

A High-Throughput Solver for Marginalized Graph Kernels on GPU

Yu-Hang Tang, Oguz Selvitopi, Doru Popovici, Aydin Buluc
Computational Research Division
Lawrence Berkeley National Laboratory

IPDPS 2020
May 21, 2020

PART 1

Marginalized graph kernel for learning on graphs

PART 2

GPU-accelerated high throughput solver

SUMMARY

PART 1


Marginalized graph kernel for learning on graphs

Scientific machine learning is key to DOE technological advances

Scientific machine learning (SciML) is a core component of artificial intelligence (AI) and a computational technology that can be trained, with scientific data, to augment or automate human skills. Across the Department of Energy (DOE), scientific machine learning (SciML) has the potential to transform science and energy research.


DOE Basic Research Needs Workshop for Scientific Machine Learning:
Core Technologies for Artificial Intelligence 2019

The successes of scientific machine learning have concentrated on select forms of data




Literature information extraction

- Text data
- Linear sequence
- Weston et al. 2019



Climate analytics


- Grid data
- Real values
- Kurth et al. 2018



Fluid Mechanics

- Mesh data
- Real values
- Raissi et al., 2020


Many scientific data and representations are beyond mere images or linear sequences



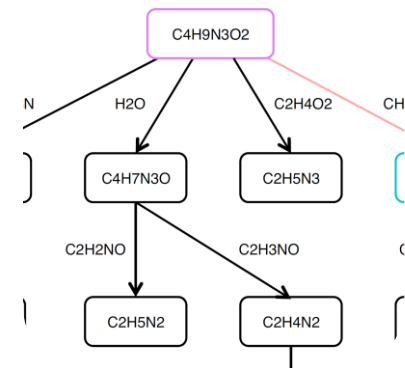
Molecules



Road network



Social network



Fragmentation tree

Variable in size

Non-sequential

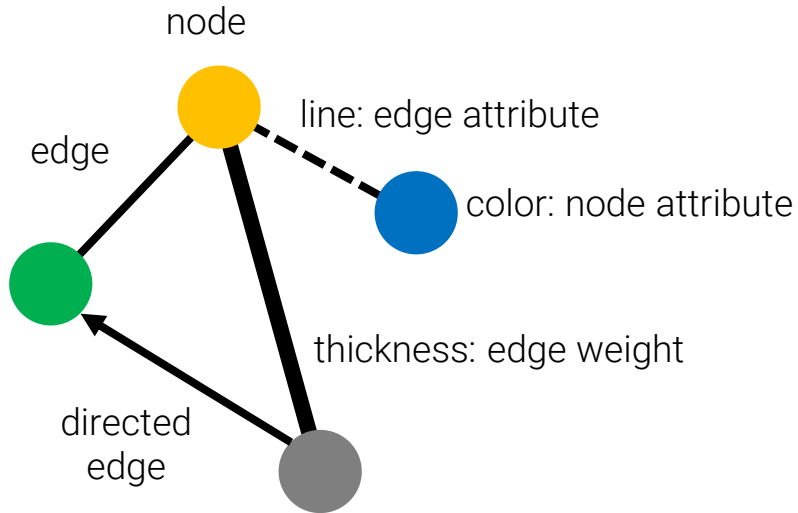
Mixed
continuous/discrete
DOFs

Existing solutions often resort to pixelating the data.

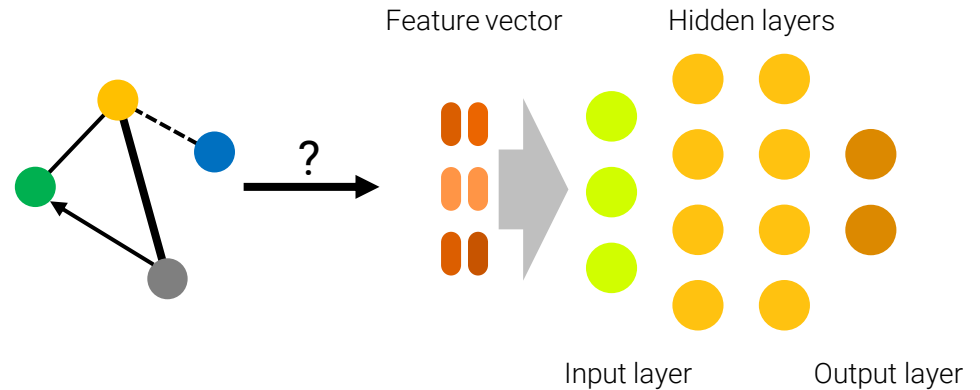
Some icons made by Freepik from www.flaticon.com

Graph is a powerful format for scientific data, but machine learning on graphs takes extra effort

- A graph is a structure that contains objects of pairwise relationships



- Most existing ML methods work on feature vectors, images, and sequences only.


