High Performance Graph Mining Systems

Xuehai Qian University of Southern California



Architecture Lab for Creative High-performance Energy-efficient Machines



Graphs are Everywhere

- A graph G = (V, E) is represented by its vertices set and edges set E:
 - E is a subset of $V \times V$, $(u, v) \in E$ iff u and v are connected by an edge
- Graphs naturally capture the relationship between entities in different applications



² <u>alchem.usc.edu</u>

Graph Mining

- Two graphs $G_0 = (V_0, E_0)$ and $G_1 = (V_1, E_1)$ are **isomorphic** iff there exists an one-to-one mapping $f: V_0 \rightarrow V_1$ such that $(u, v) \in E_0 \Leftrightarrow (f(u), f(v)) \in E_1$
 - An equivalent relation
- Graph mining: find patterns from a graph
 - Input: a large input graph; a pattern graph
 - Compute: enumerate all the subgraphs isomorphic to the given pattern—embeddings
 - Process: gather some information, depending on the application
- We consider two main types:
 - Counting: simply return the count of embeddings
 - Frequent subgraph mining (FSM)





Edge/Vertex-induced Subgraphs

- Edge-induced subgraph
 - For two graphs $g = (V_g, E_g)$, G = (V, E)such that $V_g \subseteq V, E_g \subseteq E$
- Graph G 3
- Consider edges: $u, v \in g$, $(u, v) \in G$, but $(u, v) \notin g$ —u and v are in g due to other edges
- Vertex-induced subgraph
 - For two graphs $g = (V_g, E_g)$, G = (V, E)such that $\{(u, v) | u, v \in V_g, (u, v) \in E\} = E_g$
 - Consider vertices: if two vertices u and v are in g, and there is an edge between then in G, the edge must be also in g



Edge-induced subgraph Not vertex-induced subgraph



Edge-induced subgraph Vertex-induced subgraph



4 <u>alchem.usc.edu</u>

Edge/Vertex-induced Embeddings

- Vertex-induced embedding
 - The subgraph that is isomorphic to a pattern should be a valid vertexinduced subgraph
 - The count be calculated from edge-induced embedding count
 - Example: C(vertex-induced 3chain) =C(edge-induced 3chain)-3C(edge-induced triangle).
 C=count
 - The #of vertex-induced count of 3-chain of G (on the right)=8-3×2=2

Graph G



2



Edge/Vertex-induced Embeddings

 For the vertex-set-based method (used in AutoMine*), the edge-induced embedding can be calculated with minor code modification



- 2: **for** $v_1 \in N(v_0)$ **do**
- 3: **for** $v_2 \in N(v_0) N(v_1)$ **do**
- 4: $count_{vertex-induced} \leftarrow count_{vertex-induced} + 1$
- 5: **end for**
- 6: **end for**
- 7: **end for**

Vertex-induced 3-chain



N(v): the vertex set containing all neighbors of v -N(v1): v₂ should not connect to v₁, otherwise it is triangle in the original graph—the 3-chain is not a valid vertex-induced subgraph

6

alchem.usc.edu



```
1: for v_0 \in V do
```

```
2: for v_1 \in N(v_0) do
```

- 3: **for** $v_2 \in N(v_0)$ **do**
- 4: $count_{edge-induced} \leftarrow count_{edge-induced} + 1$
- 5: **end for**
- 6: **end for**
- 7: **end for**

Edge-induced 3-chain

In the talk, we consider edge-induced embeddings

* Daniel Manwhirter, et al. AutoMine: harmonizing high-level abstraction and high performance for graph mining. SOSP'19



Graph Mining Applications

- Mining biochemical structures
- Finding biological conserved subnetworks
- Finding functional modules
- Program control flow analysis
- Intrusion network analysis
- Mining communication networks
- Anomaly detection
- Mining XML structures
- Building blocks for graph classification, clustering, compression, comparison, correlation analysis, and indexing





7

Graph Mining Systems

- While it is important, it is hard to write graph mining codes for various applications
 - Different patterns leads to different algorithms
 - The same algorithm can be implemented in various ways with different performance
- A general graph mining system can offer better programmability and high performance
 - Users simply specify the patterns
 - The system chooses the best implementations
 - A typical domain-specific system with similar motivation as graph processing systems
- Graph mining vs. graph computation
 - Graph computation: simple computation, memory bound
 - Graph mining: computational intensive



8

Existing Graph Mining Systems

- Single-machine systems
 - RStream (OSDI'18)
 - AutoMine (SOSP'19)
 - Peregrine (EuroSys'20)
 - Pangolin (VLDB'20)
 - Kaleido (ICDE'20)
- Distributed systems
 - Arabesque (SOSP'15)
 - G-Miner (Eurosys'18), G-Thinker (ICDE'20)
 - Fractal (SIGMOD'19)
 - GraphPi (SC'20)





9

Arabesque: Exhaustive Check



Teixeira et al. Arabesque: A System for Distributed Graph Mining. SOSP'15





RStream: Leveraging Relational Algebra



Wang et al. RStream: Marrying Relational Algebra with Streaming for Efficient Graph Mining on A Single Machine. OSDI'18



High Performance Graph Mining Systems



11

AutoMine: Compiler-generated Pattern Enumeration using Cost Model



Mawhirter et al. AutoMine: Harmonizing High-level Abstraction and High Performance for Graph Mining. SOSP'19



High Performance Graph Mining Systems

ALCHEM

alchem.usc.edu

12

Peregrine: Complete Symmetry Breaking



A Symmetric Pattern

The four embeddings are redundant: (v0,v1,v2,v3) (v2,v1,v0,v3) (v2,v3,v0,v1) (v0,v3,v2,v1) → should be only counted once

Symmetry breaking:

add constraints— (v0<v2) and (v1<v3) → only one embedding left



Jamshidi et al. Peregrine: A Pattern-Aware Graph Mining System. EuroSys'20





Pangolin: Application-specific Optimizations with flexible APIs

Pangolin Applications (TC, CF, MC, FSM)							
Pangolin API							
Execution Engine	Helper Routines		Embedding List Data Structure				
Galois System							
Multicore C	PU	GPU					

