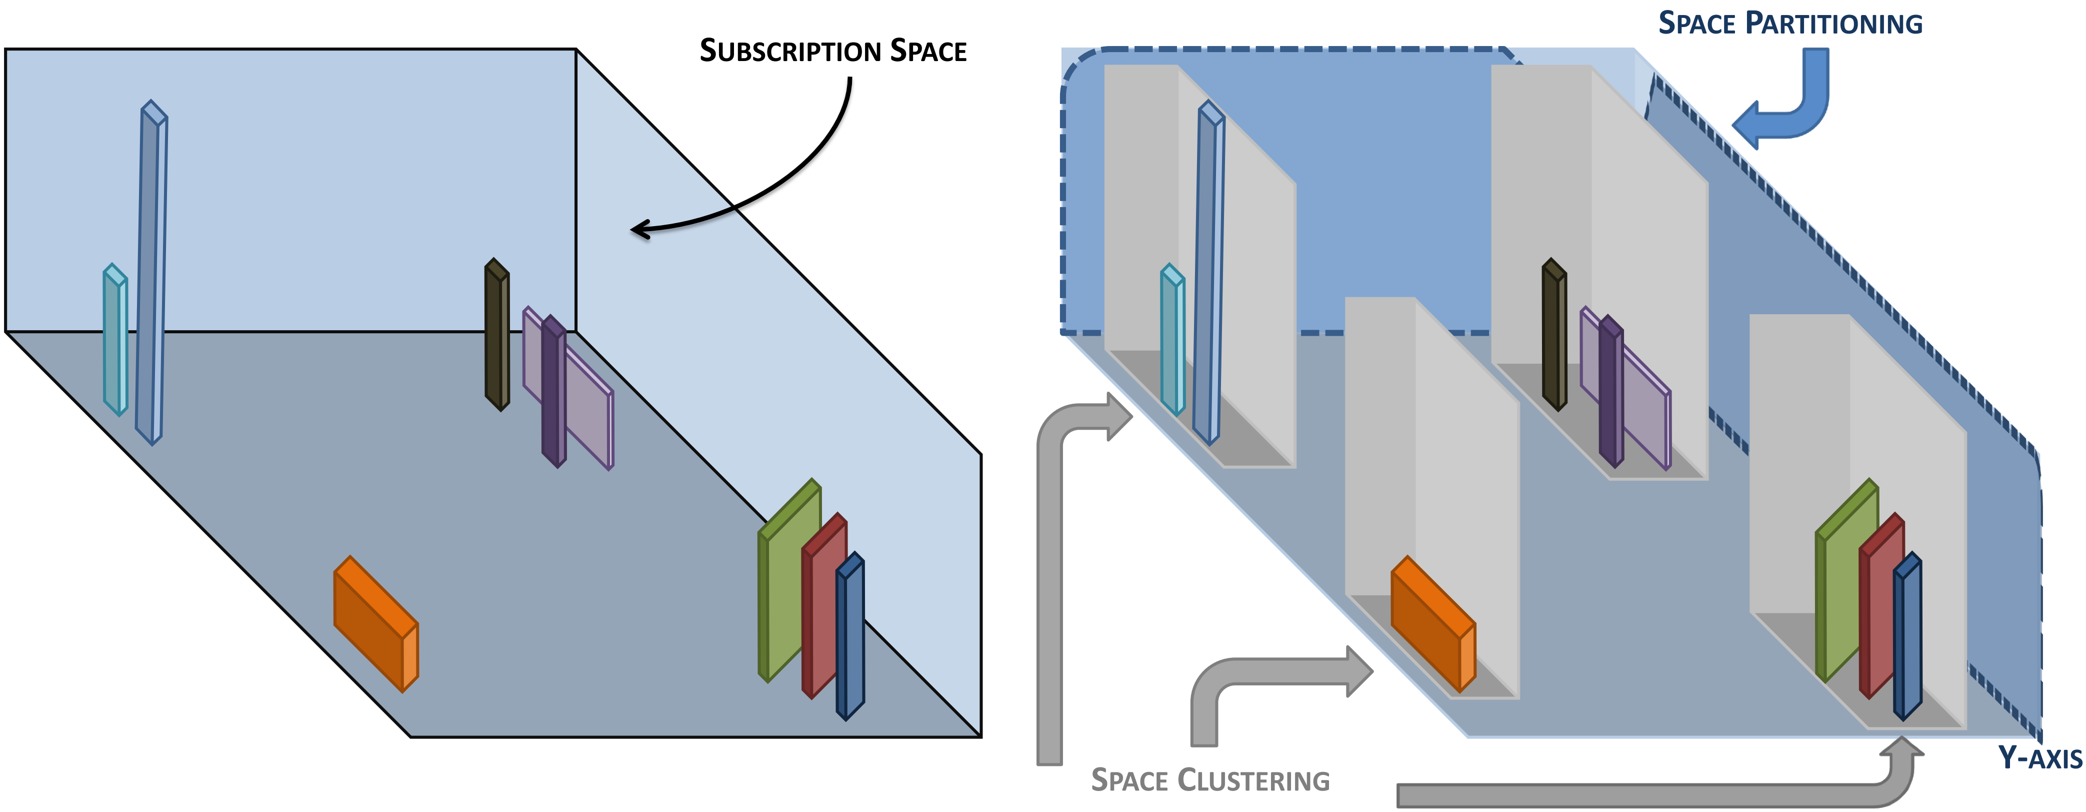

BE*-Tree Novel Features (Hierarchical Top-k Matching)