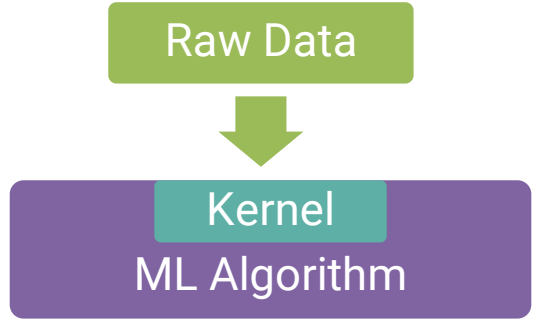


Kernel method in machine learning: what, why, and how

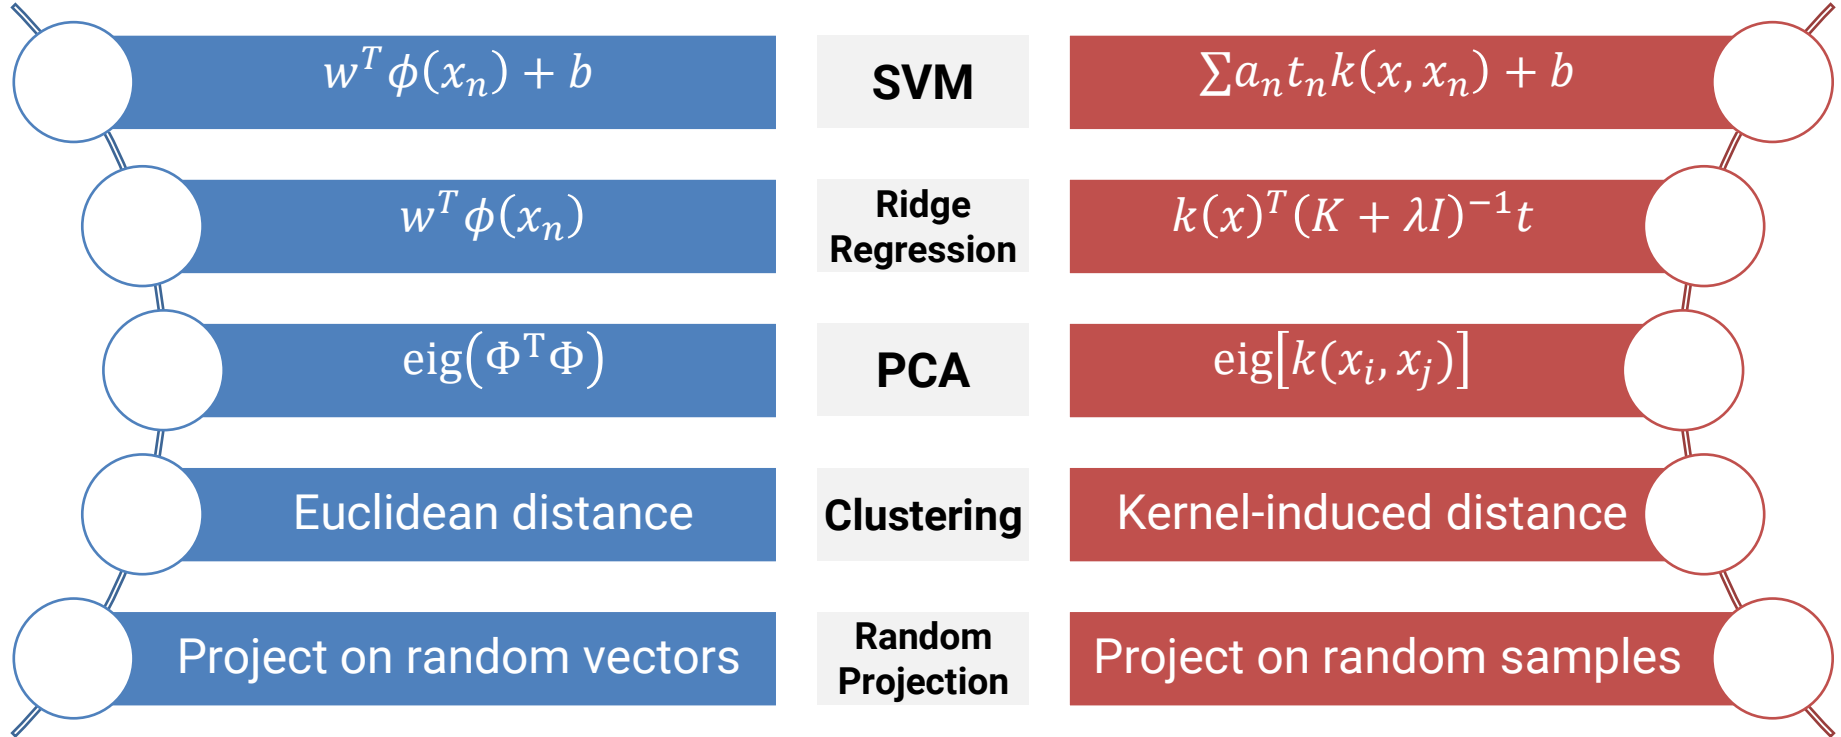
A kernel is a function that Implicitly transforms raw data into high-dimensional feature vectors via a **feature map**; and then Returns an **inner product** between the feature vectors. Must be **positive-definite**.

A kernel is useful for **Factor out** knowledge on data representation from downstream algorithms, Exploit **infinite dimensionality and nonlinear** feature spaces.

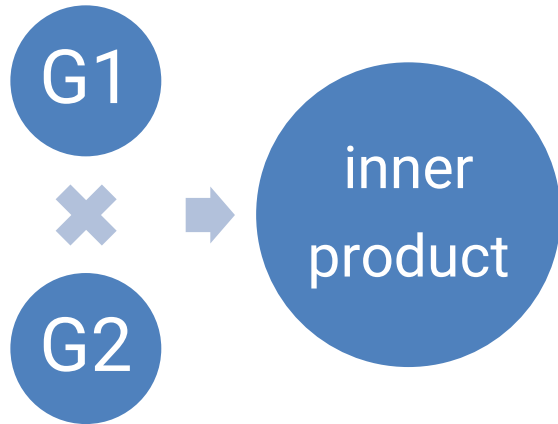
Kernels are used in Support vector machine (SVM), Gaussian process regression (GPR), Kernel principal component analysis (kPCA), etc.