- 1 bool toExtend(Embedding emb, Vertex v);
- 2 bool toAdd (Embedding emb, Vertex u)
- 3 bool toAdd (Embedding emb, Edge e)
- 4 Pattern getPattern (Embedding emb)
- 5 Pattern getCanonicalPattern(Pattern pt)
- 6 Support getSupport (Embedding emb)
- 7 Support Aggregate (Support s1, Support s2)
- 8 bool toPrune(Embedding emb);

Chen et al. Pangolin: An Efficient and Flexible Graph Mining System on CPU and GPU. VLDB'20





GraphPi: Choose the Best Symmetry Breaking



- Introduced the restriction set generator
 - Systematically explore various symmetry breaking restrictions
 - Eliminate all redundant computation due to symmetry
- Replicated distributed execution
- If the innermost K for-loops are independent, the count is calculated mathematically

Shi et al. GraphPi: High Performance Graph Pattern Matching through Effective Redundancy Elimination. SC'20



High Performance Graph Mining Systems

alchem.usc.edu

15

Existing Graph Mining Systems

Single-machine systems

- RStream (OSDI'18):
 - Relational algebra based API and implementation
- AutoMine (SOSP'19):
 - Compiler generated algorithms for pattern enumeration
- Peregrine (EuroSys'20):
 - Pattern-based programming model enabling pattern-aware optimization
- Pangolin (VLDB'20):
 - A set of flexible APIs enables powerful pattern-specific optimizations
 - The first graph mining system supporting GPU
- Kaleido (ICDE'20):
 - Succinct intermediate data representation && faster isomorphism test

Distributed systems

- Arabesque (SOSP'15):
 - Exhaustively check all subgraphs up to the pattern size
- G-Miner (Eurosys'18), G-Thinker (ICDE'20)
 - Subgraph-centric programming model with partitioned graph
- Fractal (SIGMOD'19):
 - DFS-based embedding exploration; build-from-scratch paradigm to reduce memory footprint
- GraphPi (SC'20):
 - Search for better symmetry breaking; an mathematical method to speedup counting

16

alchem.usc.edu



Further Improving Performance?

- Observation: all existing systems consider each given pattern as a while
 - Empirically, the embedding enumeration cost can increase rapidly as the pattern size grows
- We build a new graph mining system based on pattern decomposition*
 - Decompose a target pattern into several smaller subpatterns
 - Compute the count of each
 - The results of the target (original pattern) can be calculated using the subpattern counts with very low additional cost
- Our project started in Fall 2019, many new papers came out in 2020, but fortunately our performance is still better than all
 - The importance of algorithm

* Ali Pinar, et al. Escape: Efficiently counting all 5-vertex subgraphs. WWW'17



Execution time of AutoMine (our own implementation) on EmailEuCore graph. 6-chain embeddings is 19,620× compared to 3-chain enumeration





Counting with Pattern Decomposition

- We explain the algorithm using relational algebra as a mathematical tool
 - The implementation still uses vertex-set-based method
- A pattern decomposition of pattern graph $p = (V_p, E_p)$ is determined by the vertex cutting set V_c
 - A subset of V_p , of which the removal breaks p into K connected components
 - An edge-induced embedding of p can be represented by a $|V_p|$ -tuple $(v_0,v_1,v_2,\ldots,\|V_p\|-1)$
 - v_i is the vertex in the embedding (subgraph) that matches the vertex i in the pattern graph
 - Each unique embedding corresponds to M tuples due to symmetric
- We can organize such tuples in a conceptual embedding table



* Ali Pinar, et al. Escape: Efficiently counting all 5-vertex subgraphs. WWW'17





Counting with Pattern Decomposition

- The K subpatterns correspond to K embedding tables: T_1, T_2, \ldots, T_K
- T_{K+1} : relational join of all T_1, T_2, \ldots, T_K using the columns associated with the cutting set V_C as keys
 - Contains all edge-induced embeddings of the original pattern p
- However, T_{K+1} contains more tuples for two reasons:
 - Symmetric: different tuples represent the same embedding—valid embeddings counted multiple times
 - Duplicated elements: embeddings matching the subpatterns contain one or more same vertices other than for cutting sets—invalid embeddings

* Ali Pinar, et al. Escape: Efficiently counting all 5-vertex subgraphs. WWW'17





Embedding Tables



- However, T_{K+1} contains more tuples for two reasons:
 - Symmetric: different tuples represent the same embedding valid embeddings counted multiple times
 - Duplicated element: embeddings matching the subpatterns contain the same vertices other than for cutting sets—invalid embeddings

• How to eliminate them?

	7	[1		_		7	ľ2					T ₃	}	
v_0	$\boldsymbol{v_1}$	v_2	v_3		$\mathbf{v_0}$	v ₁	v ₂	v ₄		$\mathbf{v_0}$	v ₁	\mathbf{v}_2	\mathbf{v}_3	v_4
а	b	g	f		b	а	g	f		а	b	g	f	d
а	g	b	f		g	а	b	f	-	а	b	g	f	С
а	b	g	с		b	а	g	с		а	b	g	с	d
а	g	b	С	_	g	а	b	С		а	b	g	С	С
а	С	b	f		С	а	b	f		а	b	С	f	g
а	b	С	f	• •	b	а	С	f		а	b	С	f	d
а	С	b	g	\bowtie	с	а	b	g		а	b	С	g	g
а	b	С	g		b	а	с	g		а	b	С	g	d
b	g	а	d		g	b	а	d		b	а	g	d	f
b	а	g	d		а	b	g	d		b	а	g	d	С
b	g	а	с		g	b	а	С		b	а	g	С	f
b	а	g	с	_	а	b	g	С		b	а	g	С	С
b	а	с	g		а	b	С	g		b	а	с	g	f
b	С	а	g		С	b	а	g	-	b	а	С	g	g
b	а	с	d		а	b	С	d		b	а	с	d	f
b	С	а	d	_	С	b	а	d	-	b	а	С	d	g



Shrinkage Pattern

- Generated by shrinking at least two vertices in pattern graph p belonging to different subpatterns
 - Specifically construct the patterns that contain duplicated elements
 - It is proved that this method can eliminate all invalid tuples*
- After eliminating embeddings matching shrinkage patterns (invalid), handling duplicate tuples (valid counted many times) is easier



* Ali Pinar, et al. Escape: Efficiently counting all 5-vertex subgraphs. WWW'17





Is Decomposition Always Better?