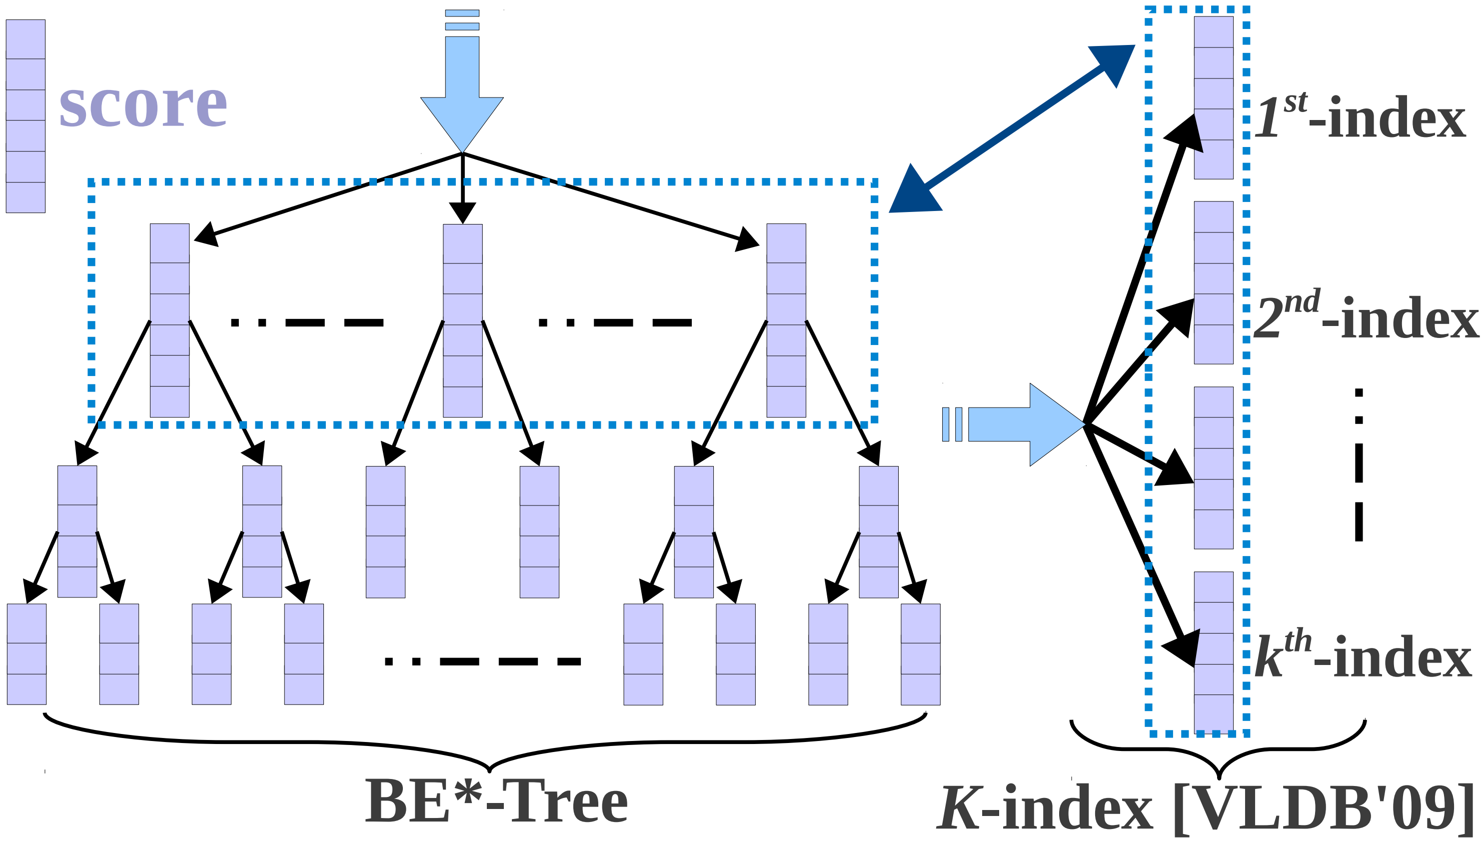

BE*-Tree continuously refining upper bound score during the matching process.

UNIVERSITY OF TORONTO

MIDDLEWARE SYSTEMS RESEARCH GROUP

Experimental Evaluation

PC-based Algorithms

1. **A-PCM:** *Parallel BE-Tree* (Sadoghi, Jacobsen)
2. **BE*:** *BE*-Tree* (Sadoghi, Jacobsen. ICDE'12)
3. **BE:** *BE-Tree* (Sadoghi, Jacobsen. SIGMOD'11)
4. **GR:** IBM *Gryphon* (Aguilera et al., PODC'99)
5. **P:** *Propagation* Algorithm (Fabret et al. SIGMOD'01)
6. *k*-**ind:** *k-index* (Whang et al. VLDB'09)
7. **SIFT:** *Counting* Algorithm (Yan et al. TODS'94)
8. **SCAN:** *Sequential Scan*

GPU-based Algorithm

1. **CLCB:** *Cuda Location-aware Content-Based Matcher* (Cugola, Margara)

Workload Configurations

Table: Synthetic and Real Workload Properties