Many ML algorithms have a kernelized counterpart



Graph kernels are kernels that act on graphs



Graph kernels

Histogram

Limited-size subgraphs [Ahmed et al. 2015]

Statistical moments [Debnath et al. 1991]

Random walk

Exponential [Vishwanathan, 2010]

Geometric [Vishwanathan, 2010]

Marginalized [Kashima et al., 2003]

[Shervashidze et al. 2011]

Weisfeiler-Lehman

[Morris et al. 2017]

Misc.

Shortest-path [Borgwardt and Kriegel, 2005]

Spanning tree [Ramon and Gärtner, 2003]

The marginalized graph kernel can seamlessly handle diverse types of graphs

- Definition: the inner product between two graphs is the statistical average of the inner product of simultaneous random walk paths on the two graphs.

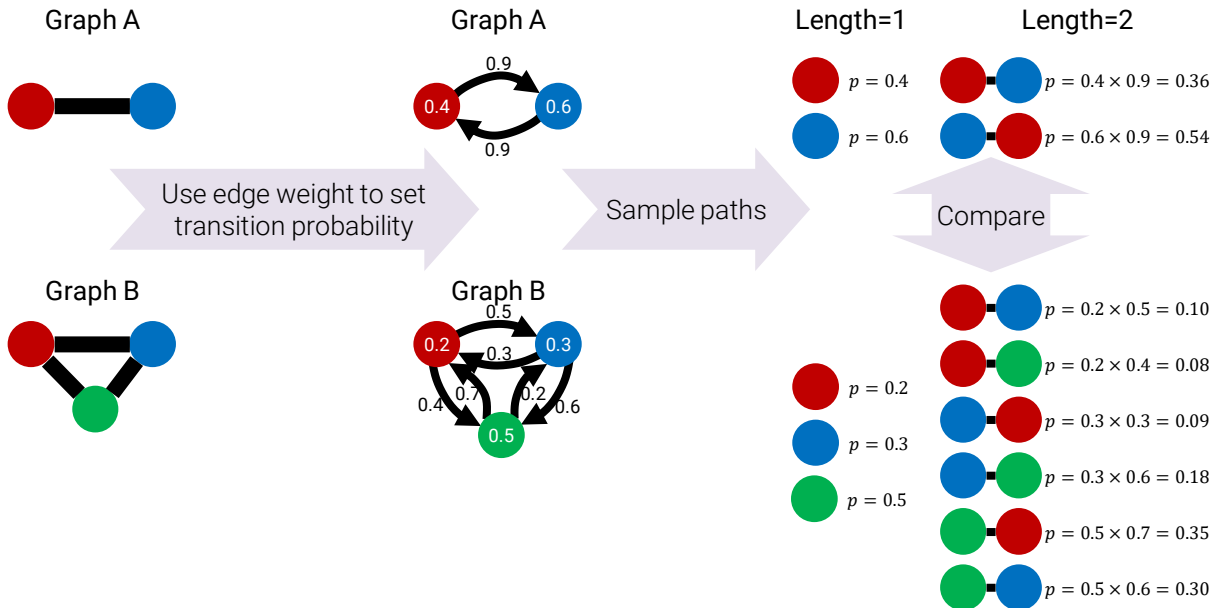
Step 1
Define random walks

$$P = D^{-1} \cdot A$$

P: transition matrix

D: degree matrix

A: adjacency matrix



The marginalized graph kernel can seamlessly handle diverse types of graphs

- Definition: the inner product between two graphs is the statistical average of the inner product of simultaneous random walk paths on the two graphs.

Step 2 Averaging path similarities

Path similarity defined
as product of base
kernel evaluations

κ_v : base kernel for nodes

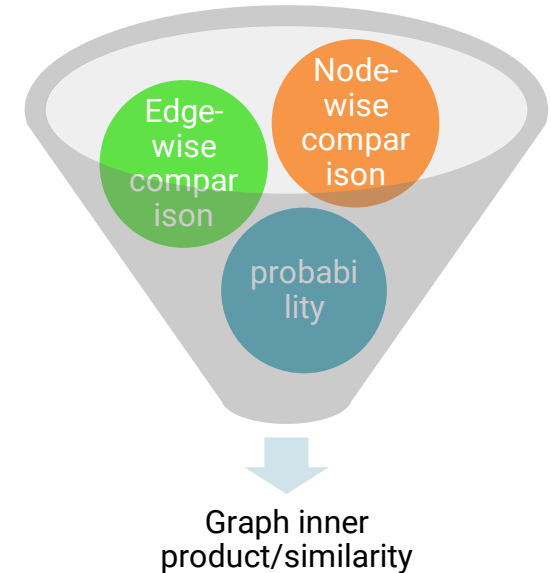
κ_e : base kernel for edges



vs



$$= \kappa_v(\text{red}, \text{red}) \cdot \kappa_e(\text{solid}, \text{dashed}) \cdot \kappa_v(\text{blue}, \text{green})$$