- The answer is NO. It does not guarantee the total runtime reduction
 - The combined number of enumerated subpatterns may be increased—we did not observe it
 - Some subpatterns after the decomposition may be very frequent
 - One of the subpatterns of a size-5 pattern is the very frequent 4-loop
- A performance model is necessary to estimate the cost of computation to avoid these cases



High Performance Graph Mining Systems

22

System Challenges

- While the decomposition-based algorithm is known, there are several challenges in building a general system
- Challenge 1: huge algorithm search space
 - A pattern specification typically has multiple patterns
 - 112 patterns for 6-motif; 823 patterns for 7-motif
 - With computation reuse, the mining of these patterns are fused together
 - Cutting set for each should be determined jointly
- Challenge 2: fast and accurate cost estimation
- Challenge 3: decomposition not compatible with symmetry breaking
- Challenge 4: beyond counting—advanced mining tasks such as frequent subgraphs mining (FSM)





23

DwarvesGraph: A Decomposition-based Graph Mining System

- We build a new graph mining system based on pattern decomposition
 - APIs to support various mining tasks
 - Approximate-mining based cost model
 - Efficient decomposition space search
 - Partial symmetry breaking



Input

24



DwarvesGraph APIs

- Users provide two programs
 - Application program: specify the major user-defined logics
 - Compilation program: specify the patterns to mine and invoke the compiler to generate the code
- APIs—both for convenient use and advanced applications:
 - High-level: int get_pattern_count();
 - Low-level: partial-embedding-centric model



High Performance Graph Mining Systems

25

Partial-Embedding Centric Model

- A new model designed for decomposition-based graph mining
 - A partial-embedding matches a subpattern
- void process_partial_embedding(PartialEmbedding pe, int count);
 - Invoked by the system when the partial-embedding pe can be "extended" to reach at least one complete embedding of the whole pattern
 - The count indicates how many complete embeddings can *be extended from the partial-embedding*
- std::vector<Embedding> materialize(PartialEmbedding pe, int num);
 - Concretize the first num embeddings of the whole pattern from the partial embedding
- These seem to be arbitrary, what are the system guarantees?

either 3 or 4. Two ways to reach the whole embedding from the partial embedding:

count=2 passed as the parameter.



High Performance Graph Mining Systems

26

Partial-Embedding Centric Model

• Complete Guarantee:

- If a partial-embedding *pe* matching a subpattern *P_{sub}* is passed to process_partial_embedding...
- Then all other partial embeddings matching P_{sub} will be also passed

• Coverage Guarantee:

- The set of subpatterns matched by the passed partialembeddings must fully cover all vertices of the pattern graph
- It is more relaxed than decomposition





A Simple Example





High Performance Graph Mining Systems

alchem.usc.edu

28

Frequent Subgraph Mining (FSM)

3-Chain

Pattern

29

alchem.usc.edu

Input graph

- **Domain** of a *pattern vertex*:
 - ullet The set of input graph vertices that can map to it ullet
 - Dom(A)={0,1};Dom(B)={0,1,2,3};Dom(C)={0,1,2,3}
- Support—a metric to quantize the *frequency* of a pattern
 - We use the minimum image-based (MINI) support definition*
 - MINI support=the size of the smallest domain across all pattern vertices
 - MINI support of the 3-chain on the input is $|\{0,1\}|=2$.
- FSM aims to discover frequent patterns
 - The application considers all patterns \leq a certain size
 - Return a pattern if its support is no less than a user-specified threshold





FSM in DwarvesGraph





returned tuple: (1,0,2,*)
vertex D in 4 chain is

alchem.usc.edu

30

- vertex D in 4-chain is UNDETERMINED
- domain[A][1]=1
- domain[B][0]=1
- domain[C][2]=1
- count is not used

- The correctness is ensured by
 - Complete guarantee: all partial-embeddings of a subpattern are returned
 - Coverage guarantee: all vertices in the pattern graph are covered



Efficient Implementation

- The users are not aware of the decomposition-based algorithm implementation
- The DwarvesGraph compiler generates the implementations based on decomposition
- The partial-embedding centric model is general
 - The complete embedding can be considered as a special case for partial-embedding
 - If the implementation does not use decomposition, our API can still work—each embedding matches the complete pattern
- The compiler efficiently generates the procedure explained with embedding tables without expensive relational algebra
 - Unlike RStream, which implements relational algebra



High Performance Graph Mining Systems

31

Efficient Implementation

T₄ $\mathbf{v}_0 \ \mathbf{v}_1 \ \mathbf{v}_2 \ \mathbf{v}_3 \ \mathbf{v}_4$

c f

alchem.usc.edu



Summary

- The partial-embedding centric model is not tied to the decomposition
- The complete and coverage guarantee can ensure the correctness
- The system implementation based on decomposition ensures the stronger property:
 - All subpatterns share the cutting set V_C
 - The coverage guarantee just requires all vertices are covered—disjoint subpatterns that can cover all vertices also satisfy it—but not decomposition method
- The algorithm we described is a *template* for the compiler to generate codes for the given pattern graph
 - The cutting set determined by algorithm generate engine



We Still Need to Solve...

- The cost model for algorithm generation engine to evaluate different choices of cutting set
- Efficiently search the cutting sets across multiple patterns
- How to make symmetric breaking work for decomposition as much as possible?