	Workload Size	Number of Dimensions	Dimension Cardinality	Predicate Selectivity	Dimension Selectivity	Sub/Event Size	% Equality Pred	Match Prob		DBLP (Author)	DBLP (Title)	Match Prob (Author)	Match Prob (Title)		Location Workload	Parallel Matching
Size	100K-1M	1M	100K	100K	100K	100K	1M	1M		100-760K	50-250K	400K	150		2.5M	5M
Number of Dim	400	50-1400	400	400	400	400	400	400		677	677	677	677		100	128
Cardinality	48	48	48-150K	48	2-10	48	48	48		26	26	26	26		65K	48
Avg. Sub Size	7	7	7	7	7	5-66	7	7		8	35	8	30		4	7
Avg. Event Size	15	15	15	15	15	13-81	15	15		8	35	16	43		4	15
Pred Avg. Range Size %	12	12	12	6-50	—	12	12	12		—	—	12	12		—	12
% Equality Pred	0.3	0.3	0.3	0.3	1.0	0.3	0.2-1.0	0.3		1.0	1.0	0.3	0.3		0.25	0.4
Op Class	Med	Med	Med	Med	Min	Med	Med	Lo-Hi		Min	Min	Lo-Hi	Lo-Hi		Hi	Hi
Match Prob %	1	1	1	1	—	1	1	0.01-9		—	—	0.01-9	0.01-9		≈ 0	$\approx$ 0-1

The experimental results were verified by the SIGMOD'11 repeatability committee.

BEGen

Our comprehensive Boolean expression workload generator: `http://msrg.org/datasets/BEGen`.