Wider adoption of the marginalized graph kernel was hindered due to practical challenges

$$K(G, G') = \sum_{l=1}^{\infty} \sum_{\mathbf{h}} \sum_{\mathbf{h}'} p_s(h_1) p'_s(h'_1) \mathbf{K}_v(v_{n_1}, v'_{h'_1}) \prod_{i=2}^l p_t(h_i | h_{i-1}) p_q(h_l) \prod_{j=2}^l p'_t(h'_j | h'_{j-1}) p'_q(h'_l) \prod_{k=2}^l \mathbf{K}_e(e_{n_{k-1}h_k}, e_{n'_{k-1}h'_k}) \mathbf{K}_v(v_{n_k}, v'_{h'_k})$$

Cost of computation could be high

- direct summation is **intractable**

Efficient training involving composite base kernels K_v , K_e is non-trivial

- **Analytic derivative** of the kernel is difficult to derive and implement

Linear algebra reformulation simplifies computations and reveals opportunities for optimization

$$K(G, G') = \sum_{l=1}^{\infty} \sum_h \sum_{h'} p_s(h_1) p'_s(h'_1) \mathbf{K}_v(v_{h_1}, v'_{h'_1}) \prod_{i=2}^l p_t(h_i | h_{i-1}) p_q(h_l) \prod_{j=2}^l p'_t(h'_j | h'_{j-1}) p'_q(h'_l) \prod_{k=2}^l \mathbf{K}_e(e_{h_{k-1}h_k}, e_{h'_{k-1}h'_k}) \mathbf{K}_v(v_{h_k}, v'_{h'_k})$$

- According to Kashima & Tsuda, the above computation can be simplified into

$$K(G, G') = \sum_h \sum_{h'} s(h_1, h'_1) R_{\infty}(h_1, h'_1)$$

where

$$s(h_1, h'_1) = p_s(h_1) p'_s(h'_1)$$

$$R_{\infty}(h_1, h'_1) = r_1(h_1, h'_1) + \sum_{i,j} t(i, j, h_1, h'_1) R_{\infty}(h_1, h'_1)$$

with

$$t(i, j, h_1, h'_1) = p_t(i | h_1) p'_t(j | h'_1) K_v(v_i, v'_j) K_e(e_{ih_1}, e_{jh'_1})$$

- We showed that the formulation is equivalent to the following tensor product linear system:

$$K(G, G') = \mathbf{p}_x \cdot \mathbf{R}_{\infty}$$

where R_{∞} can be solved from

$$[\mathbf{D}_x \mathbf{V}_x^{-1} - \mathbf{A}_x \odot \mathbf{E}_x] \mathbf{R}_{\infty} = \mathbf{D}_x \mathbf{q}_x.$$

$$p_{xij} = p_s(i) p'_s(j), q_{xij} = p_q(i) p'_q(j)$$

$$\text{diag}(D_x)_{ij} = \text{deg}(v_i) \text{deg}(v'_j)$$

$$\text{diag}(V_x)_{ij} = K_v(v_i, v'_j)$$

$$A_{xijkl} = w_{ij} w'_{kl}$$

$$E_{xijkl} = K_e(e_{ij}, e'_{kl})$$

multi-index: ij : element at $i \cdot n' + j$ $ijkl$: element at (ij, kl)

Linear algebra reformulation simplifies computations and reveals opportunities for optimization

$$K(G, G') = \sum_{l=1}^{\infty} \sum_h \sum_{h'} p_s(h_1) p'_s(h'_1) K_v(v_{h_1}, v'_{h'_1}) \prod_{i=2}^l p_t(h_i | h_{i-1}) p_q(h_l) \prod_{j=2}^l p'_t(h'_j | h'_{j-1}) p'_q(h'_l) \prod_{k=2}^l K_e(e_{h_{k-1}h_k}, e_{h'_{k-1}h'_k}) K_v(v_{h_k}, v'_{h'_k})$$

- The marginalized graph kernel in linear algebra form represents a modified graph Laplacian

$$K(G, G') = \mathbf{p}_\times^\top \left(\mathbf{D}_\times \mathbf{V}_\times^{-1} - \mathbf{A}_\times \odot \mathbf{E}_\times \right)^{-1} \mathbf{D}_\times \mathbf{q}_\times$$

SPD system to solve

degree vertex label adjacency edge label

Tang & de Jong, J Chem Phys, 2019: Prediction of atomization energy using graph kernel and active learning
<https://doi.org/10.1063/1.5078640>

Linear algebra reformulation simplifies derivation of analytic derivatives

- The gradient of the marginalized graph kernel is crucial for efficient training
- It can be derived using matrix calculus:

$$K(G, G') = \mathbf{p}_x^T [\mathbf{D}_x \mathbf{V}_x^{-1} - \mathbf{A}_x \odot \mathbf{E}_x]^{-1} \mathbf{D}_x \mathbf{q}_x$$

Denote

$$\mathbf{Y} = \mathbf{D}_x \mathbf{V}_x^{-1} - \mathbf{A}_x \odot \mathbf{E}_x$$

Then

$$\frac{\partial K}{\partial \theta} = \text{tr} \left[\frac{\partial K}{\partial \mathbf{Y}} \cdot \frac{\partial \mathbf{Y}}{\partial \theta} \right] = (\mathbf{Y}^{-1} \mathbf{p}_x)^T \frac{\partial \mathbf{Y}}{\partial \theta} (\mathbf{Y}^{-1} \mathbf{D}_x \mathbf{q}_x)$$

Differentiation w.r.t. other hyperparameters can be derived similarly.