High Performance Graph Mining Systems

34

Cost Model

- We need to quickly evaluate the performance of generated subpattern enumeration algorithms
 - Executing of algorithms on real datasets/machines is too expensive
- Pattern enumeration is a set of nested for-loops
 - The key problem: estimate the cost of each loop
- AutoMine* is the first system that uses a cost model to select pattern matching schedules for better performance
 - Problem: its cost model is over-simplified
 - Assumes that the algorithm runs on a random graph with n vertices, each vertex pair is connected by a fixed probability p
 - For counting k-clique, #iteration of 1st,2nd,3rd,...,kth loop are n,np,np²,...,np^{k-1}. With k=5, line 6 is should be executed n⁵p¹⁰ times
 - Patents graph: n=3.8M, avg_deg=8.76, p= 2.3×10^{-6} , line 5 is estimated to execute 3.28×10^{-24} times
 - In reality, Patents graph has 3M 5-cliques, line 5 executed for $3M \times 5!$ times

Mawhirter et al. AutoMine: Harmonizing High-level Abstraction and High Performance for Graph Mining. SOSP'19



```
Counting k-clique
```

35



High Performance Graph Mining Systems

ALCHEM alchem.usc.edu

A New Cost Model

- Key insight: every iteration corresponds to a match of a pattern
- The problem is converted to the pattern count estimation of the input graph
 - Can be *approximate*
 - Only need to be *relative*
- A new cost model based on approximate graph mining
 - Generate a reduced graph by sampling input graph
 - At most 32M edges
 - Run neighborhood sampling in ASAP* to get the approximation of the patterns up to certain size, store the results in table persisted in disk
 - During algorithm search, query the table to get the cost of loop based on the count of the corresponding pattern
- Obtain the count of frequent patterns accurately, while underestimating that of the infrequent ones



* lyer et al. ASAP: Fast, approximate graph pattern mining at scale. OSDI'18



High Performance Graph Mining Systems

1: $Cnt \leftarrow 0$				
2: for $i \leftarrow 1 \dots NumSamples$ do				
3: $v_0 \leftarrow UniformSample(V(G))$				
4: $v_1 \leftarrow UniformSample(N(v_0))$	Graph	Profiling Time (s)		
5: $v_2 \leftarrow UniformSample(N(v_1))$	CiteSeer	1.96		
6: if $v_2 \in N(v_0)$ then	MiCo	3.50		
$\begin{array}{ll} n: & Cnt \leftarrow Cnt + v(G) \cdot N(v_0) \cdot N(v_1) \\ \text{s. end if} \end{array}$	Patents	6.64		
9: end for	LiveJournal	7.14		
10: $Cnt \leftarrow Cnt / NumSamples / 6$	Friendster	7.10		

36
The New Cost Model Effectiveness



AutoMine

Our Method

alchem.usc.edu

37



Decomposition Space Search

- The graph mining applications need to handle multiple patterns
 - Motif Counting (MC) aims at counting all connected patterns with a particular size
- Need to select a cutting set for each pattern
- With computation reuse, enumeration of multiple patterns can be fused, the search becomes joint
- We propose the circulant tuning method with fast convergence





High Performance Graph Mining Systems

38

Circulant Tuning



Partial Symmetry Breaking

 Problem: with symmetry breaking for subpatterns, the complete embeddings cannot be correctly joined





High Performance Graph Mining Systems

AI CHFM

alchem.usc.edu

DwarvesGraph Evaluation

• System:

- Each node has two 8-core Intel Xeon E5-2630 CPUs (hyperthreading disabled) and 64GB DRAM
- GPU (Pangoline): NVIDIA V100 GPU with 32GB memory
- Arabeque and Fractal (distributed) use 8 nodes. Arabeque uses Hadoop 2.7.7, Fractal uses Apache Spark 2.2.0

• Applications:

- Motif Counting (MC): count all connected vertex-induced patterns with a particular size
- Pseudo Clique Mining (PC): A vertex-induced pattern is a pseudo clique if the number of its edges is no less than n(n-1)/2 k, n is the #vertex and k is a parameter
- Frequent Subgraph Mining (FSM)

• Other systems:

- In-house AutoMine implementation
- RStream (OSDI'18)
- Arabesque (SOSP'15)
- Peregrine (EuroSys'20)
- Pangoline (CPU/GPU) (VLDB'20)
- Fractal (SIGMOD'19)
- GraphPi (SC'20)

App		Graph	Our Impl.	Original Impl.
		wk	27.3ms	34.5ms
	3 MC	mc	161ms	230ms
	3-MC	pt	0.9s	9s 1.9s 0s 13.4s 0s 11.5s
		lj	9.0s	13.4s
	4-MC	wk	7.0s	11.5s
		mc	31.7s	45.2s
		pt	24.3s	82.1s
		lj	457m	367m
		wk	4345s	5300s
	5-MC	mc	2.91h	5.56h
		pt	54m	117m

Graph	Abbr.	IVI	IEI	ILI
CiteSeer [5, 18, 43]	cs	3.3K	4.5K	6
EmailEuCore [28, 54]	ee	1.0K	16.1K	42
WikiVote [26]	wk	7.1K	100.8K	N/A
MiCo [16]	mc	96.6K	1.1M	29
Patents [27]	pt	3.8M	16.5M	N/A
Labeled-Patents [27]	lpt	2.7M	14.0M	37
LiveJournal [4, 30]	lj	4.8M	42.9M	N/A
Friendster [53]	fr	65.6M	1.8B	N/A
RMAT-100M [10]	rmat	100M	1.6B	N/A

