

GraphChallenge.org: Raising the Bar on Graph Analytic Performance

Siddharth Samsi, Vijay Gadepally, Michael Hurley, Michael Jones, Edward Kao, Sanjeev Mohindra, Paul Monticciolo, Albert Reuther, Steven Smith, William Song, Diane Staheli, Jeremy Kepner
MIT Lincoln Laboratory, Lexington, MA

Abstract—The rise of graph analytic systems has created a need for new ways to measure and compare the capabilities of graph processing systems. The MIT/Amazon/IEEE Graph Challenge has been developed to provide a well-defined community venue for stimulating research and highlighting innovations in graph analysis software, hardware, algorithms, and systems. GraphChallenge.org provides a wide range of pre-parsed graph data sets, graph generators, mathematically defined graph algorithms, example serial implementations in a variety of languages, and specific metrics for measuring performance. Graph Challenge 2017 received 22 submissions by 111 authors from 36 organizations. The submissions highlighted graph analytic innovations in hardware, software, algorithms, systems, and visualization. These submissions produced many comparable performance measurements that can be used for assessing the current state of the art of the field. There were numerous submissions that implemented the triangle counting challenge and resulted in over 350 distinct measurements. Analysis of these submissions show that their execution time is a strong function of the number of edges in the graph, N_e , and is typically proportional to $N_e^{4/3}$ for large values of N_e . Combining the model fits of the submissions presents a picture of the current state of the art of graph analysis, which is typically 10^8 edges processed per second for graphs with 10^8 edges. These results are 30 times faster than serial implementations commonly used by many graph analysts and underscore the importance of making these performance benefits available to the broader community. Graph Challenge provides a clear picture of current graph analysis systems and underscores the need for new innovations to achieve high performance on very large graphs.

I. INTRODUCTION

The importance of graph analysis has dramatically increased and is critical to a wide variety of domains that include the analysis of genomics [1]–[6], brain mapping [7], computer networks [8]–[11], social media [12], [13], cybersecurity [14], [15], and sparse machine learning [16]–[20].

Many graph processing systems are currently under development. These systems are exploring innovations in algorithms [21]–[31], software architecture [32]–[41], software standards [42]–[46], and parallel computing hardware [47]–[53]. The rise of graph analysis systems has created a need for new ways to measure and compare the capabilities of these systems. The MIT/Amazon/IEEE Graph Challenge has been developed

This material is based upon work supported by the Assistant Secretary of Defense for Research and Engineering under Air Force Contract No. FA8721-05-C-0002 and/or FA8702-15-D-0001. Any opinions, findings, conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the Assistant Secretary of Defense for Research and Engineering.

to provide a well-defined community venue for stimulating research and highlighting innovations in graph analysis software, hardware, algorithms, and systems. GraphChallenge.org provides a wide range of pre-parsed graph data sets, graph generators, mathematically defined graph algorithms, example serial implementations in a variety of languages, and specific metrics for measuring performance.

Scale is an important driver of the Graph Challenge and graphs with billions to trillions of edges are of keen interest. The Graph Challenge is designed to work on arbitrary graphs drawn from both real-world data sets and simulated data sets. Examples of real-world data sets include the Stanford Large Network Dataset Collection (see <http://snap.stanford.edu/data>), the AWS Public Data Sets (see aws.amazon.com/public-data-sets), and the Yahoo! Webscope Datasets (see webscope.sandbox.yahoo.com). These real-world data sets cover a wide range of applications and data sizes. While real-world data sets have many contextual benefits, synthetic data sets allow the largest possible graphs to be readily generated. Examples of synthetic data sets include Graph500, Block Two-level Erdos-Renyi graph model (BTER) [54], Pure Kronecker Graphs [55], and Perfect Power Law graphs [56], [57]. The focus of the Graph Challenge is on graph analytics. While parsing and formatting complex graph data are necessary in any graph analysis system, these data sets are made available to the community in a variety of pre-parsed formats to minimize the amount of parsing and formatting required by Graph Challenge participants. The public data are available in a variety of formats, such as linked list, tab separated, and labeled/unlabeled.