Marginalized graph kernel has found successful applications in a variety of ML tasks

- Prediction of molecular atomization energy
 - nodes = atoms, edges = interatomic interactions
 - Jump probabilities proportional to edge weights, which decay with interatomic distance


$$\text{e.g. } w_{ij} = \left(1 - \frac{r_{ij}}{r_c}\right)^n$$

- Kronecker delta kernel on nodes labeled with chemical elements

$$\text{e.g. } \kappa_v(v_1, v_2) = \begin{cases} 1, & \text{if } v_1 = v_2 \\ h, & \text{otherwise} \end{cases} \text{ etc.}$$

- Gaussian kernel on edges labeled by interatomic distance


$$\text{e.g. } \kappa_e(l_1, l_2) = \exp\left[-\frac{1}{2} \frac{(l_1 - l_2)^2}{\sigma^2}\right]$$



Tang & de Jong, J Chem Phys, 2019: Prediction of atomization energy using graph kernel and active learning
<https://doi.org/10.1063/1.5078640>

Marginalized graph kernel has found successful applications in a variety of ML tasks


- Quality assurance on noisy chromatography data



Tang et al. Uncertainty Quantification and Outlier Detection on Noisy Data. Manuscript in preparation.

Marginalized graph kernel has found successful applications in a variety of ML tasks

- Protein function prediction



Borgwardt, K. M., Ong, C. S., Schönauer, S., Vishwanathan, S. V. N., Smola, A. J., & Kriegel, H.-P. (2005). Protein function prediction via graph kernels. *Bioinformatics*, 21(suppl_1), i47–i56. <https://doi.org/10.1093/bioinformatics/bti1007>


PART 2

GPU-accelerated high throughput solver for marginalized graph kernel

The marginalized graph kernel equation can be efficiently solved using conjugate gradient

- The conjugate gradient algorithm can be used to **iteratively** solve the marginalized graph kernel equation
 - V and E are not necessarily real matrices
 - κ can be complex functions

$$K(G, G') = \mathbf{p}_x^T \left(\mathbf{D}_x \mathbf{V}_x^{-1} - \mathbf{A}_x \odot \mathbf{E}_x \right)^{-1} \mathbf{D}_x \mathbf{q}_x$$




```

1 function CG4GK(d,d',v,v',A,A',E,E', q,q')
2   M ← diag [ (d ⊗ d') ⊙ (v ⊗ v')κ ]-1
3   x ← 0
4   r ← (d ⊗ d') · (q ⊗ q')
5   z ← v ⊗ v'κ
6   p ← z
7   ρ ← rTz
8   repeat
9     a ← (d ⊗ d') ⊙ (v ⊗ v')κ · p
10    + (A ⊗ A') ⊙ (E ⊗ E') · p
11    α ← ρ / (pTa)
12    x ← x + αp
13    r ← r - αa
14    z ← M-1r
15    ρ' ← rTz
16    β ← ρ' / ρ
17    p ← z + βp
18    ρ ← ρ'
19  until rTr < ε
20  return x
    
```

Naïve CG on precomputed matrices can only handle small graphs

- Due to the tensor product structure of the linear system, memory usage grows in quartic order

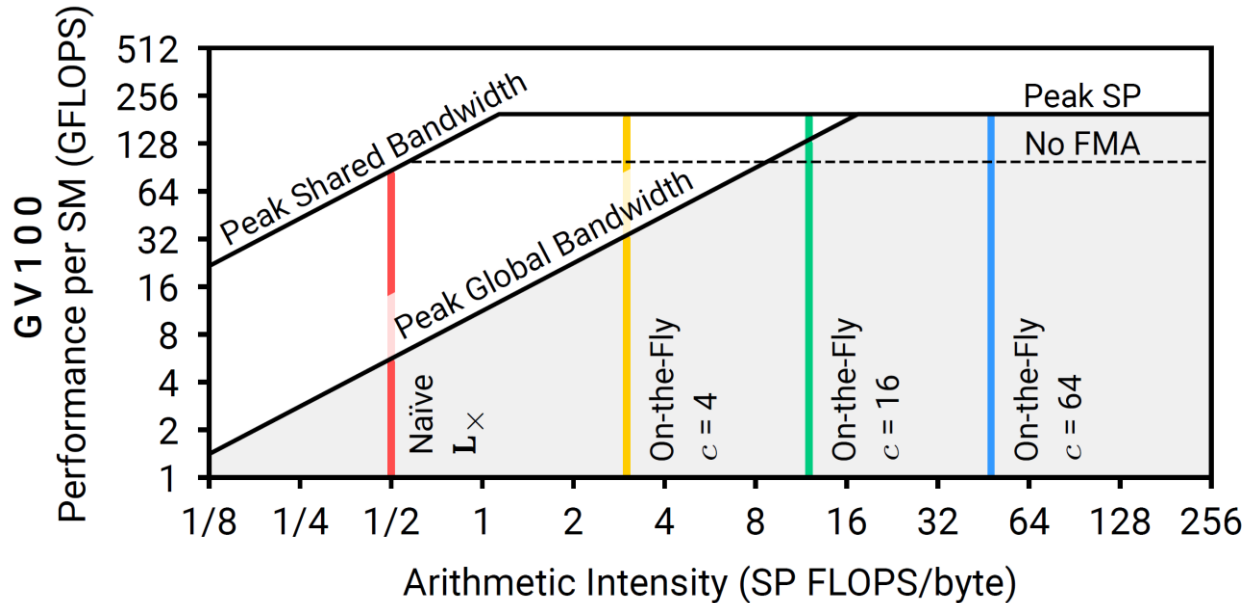
To compute similarity between a pair of 1000-node graphs, a system of **1000000×1000000 (4TB)** is involved.



```
1 function CG4GK(d,d',v,v',A,A',E,E', q,q')
2   M ← diag [ (d ⊗ d') ⊙ (v ⊗ v')κ ]-1
3   x ← 0
4   r ← (d ⊗ d') · (q ⊗ q')
5   z ← v ⊗ v'
6   p ← z
7   ρ ← rTz
8   repeat
9     a ← (d ⊗ d') ⊙ (v ⊗ v')κ -1 · p
10    + (A ⊗ A') ⊙ (E ⊗ E') · p
11    α ← ρ / (pTa)
12    x ← x + αp
13    r ← r - αa
14    z ← M-1r
15    ρ' ← rTz
16    β ← ρ' / ρ
17    p ← z + βp
18    ρ ← ρ'
19  until rTr < ε
20  return x
```