Graph datasets

41

alchem.usc.edu



Comparing with AutoMine, RStream, and Arabesque

App.	G	DwarvesGraph	AutomineInHouse	RStream	Arabesque
	cs	0.16ms	0.16ms(1.0x)	142ms (888x)	10.1s (63.125 x)
	ee	0.8ms	7.3ms (8.9x)	21.0s(25.471x)	10.2s (12.352x)
IC	wk	7.6ms	27.3 ms (3.6 x)	17.9m(141.437x)	12.1s(1.586x)
3-N	pt	335.7ms	931 ms (2.8 x)	104.1m (18.611x)	96.4s (287x)
	mc	48.0ms	161 ms (3.4 x)	144.8m(181.051x)	21.1s (440x)
	li	2.88	9.0s(3.3x)	T	24.3m (529x)
	CS	0.2ms	4.8ms(23x)	3.78(17.647x)	9.98(46.794x)
	ee	9.4ms	920ms (98x)	132.4m (842.186x)	19.1s (2.023x)
IC	wk	60.0ms	7.0s(117x)	T	402.2s(6.704x)
4-V	pt	1.58	24.3s(16x)	T	68.3m(2.711x)
7	mc	1.3s	31.7s(24x)	T	42.8m(1.942x)
	li	32.8s	456.5m (836x)	T	C
	cs	1.4ms	332ms (229x)	146.4s (101.124x)	11.4s (7.843x)
	ee	360.3ms	104.8s (291x)	T	19.4m (3.233x)
1C	wk	5.3s	72.4m(823x)	T	C
5-N	pt	32.68	53.9m (99x)	T	Ċ
	mc	114.7s	174.6m (91x)	T	C
	li	167.7m	Т	Т	С
	cs	247.0ms	35.9s (145x)	108.7m (26,403x)	48.7s (197x)
4C	ee	91.3s	259.0m (170x)	Т	Ċ
Q-9	wk	38.7m	T	Т	C
_	pt	57.9m	Т	Т	C
	cs	0.3ms	0.5ms (1.7x)		
SC	ee	719ms	67.1s (93x)		
1-L	wk	735ms	90.8s (24x)		
	pt	499ms	15.7s (31x)		
	cs	0.3ms	0.5ms (1.7x)		
PC	ee	1.3s	433.1s (322x)		
8-]	wk	1.2s	463.0s (387x)		
	pt	582ms	86.2s (148x)		
0	cs	0.2ms	0.3ms (1.5x)	522ms (2,609x)	10.3s (51,315x)
-3(ee	0.2ms	0.2ms (1.0x)	3.6s (18,090x)	9.6s (48,235x)
SM	lpt	20.8s	20.3s (0.98x)	4,713.5s (226x)	C
н	mc	308ms	441ms (1.4x)	149.1m (29,013x)	C
X	cs	0.6ms	0.6ms (1.0x)	77.9ms (130x)	9.6s (15,931x)
1-3	ee	0.2ms	0.2ms (1.0x)	210ms (1,049x)	9.8s (48,985x)
SN	lpt	18.6s	18.1s (0.98x)	89.0m (287x)	C
H	mc	124ms	300ms (2.4x)	141.9m (68,813x)	157.9s (1,276x)

Graph	Abbr.	IVI	IEI	ILI
CiteSeer [5, 18, 43]	cs	3.3K	4.5K	6
EmailEuCore [28, 54]	ee	1.0K	16.1K	42
WikiVote [26]	wk	7.1K	100.8K	N/A
MiCo [16]	mc	96.6K	1.1M	29
Patents [27]	pt	3.8M	16.5M	N/A
Labeled-Patents [27]	lpt	2.7M	14.0M	37
LiveJournal [4, 30]	lj	4.8M	42.9M	N/A
Friendster [53]	fr	65.6M	1.8B	N/A
RMAT-100M [10]	rmat	100M	1.6B	N/A
Friendster [53] RMAT-100M [10]	fr rmat	65.6M 100M	1.8B 1.6B	N/ N/

Graph datasets



High Performance Graph Mining Systems

Comparing with Peregrine, Pangolin, and Fractal

App.	G	DwarvesGraph	Peregrine	Pangolin(CPU/GPU)	Fractal
	cs	0.16ms	5.8ms	5.0ms / 0.1ms	5.9s
3-MC	pt	0.3s	1.4s	1.4s / 0.2s	79.7s
	mc	48ms	60ms	280ms / 14.1ms	12.9s
	cs	0.2ms	21.2ms	15.3ms / 0.7ms	6.0s
4-MC	pt	1.5s	11.2s	329.5s / 8.0s	141.6s
	mc	1.3s	5.3s	242.7s / 3.7s	58.4s
	cs	1.4ms	41.7ms	688.3ms / 1.3ms	6.1s
5-MC	pt	32.6s	513.6s	C/C	4517.0s
	mc	114.7s	5,635.1s	C/C	1240.0s
6 MC	cs	0.2s	0.8s	14.9s / C	4.6s
0-MC	pt	3,472.6s	Т	C/C	Т
FSM-100		14.0s	C	C/C	346.6s
FSM-300		9.6s	C	C/C	280.2s
FSM-1K	mc	2.5s	1,782.2s	C/C	169.1s
FSM-3K		0.5s	189.3s	C/C	109.4s
FSM-1K		1,511.5s	Т	C/C	Т
FSM-10K	Int	71.4s	34,403.6s	C/C	Т
FSM-20K	Ipt	9.0s	4,781.0s	333.3s / C	270.1s
FSM-25K		2.7s	1,353.3s	126.5s / C	250.7s

Graph	Abbr.	IVI	IEI	ILI
CiteSeer [5, 18, 43]	cs	3.3K	4.5K	6
EmailEuCore [28, 54]	ee	1.0K	16.1K	42
WikiVote [26]	wk	7.1K	100.8K	N/A
MiCo [16]	mc	96.6K	1.1M	29
Patents [27]	pt	3.8M	16.5M	N/A
Labeled-Patents [27]	lpt	2.7M	14.0M	37
LiveJournal [4, 30]	lj	4.8M	42.9M	N/A
Friendster [53]	fr	65.6M	1.8B	N/A
RMAT-100M [10]	rmat	100M	1.6B	N/A

Graph datasets

Pangolin's GPU performance is competitive, but achieved with a significantly more expensive device (NVIDIA V100-32GB).