Effect of Workload Size on Matching (Log Scale)

Table: Comparing BE-Tree (PC) and CLCB (GPU)

Workload Type	BE-Tree 1.1	BE-Tree 1.3	CLCB
without location	0.081 *ms*	0.045 *ms*	N/A
with location	0.144 *ms*	0.067 *ms*	0.306 *ms*

UNIVERSITY OF TORONTO

MIDDLEWARE SYSTEMS RESEARCH GROUP

Effect of Workload Size on Matching (Log Scale)

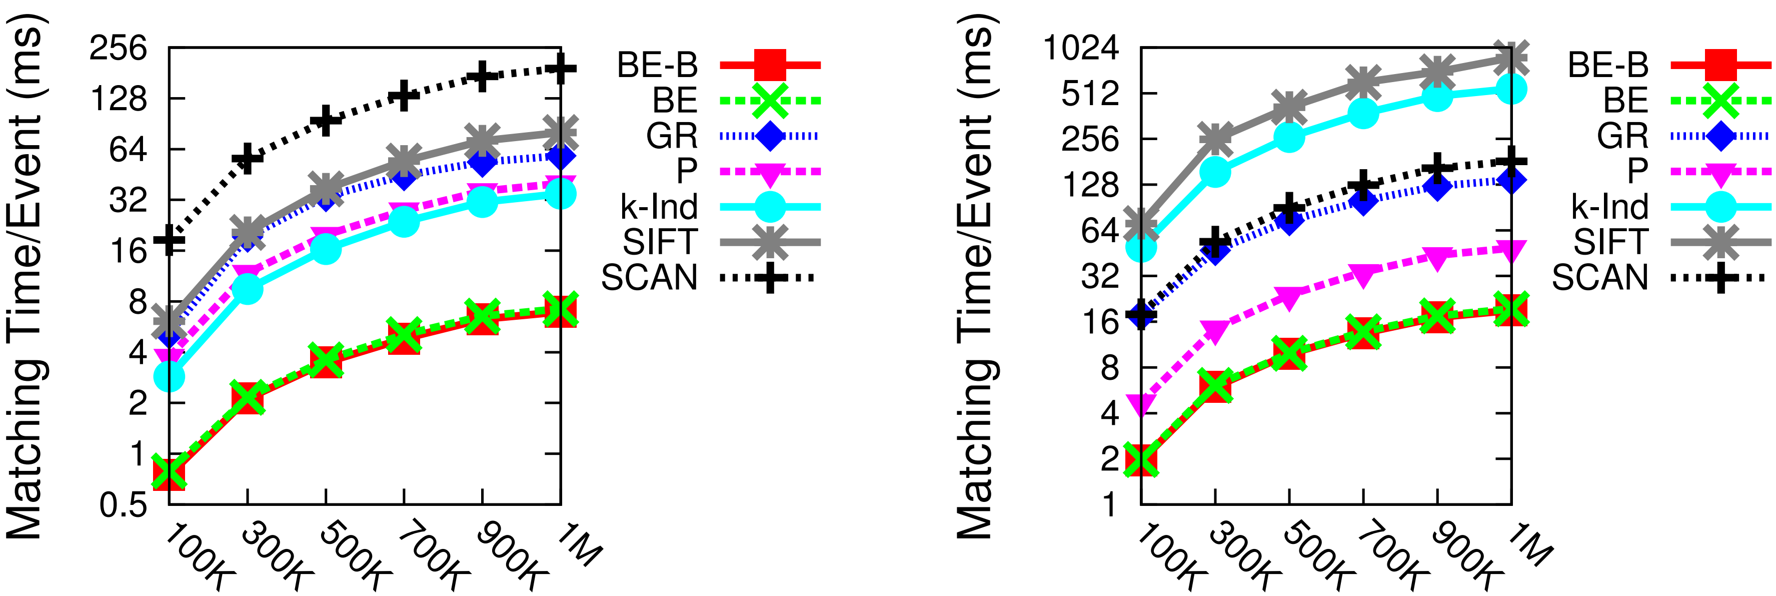

(c) Uniform: Workload Size

(d) Zipf: Workload Size

Figure: Varying Workload Size

UNIVERSITY OF TORONTO

MIDDLEWARE SYSTEMS RESEARCH GROUP

Effect of Matching Prob. on Top-k Matching (Log Scale)