Naïve CG on precomputed matrices is also memory-bound on GPUs

- NVIDIA Volta GPU requires more than 16 FLOPS per byte (64 FLOPS per float) arithmetic intensity to achieve peak performance



On-the-fly Kronecker matrix-vector multiplication (XMV) can overcome storage and memory bandwidth difficulties

On-the-fly Kronecker matrix-vector multiplication (OTF XMV)


- Regenerates the product linear system on the fly by streaming 8-by-8 submatrices (tiles).
- Tiles staged in shared memory.
- Trade FLOPS for GB/s, but asymptotic arithmetic complexity stays the same.

```

8   repeat
9    $\mathbf{a} \leftarrow (\mathbf{d} \otimes \mathbf{d}') \odot (\mathbf{v} \otimes \mathbf{v}')^{-1} \cdot \mathbf{p}$ 
10   $+(\mathbf{A} \otimes \mathbf{A}') \odot (\mathbf{E} \otimes \mathbf{E}') \cdot \mathbf{p}$ 
11   $\alpha \leftarrow \rho / (\mathbf{p}^\top \mathbf{a})$ 
12   $\mathbf{x} \leftarrow \mathbf{x} + \alpha \mathbf{p}$ 

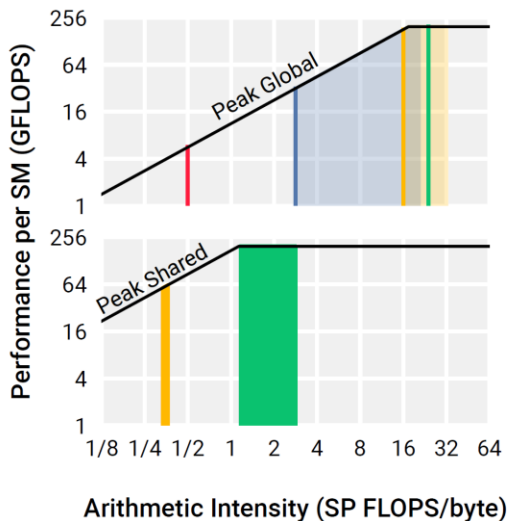
```

Tang, Selvitopi, Popovici & Buluc, IPDPS 2020: A High-Throughput Solver for Marginalized Graph Kernels on GPU <https://arxiv.org/abs/1910.06310>




OTF XMV achieves much higher FLOPS on dense graphs

- Microbenchmark on V100 with a dot product base kernel



Metrics on 5120 pairs of graphs each with 72 nodes

Naive Shared Tiling Register Blocking Tiling + Blocking




A 2-level hierarchical sparse matrix format ensures efficient memory usage

2-Level sparsity exploitation

- Outer level: retain only non-empty tiles
- Inner level: use bitmap + compact storage format

- **Packing** into compact format: performed on CPU as a preprocessing step
- **Unpacking** for OTF XMV: performed in parallel on GPU using bit magic + warp intrinsics

EMPTY
TILE
DISCARDED



NON-EMPTY
TILE
COMPRESSED

				K					
			H	L					R
		E		M		Q			
		F				O			
			I						
				N					
A	C	G	J		P				
B	D								S

DENSE STORAGE

A B C D E F G H I J K L M N O P Q R S

BITMAP

0	0	0	0	1	0	0	0	0	0
0	0	0	1	1	0	0	0	1	0
0	0	1	0	1	0	1	0	0	0
0	0	1	0	0	1	0	0	0	0
0	0	0	1	0	0	0	0	0	0
0	0	0	0	1	0	0	0	0	0
1	1	1	1	0	1	0	0	0	0
1	1	0	0	0	0	0	0	1	0

64-bit integer **nzmask**

0b100000010000001000100100000010011101010010010011001100000011000000

0x0303324AE4122041


A 2-level hierarchical sparse matrix format ensures efficient memory usage

2-Level sparsity exploitation

- Outer level: retain only non-empty tiles
- Inner level: use bitmap + compact storage format

- Heuristics for dynamic code path selection:
 - If both tiles contain more than certain number of non-zero elements, treat them as dense matrices.
 - Otherwise, compute only the non-zeros.