Comparing with GraphPi



GraphPi only handles the individual pattern and does not support FSM

Graph	Abbr.	IVI	IEI	ILI
CiteSeer [5, 18, 43]	cs	3.3K	4.5K	6
EmailEuCore [28, 54]	ee	1.0K	16.1K	42
WikiVote [26]	wk	7.1K	100.8K	N/A
MiCo [16]	mc	96.6K	1.1M	29
Patents [27]	pt	3.8M	16.5M	N/A
Labeled-Patents [27]	lpt	2.7M	14.0M	37
LiveJournal [4, 30]	lj	4.8M	42.9M	N/A
Friendster [53]	fr	65.6M	1.8B	N/A
RMAT-100M [10]	rmat	100M	1.6B	N/A

Graph datasets



Decomposition Space Search Methods

Circulant tuning is slower than separate tuning

R: random; S: separate tuning; C: circulant tuning RT: runtime; ST: search time

App.	Graph	R-RT	S-RT	S-ST	C-RT	C-ST
	CS	5.3ms	2.2ms	5.0ms	1.4ms	0.6s
ИС	ee	917ms	392ms	5.0ms	360ms	0.4s
5-N	wk	14.4s	8.2s	5.0ms	5.3s	0.8s
	pt	69.3s	36.9s	1.7ms	32.6s	0.7s
6-MC	cs	576ms	280ms	37.2ms	247ms	325s
	ee	437.7s	99.8s	35.8ms	91.3s	252s
	wk	_	2,515.9s	40.1ms	2,320.0	409s
	pt	—	3,688.5s	43.7ms	3,472.6	222s

Circulant tuning achieves up to 1.57x speedup. For large graph, the benefit is more and search time can be amortized.



High Performance Graph Mining Systems

Partial Symmetry Breaking and Decomposition



p0 — p19 are all size-5 patterns except for 5-clique



High Performance Graph Mining Systems

alchem.usc.edu

Large Graphs and Large Patterns

Graph	#Vertices	#Edges	App.	Runtime (s)
fr	65.6M	1 9 D	4-Motif	4,301
11	05.0141	1.0D	4-Chain	862
rmot	100M	1 6 D	4-Motif	5,900
Innat	100101	1.0D	4-Chain	800

None of the previous exact graph mining systems have reported 4-motif results on graphs at this scale



k-CHM: k-chain mining

Keep increasing the size of the pattern until the task cannot finish within 24 hours



High Performance Graph Mining Systems

alchem.usc.edu

DwarvesGraph: A Decomposition-based Graph Mining System

- We build a new graph mining system based on pattern decomposition
 - APIs to support various mining tasks
 - Approximate-mining based cost model
 - Efficient decomposition space search
 - Partial symmetry breaking
- The results show that our system is faster than all existing systems and can likely scale to large patterns



High Performance Graph Mining Systems

48

Why Distributed Graph Mining System?

- Single-machine shared memory systems (Peregrine, Pangolin, AutoMine,DwarvesGraph)
 - Both #cores and amount of memory is limited to one machine
- Single-machine out-of-core systems (RStream)
 - Scale to large graph with external storage
 - #cores is still limited to one machine
 - Sacrifice efficiency: to fully utilize disk bandwidth, a less efficient algorithm with graph streaming and relational join—much slower than recent systems
- Distributed systems with graph replication (Arabesque, Fractal, GraphPi)
 - The complete graph data replicated in each machine
 - Scaling #cores, but not memory
- Distributed systems with graph partition (G-Miner, G-Thinker)
 - Both #cores and amount of memory can scale
 - Current systems sacrifice efficiency and programmability



High Performance Graph Mining Systems

49

Graph Mining Systems with Graph Partition

- Poor programmability: complicated task-based subgraph-centric model
 - Users responsible for dividing the enumeration process into a number of subgraph-centric tasks
 - Each task: users specify the subgraph containing all data needed
 - Example: clique, task—"counting the number of k-cliques containing a given vertex", subgraph—an induced subgraph including all 1-hope neighbors of the vertex
 - Users need to handle system problems such as stragglers
- Inefficient system for communication/computation scheduling
 - Reference-counting based SW cache with GC for fetched remote data
 - Triangle counting on Patents dataset (3.8M vertices)
 - G-Thinker (8 nodes: 16 cores each, 128 cores in total): 285.3s
 - A simple single-thread implementation (reported in AutoMine): 6.2s
 - Peregrine with 16 cores single machine: 1.1s



High Performance Graph Mining Systems

50

Goals & Problems

- Khuzdul: A distributed graph mining system with graph partition with simple programming interface and high performance
- Problem 1: Can users just specify the patterns without considering all other system issues?
- Problem 2: How to efficiently control scheduling of computation and communication?
- Problem 3: How to achieve efficient implementation?



High Performance Graph Mining Systems

51

Khuzdul: Key Ideas

- Breaking down the vertex-setbased enumeration process into smaller operations that can be expressed with vertex functions
 - Vertex functions transparent to users, unlike graph processing—"think like the vertex"
- The inter-loop dependency among vertex functions
- Efficient multi-level scheduler designed for inter-loop dependency
- Optimizations to reduce data movements