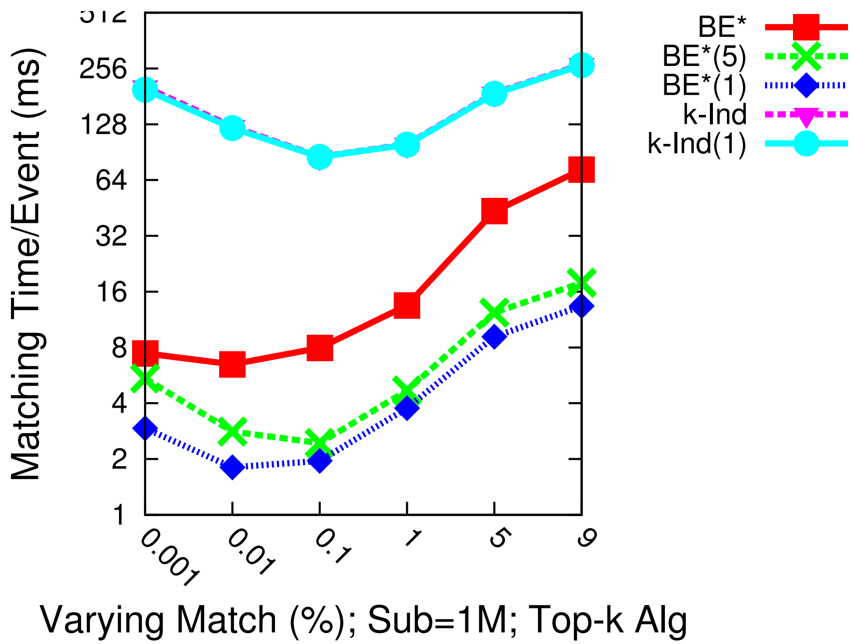

(a) Zipf Workload

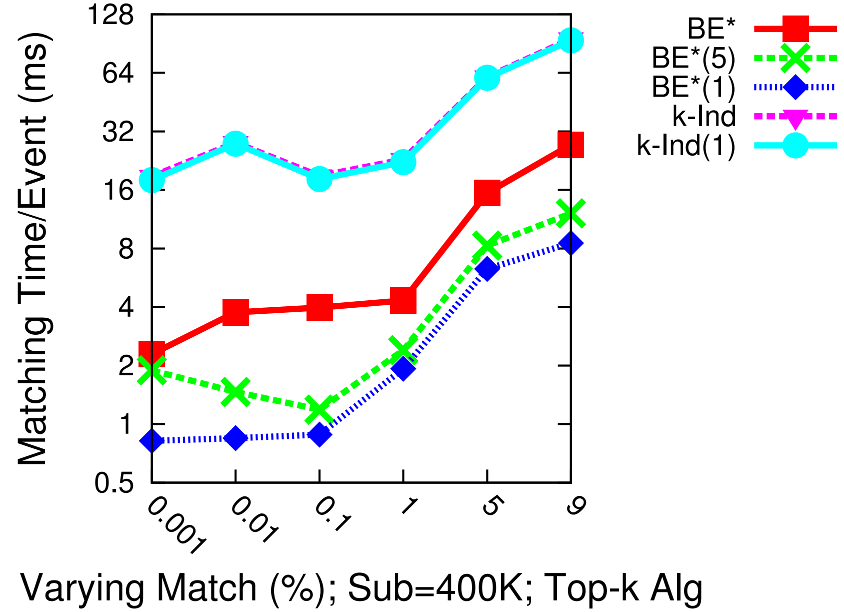

(b) DBLP Author Workload

Figure: Varying % of Matching Probability Predicates

Effect of Parallel Matching (Log Scale)