EMPTY
TILE
DISCARDED



NON-EMPTY
TILE
COMPRESSED

				K					
			H	L					R
		E		M				Q	
		F					O		
			I						
				N					
A	C	G	J		P				
B	D								S

DENSE STORAGE

A B C D E F G H I J K L M N O P Q R S

BITMAP

0	0	0	0	1	0	0	0	0	0
0	0	0	1	1	0	0	0	1	0
0	0	1	0	1	0	1	0	0	0
0	0	1	0	0	1	0	0	0	0
0	0	0	1	0	0	0	0	0	0
0	0	0	0	1	0	0	0	0	0
1	1	1	1	0	1	0	0	0	0
1	1	0	0	0	0	0	0	1	0

64-bit integer **nzmask**

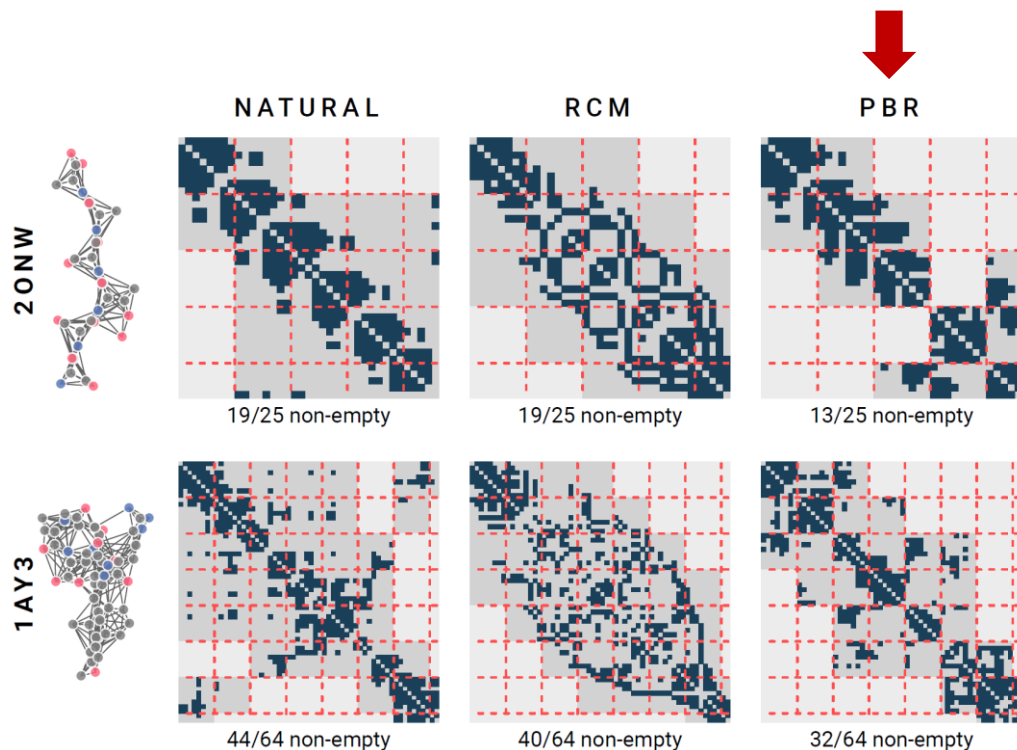
0b100000100000010001001000001001110101001001001100110000011000000

0x0303324AE4122041

Specialized graph reordering algorithm improves efficiency of OTF XMV and 2-level sparse format


Partition-based graph reordering (PBR)

- Reduces # of non-empty sparse tiles
- Improves density of non-empty tiles
- Cost easily amortized by repeated pairwise graph kernel computations.



The On-the-Fly GPU Solver Achieves Four Orders of Magnitude Speedup Over Existing Packages

- GraKeL: Cython, multi-threading
- GraphKernels: Python, no parallelization



Marginalized graph kernel has found successful applications in a variety of ML tasks