52

alchem.usc.edu



Inter-loop Dependent Vertex Function



High Performance Graph Mining Systems

alchem.usc.edu

Inter-loop Dependent Vertex Function

 We can illustrate the idea conceptually with Python-like pseudocode





High Performance Graph Mining Systems

alchem.usc.edu

Inter-loop Dependent Vertex Function

Our system generates C++ codes based on the userspecified patterns

Aggregator<uint64_t> triangle_cnt_agg;

```
class ProcessSecondTriangleVertex: public VertexFunction {
 Set<VertexFunction*> process vertex(VertexId v. VertexSet neighbors, Objects shared_objs, Buffer workspace) {
   VertexSet * v_0 nbrs = shared_ojbs->get("v_0_nbrs"); // obtain the shared N(v_0) from the previous vertex-function
   VertexSet intersection_result(workspace); // use the workspace buffer to create the vertex set containing N(v_0)
   intersect N(v 1)
   // cnt += |N(v_0)| intersect N(v_1)
   VertexSet::intersect(v_0_nbrs) &neighbours, &intersection_result);
   triangle_cnt_agg.add(intersection_result.size());
 return NULL:
JInter-loop data
dependence
class ProcessFirstTriangleVertex: public VertexFunction {
 Set<VertexFunction*> process_vertex(VertexId v, VertexSet neighbors, Objects shared_objs, Buffer workspace) {
   Set<VertexFunction*> S;
                                                                                      Implemented Triangle-Counting Algorithm:
   // specify the objects to be shared with new vertex-functions
                                                                                      N(v): the neighbour vertex set of v;
   Objects objs to share;
   objs_to_share.put("v_0_nbrs", &neighbours);
                                                                                      cnt = 0;
                                                                                      for v_0 in V(G):
   // for v_1 in N(v_0):
                                                                                         for v_1 in N(v_0):
   for (VertexId u in neighbours) {
                                                                                             cnt += |N(v_0) \setminus intersect N(v_1)|
     // allocate a vertex-function to calculate 'N(v_0) \intersect N(v_1)'
     VertexFunction * f = allocate vertex Manction<ProcessSecondTriangleVertex>(u);
     f->set_shared_objs(objs_to_share);
     // the buffer needed by the new vertex-function stores N(v_0) \intersect N(v_1)
     // whose size cannot exceed min(|N(v_0)|, |N(v_1)|)
     f->set_workspace_size(min(get_degree(u), get_degree(v)) * sizeof(VertexId));
     S.add(f); Level 0 → Level 1
   return S;
 3
}:
                                                                                                           High Performance Graph Mining Systems
                                                                                                     55
                                                                                                           alchem.usc.edu
```

Abstract Execution Model

How to pick up the next vertex-function determines the concrete execution model. FIFO: breadth-first schedule FILO: depth-first schedule

1: $F_v \leftarrow$ a set of user-provided initial vertex-functions

- 2: while F_v is not empty do
- 3: pick up a vertex-function f_v from F_v consume item
- 4: remove f_v from F_v

5:
$$New_{-}F_{v} \leftarrow \operatorname{execute}(f_{v})$$

- 6: for all $f_{\boldsymbol{v}} \in New_{-}F_{\boldsymbol{v}}$ do
- 7: add $f_{\boldsymbol{v}}$ to $F_{\boldsymbol{v}}$

end for

9: end while

The execution of a vertexfunction may trigger new vertex-functions in the next loop level (in the algorithm).

56

alchem.usc.edu

8:

produce

new item

Problems of Conventional Scheduler



High Performance Graph Mining Systems

alchem.usc.edu



• Same level

- Communication batching and overlapping with computation
- Shuffle the vertex functions into different groups, each group only fetches the data from one node
- Each vertex function only accesses the neighbors of a given vertex—all in one node

Different level

• Avoid memory fragmentation and reference counting





Reducing Communication with SW Cache







Communication Merging



 Efficient hash-table based implementation that allows false negative in merging





NUMA Subpartition



- Each NUMA socket maintains an execution engine
 - All the buffers of this engine (e.g., fetched graph data, workspace buffer) are NUMA-local
 - Avoid expensive cross-socket memory accesses



High Performance Graph Mining Systems

61

Khuzdul Evaluation

• System:

- Each node has two 8-core Intel Xeon E5-2630 CPUs (hyperthreading disabled) and 64GB DRAM
- Network: 56GBps InfiniBand
- Applications: triangle, 3-motif, 4-clique, 5-clique
- Single-machine systems
 - In-house AutoMine implementation
 - Peregrine (EuroSys'20)
 - Pangoline (CPU/GPU) (VLDB'20)

Distributed Systems

- Partitioned graph: G-thinker (ICDE'20) Graph datasets
- Replicated graph: GraphPi (SC'20)



62

Comparing with G-Thinker: Partitioned Graph Triangle Counting

Graph	Khuzdul	G-thinker
wk	24.4ms	1.0s (41.0x)
mc	40.0ms	2.1s (52.5x)
pt	254.9ms	285.3s (1,119.3x)
lj	826.8ms	31.6s (38.2x)
uk	690.1s	CRASHED
tw	2171.2s	CRASHED
fr	81.2s	CRASHED





Comparing with GraphPi: Replicated Graph

Application	Graph	Khuzdul (8-node)	GraphPi (8-node)
Triangle	wk	24.4ms	534.3ms
	Мс	40.0ms	704.4ms
	pt	257.5ms	6.7s
	lj	826.8ms	9.8s
	uk	690.1s	1268.4s
	tw	2171.2s	2886.5s
	Fr	81.2s	169.2s
3-motif	wk	24.8ms	1.1s
	mc	52.8ms	1.5s
	pt	429.3ms	13.8s
	lj	1.9s	20.1s
	uk	3,005.1s	1,380.7s
	tw	9401.0s	3,032.1s
	Fr	190.5s	388.5s



Comparing with GraphPi: Replicated Graph

Application	Graph	Khuzdul (8-node)	GraphPi (8-node)
4-clique	wk	43.0ms	500.4ms
	mc	321.5ms	844.0ms
	pt	412.7ms	6.7s
	lj	5.3s	12.8s
	uk	17,241.0s	31,008.6s
	tw	18,817.0s	TIMEOUT
	fr	190.5s	177.8s
5-clique	wk	78.1ms	522.1ms
	mc	11.1s	8.2s
	pt	822.5ms	6.8s
	lj	188.6s	174.7s
	Fr	220.3s	260.0s





Application	Graph	Khuzdul (8-node / single-node)	Automine	Peregrine	Pangolin
Triangle	wk	24.4ms / 23.4ms	9.9ms	8.3ms	7ms
	mc	40.0ms / 100.2ms	52.3ms	68.7ms	56ms
	pt	257.5ms / 1.6s	330.7ms	1.1s	289ms
	lj	826.8ms / 5.5s	2.8s	3.8s	2.2s
	uk	690.1s / 6132.8s	7305.0s	4667.0s	26.6s
	tw	2171.2s / 16806.9s	30866.1s	20605.5s	747.7s
	Fr	81.2s / 515.3s	378.3s	305.2s	384.6s





Application	Graph	Khuzdul (8-node / single-node)	Automine	Peregrine	Pangolin
3-motif	Wk	24.8ms / 39.8ms	27.4ms	26.9ms	45ms
	mc	52.8ms / 255.2ms	160.3ms	84.7ms	288ms
	pt	429.3ms / 2.8s	930.9ms	1.7s	1.5s
	lj	1.9s / 12.8s	8.9s	4.6s	29.2s
	uk	3,005.1s / 24700.190s	TIMEOUT (>10 hours)	4660.9s	TIMEOUT
	tw	9401.0s / TIMEOUT	TIMEOUT	20477.6s	TIMEOUT
	fr	190.5s / 1420.8s	1206.8s	316.1s	6305.9s