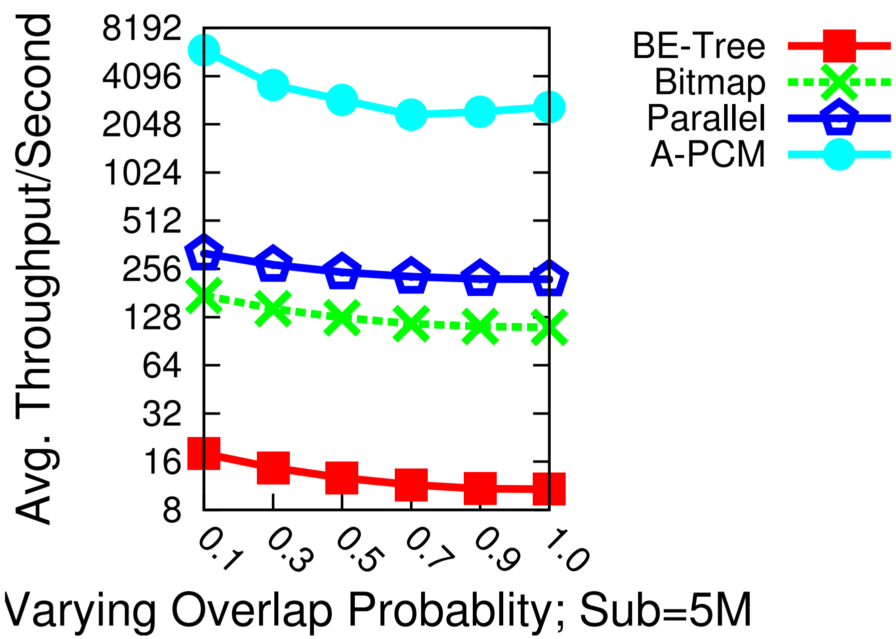

(a) Matching Probability $m = 1\%$

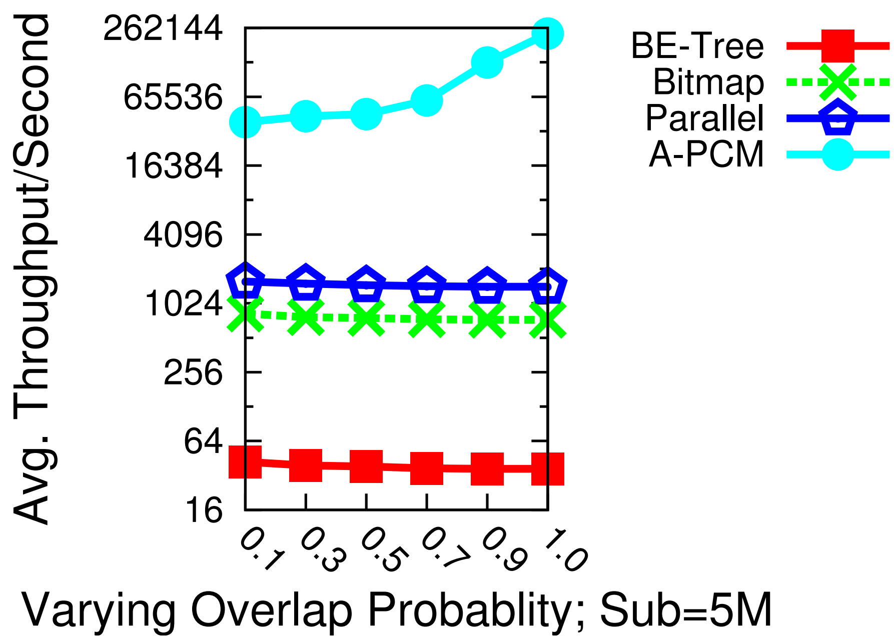

(b) Matching Probability $\approx 0\%$

Figure: Varying % of Stream Similarity

Conclusions

1. *BE-Tree* is a major step forward in addressing notable challenges such as scalability, expressiveness, dynamic construction and adaptation, by proposing a novel, self-adjusting index structure [SIGMOD'11].
2. *BE-Tree* also solves the problem of location-based matching (contrary to the claim that specialized algorithm is a must for location-based matching) [SIGMOD'11].
3. *BE-Tree* provably outperforms existing prominent approaches presented in the scientific literature. Our results were verified by the SIGMOD'11 repeatability committee.
4. *BE*-Tree* has potential to impact the design of computational advertising engines, in which click streams and user profile information is matched against advertisement inventory to serve the most advertisements [ICDE'12].
5. Our hardware acceleration can play an essential role in the design of high-throughput and low-matching latency requiring event processing engines for real-time data analysis, e.g., algorithmic trading [VLDB'10, DEBS'11, DaMoN'11, ICDE'12].