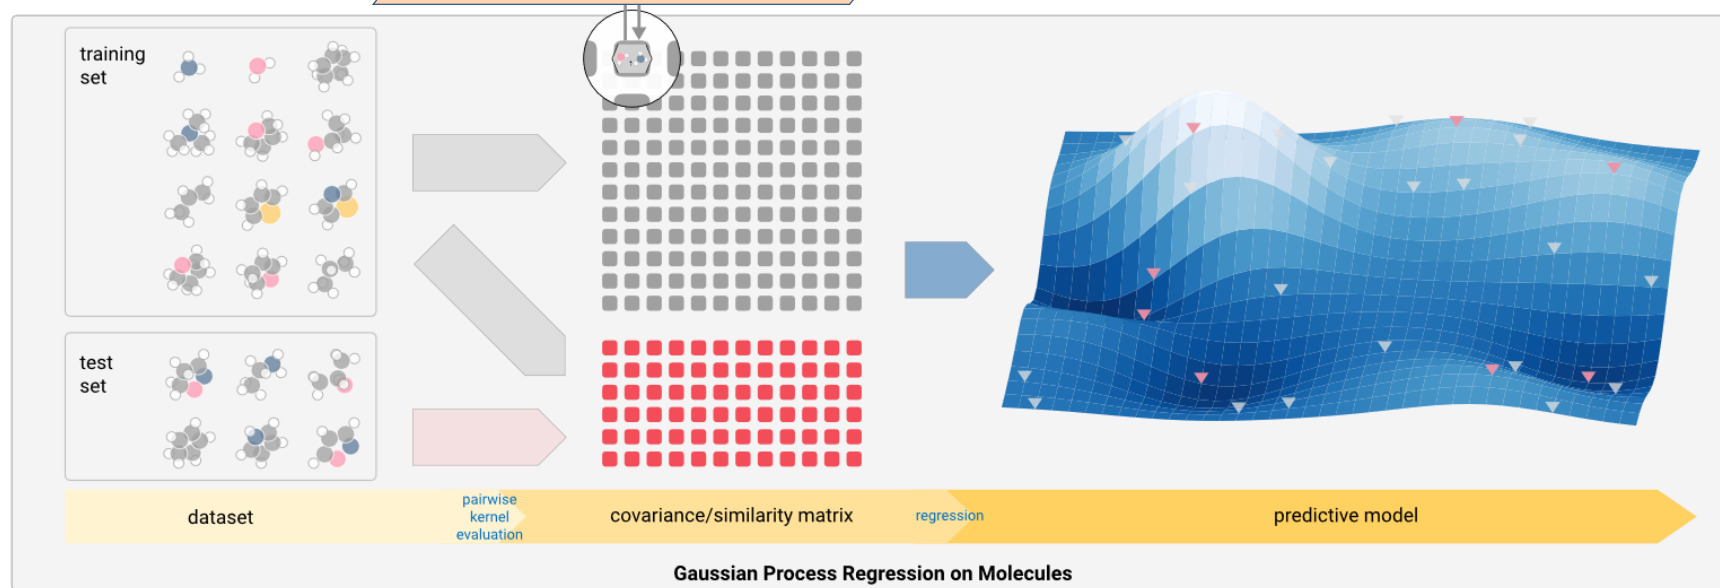
- Prediction of molecular atomization energy

Tang & de Jong, J Chem Phys, 2019

Prediction of atomization energy using graph kernel and active learning

<https://doi.org/10.1063/1.5078640>

Graph kernel evaluation



Marginalized graph kernel enables active learning of atomization energy in orders of magnitude less time than NN

- QM7: 7165 small organic molecules consisting of H, C, N, O, S, up to 23 atoms
 - From scratch training time: N = 1000: 10 s training, 0.018 s/sample predicting, N = 2000: 40 s training, 0.034 s/sample predicting

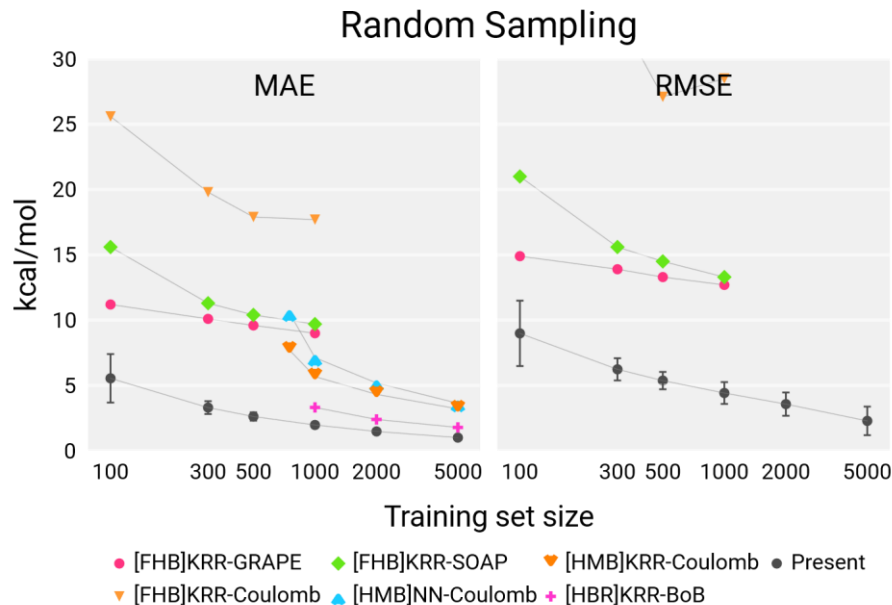
MAE: Mean Average Error

RMSE: Root-Mean Square Error

KRR: Kernel Ridge Regression

NN: Neural Network

GRAPE, SOAP, Coulomb, BoB: fingerprint algorithms



SUMMARY

Summary & Acknowledgement

Graphs are useful data structures for representing scientific datasets.

The **marginalized graph kernel** is a very generic tool for machine learning on graphs.

Marginalized graph kernel can be computed very **efficiently on GPUs**

- LBNL LDRD Project “Active Learning of Ab Initio Force Fields with Applications to Large-Scale Simulations of Materials and Biophysical Systems”
- Also supported in part by the Applied Mathematics program of the DOE Office of Advanced Scientific Computing Research under Contract No. DE-AC02-05CH11231, and in part by the Exascale Computing Project (17-SC-20-SC), a collaborative effort of the U.S. Department of Energy Office of Science and the National Nuclear Security Administration.

Thank You!

```
pip install graphdot
```

Tang, Selvitopi, Popovici, Buluc, IPDPS 2020: A High-Throughput Solver for Marginalized Graph Kernels on GPU.

<https://arxiv.org/abs/1910.06310>

Manuscript in preparation: GraphDot: A GPU-Accelerated Python Package for Graph-Based Machine Learning.