High Performance Graph Mining Systems

ALCHEM

alchem.usc.edu

Application	Graph	Khuzdul (8-node / single-node)	Automine	Peregrine	Pangolin
4-clique	wk	43.0ms / 77.0ms	47.2ms	104.6ms	47ms
	mc	321.5ms / 1.9s	1.2s	1.8s	2.8s
	pt	412.7ms / 2.6s	381.0ms	1.3s	773ms
	lj	5.3s / 37.6s	31.3s	49.6s	54.7s
	uk	17,241.0s / TIMEOUT	TIMEOUT	TIMEOUT	MEM
	tw	18,817.0s / TIMEOUT	TIMEOUT	TIMEOUT	MEM
	Fr	157.9s / 887.6s	570.3s	1237.5s	MEM



High Performance Graph Mining Systems

Application	Graph	Khuzdul (8-node / single-node)	Automine	Peregrine	Pangolin
5-clique	Wk	78.1ms / 226.3ms	124.7ms	477.2ms	146ms
	mc	11.1s / 71.2s	46.8s	78.0s	132.0s
	pt	822.5ms / 5.3s	408.4ms	1.5s	967ms
	lj	188.6s / 1385.8s	982.9s	2076.6s	MEM
	Fr	220.3s / 1593.9s	900.2s	3032.8s	MEM



Communication/Computation Overlapping

Application	Graph	Runtime / Communication Time on the Critical Path (with overlap)	Runtime / Communication Time on the Critical Path (without overlap)
3-motif	mc	52.8ms / 2.4ms	80.7ms / 24.9ms
	pt	429.3ms / 144.6ms	556.7ms / 276.0ms
	lj	1.86s / 0.13s	2.81s / 1.03s
	Fr	190.5s / 11.0s	296.6s / 116.4s
5-clique	mc	11.1s / 0.017s	11.8s / 0.6s
	pt	822.5ms / 171.7ms	938.1ms / 305.0ms
	lj	188.6s / 0.56s	197.9s / 13.4s
	Fr	220.3s / 15.0s	343.8s / 138.3s



High Performance Graph Mining Systems

alchem.usc.edu

Duplicated Request Merging

Application	Graph	Runtime/Communication Volume (with request merging)	Runtime/Communication Volume (without request merging)
3-motif	mc	52.8ms / 338.7MB	61.1ms / 623.3MB
	pt	429.3ms / 1.8GB	440.6ms / 2.1GB
	lj	1.86s / 16.9GB	2.31s / 37.8GB
	Fr	190.5s / 3.2TB	192.0s / 3.4TB
5-clique	mc	11.1s / 12.9GB	18.5s / 281.0GB
	pt	822.5ms / 1.3GB	861.3ms / 1.8GB
	lj	188.6s / 79.5GB	223.4s / 3.5TB
	Fr	220.3s / 744.7GB	722.0s / 13.9TB





SW Graph Cache

Application	G	Runtime / Communication Time on the Critical Path / Communication Volume (with cache)	Runtime / Communication Time on the Critical Path / Communication Volume (without cache)
3-motif	mc	52.8ms / 2.4ms / 0.34GB	52.3ms / 1.7ms / 0.34GB
	pt	429.3ms / 144.6ms / 1.8GB	417.1ms / 126.8ms / 2.0GB
	lj	1.86s / 0.13s / 16.9GB	1.92s / 0.18s / 18.3GB
	Fr	190.5s / 11.0s / 3.2TB	200.7s / 22.5s / 4.2TB
	uk	3,005.126s / 1.9s / 829.2GB	4595.6s / 638.5s / 83.2TB
5-clique	mc	11.1s / 0.017s / 12.9GB	11.1s / 0.019s / 16.6GB
	pt	822.5ms / 171.7ms / 1.3GB	855.9ms / 209.5ms / 1.8GB
	lj	188.6s / 0.56s / 79.5GB	187.1s / 0.56s / 121.9GB
	Fr	220.3s / 15.0s / 744.7GB	255.4s / 48.6s / 3.7TB




NUMA-aware Graph Subpartition

Application	Graph	Runtime with sub-partitioning	Runtime without sub-partitioning
3-motif	mc	52.8ms	74.5ms
	pt	429.3ms	970.5ms
	lj	1.86s	3.1s
	Fr	190.5s	327.2s
5-clique	mc	11.1s	16.1s
	pt	822.5ms	2.0s
	lj	188.6s	319.8s
	Fr	220.3s	447.5s



Scalability: LiveJournal Graph

#Nodes	Triangle	3-Motif	4-Clique	5-Clique
1	5.5s	12.8s	37.6s	1385.8s
2	2.8s	6.5s	19.1s	704.7s
4	1.5s	3.4s	9.9s	361.4s
8	826.8ms	1.9s	5.3s	188.6s
8-node speedup over 1-node	6.65x	6.74x	7.09x	7.35x



Scalability: Friendster Graph

#Nodes	Triangle	3-Motif	4-Clique	5-Clique
1	515.3s	1420.8s	887.6s	1593.9s
2	257.3s	709.3s	494.7s	816.7s
4	150.4s	364.8s	289.7s	421.3s
8	81.2s	190.5s	157.9s	220.3s
8-node speedup over 1-node	6.35x	7.46x	5.62x	7.24x





Khuzdul: Key Ideas

- Breaking down the vertex-setbased enumeration process into smaller operations that can be expressed with vertex functions
 - Vertex functions transparent to users, unlike graph processing—"think like the vertex"
- The inter-loop dependency among vertex functions
- Efficient multi-level scheduler designed for inter-loop dependency
- Optimizations to reduce data movements



76

alchem.usc.edu



High Performance Graph Mining Systems

Xuehai Qian University of Southern California



Architecture Lab for Creative High-performance Energy-efficient Machines