UNIVERSITY OF TORONTO

MIDDLEWARE SYSTEMS RESEARCH GROUP

References

- Amer Farroukh, Mohammad Sadoghi, and Hans-Arno Jacobsen.
  Towards vulnerability-based intrusion detection with event processing.
  In *Proceedings of the 5th ACM international conference on Distributed event-based system*, DEBS'11, pages 171–182, New York, New York, USA, 2011. ACM.
- Gianpaolo Cugola and Alessandro Margara.
  High-performance location-aware publish-subscribe on GPUs.
  In *Proceedings of the ACM/IFIP/USENIX 13th International Middleware Conference*, volume 7662 of *Lecture Notes in Computer Science*, pages 312–331, Montreal, QC, Canada, 2012. Springer.
- Mohammad Sadoghi.
  Towards an extensible efficient event processing kernel.
  In *Proceedings of the on SIGMOD/PODS 2012 PhD Symposium*, PhD '12, pages 3–8, Scottsdale, Arizona, USA, 2012. ACM.
- Mohammad Sadoghi, Ioana Burcea, and Hans-Arno Jacobsen.
  GPX-Matcher: a generic Boolean predicate-based XPath expression matcher.
  In *Proceedings of the 14th International Conference on Extending Database Technology*, EDBT/ICDT '11, pages 45–56, Uppsala, Sweden, 2011. ACM.
- Mohammad Sadoghi and Hans-Arno Jacobsen.
  Indexing Boolean expression over high-dimensional space.
  Technical Report CSRG-608, University of Toronto'10, 2010.
- Mohammad Sadoghi and Hans-Arno Jacobsen.
  BE-Tree: an index structure to efficiently match Boolean expressions over high-dimensional discrete space.
  In *Proceedings of the 2011 ACM SIGMOD International Conference on Management of data*, SIGMOD '11, pages 637–648, Athens, Greece, 2011. ACM.
- Mohammad Sadoghi and Hans-Arno Jacobsen.
  Relevance matters: Capitalizing on less (top-k matching in publish/subscribe).
  In *IEEE 28th International Conference on Data Engineering*, ICDE'12, pages 786–797, Arlington, Virginia, USA, 2012. IEEE Computer Society.
- Mohammad Sadoghi, Hans-Arno Jacobsen, Martin Labrecque, Warren Shum, and Harsh Singh.
  Efficient event processing through reconfigurable hardware for algorithmic trading.
  Proceedings of the VLDB Endowment, 3(2):1525–1528, 2010.
- Mohammad Sadoghi, Rija Javed, Naif Tarafdar, Harsh Singh, Rohan Palaniappan, and Hans-Arno Jacobsen.
  Multi-query stream processing on fpgas.
  In *Proceedings of the 2012 IEEE 28th International Conference on Data Engineering*, ICDE '12, pages 1229–1232, Washington, DC, USA, 2012. IEEE Computer Society.
- Mohammad Sadoghi, Harsh Singh, and Hans-Arno Jacobsen.
  fpga-ToPSS: line-speed event processing on FPGAs.
  In *Proceedings of the 5th ACM international conference on Distributed event-based system*, DEBS '11, pages 373–374, New York, New York, USA, 2011. ACM.
- Mohammad Sadoghi, Harsh Singh, and Hans-Arno Jacobsen.
  Towards highly parallel event processing through reconfigurable hardware.
  In *Proceedings of the Seventh International Workshop on Data Management on New Hardware*, DaMoN '11, pages 27–32, Athens, Greece, 2011. ACM.

UNIVERSITY OF TORONTO

MIDDLEWARE SYSTEMS RESEARCH GROUP