Graph Challenge 2017 received a large number of submissions that highlighted innovations in hardware, software, algorithms, systems, and visualization. These submissions produced many comparable performance measurements that can be used for assessing the current state of the art of the field. The goal of this paper is to analyze and synthesize these measurements to provide a picture of the current state of the art of graph analysis systems. The organization of this paper is as follow. First, a recap of triangle counting is provided, along with a few standard algorithms. Next, an overview is presented of the Graph Challenge 2017 submissions. The core of the paper is the section on the analysis of the 12 submission that all performed the triangle counting challenge. Based on this analysis, these results are synthesized to provide a picture of the current state of the art. Finally, the conclusions are

provided along with a discussion of further work.

II. TRIANGLE COUNTING

The Graph Challenge consists of three challenges

- Pre-challenge: PageRank pipeline [58]
- Static graph challenge: subgraph isomorphism [59]
- Streaming graph challenge: stochastic block partition [60]

The static graph challenge is further broken down into triangle counting and k-truss. Triangle counting received the most submissions and is the focus of this paper.

Triangles are the most basic, trivial sub-graph. A triangle can be defined as a set of three mutually adjacent vertices in a graph. As shown in Figure 1, the graph G contains two triangles comprising nodes $\{a,b,c\}$ and $\{b,c,d\}$. The number of triangles in a graph is an important metric used in applications such as social network mining, link classification and recommendation, cyber security, functional biology, and spam detection [61].




Fig. 1. The graph shown in this example contains two triangles consisting of nodes $\{a,b,c\}$ and $\{b,c,d\}$.

The number of triangles in a given graph G can be calculated in several ways. We highlight two algorithms based on linear algebra primitives. The first algorithm proposed by Wolf et al [62] uses an overloaded matrix multiplication approach on the adjacency and incidence matrices of the graph and is shown in Algorithm 1. The second approach proposed by Burkhardt et al [63] uses only the adjacency matrix of the given graph and is shown in Algorithm 2.

Algorithm 1: Array based implementation of triangle counting algorithm using the adjacency and incidence matrix of a graph [62].

Data: Adjacency matrix A and incidence matrix E

Result: Number of triangles in graph G

initialization;

$$C = AE$$

$$n_T = nnz(C)/3$$

Multiplication is overloaded such that

$$C(i, j) = \{i, x, y\} \text{ iff}$$

$$A(i, x) = A(i, y) = 1 \ \& \ E(x, j) = E(y, j) = 1$$

Another algorithm for triangle counting based on a masked matrix multiplication approach has been proposed by Azad et al [64]. The serial version of this algorithm based on the MapReduce implementation by Cohen et al [65] is shown in Algorithm 3. Finally, a comparison of triangle counting algorithms can be found in [66].

Algorithm 2: Array based implementation of triangle counting algorithm using only the adjacency matrix of a graph [63].

Data: Adjacency matrix A

Result: Number of triangles in graph G

initialization;

$$C = A^2 \circ A$$

$$n_T = \sum_{ij} (C)/6$$

Here, \circ denotes element-wise multiplication

Algorithm 3: Serial version of triangle counting algorithm based on MapReduce version by Cohen et al [65] and [64].

Data: Adjacency matrix A

Result: Number of triangles in graph G

initialization;

$$(L, U) \leftarrow A$$

$$B = LU$$

$$C = A \circ B$$

$$n_T = \sum_{ij} (C)/2$$

Here, \circ denotes element-wise multiplication

III. COMMUNITY SUBMISSIONS

Graph Challenge 2017 received 22 submissions by 111 authors from 36 organizations [24]–[31], [34]–[41], [48]–[53]. The submissions were judged by a panel of experts on their effectiveness at using Graph Challenge to highlight innovations in graph algorithms, hardware, software, and systems. Because of the technical expertise required to submit, submitting successfully is a significant accomplishment, and potential participants who were unable to effectively complete the challenge simply did not submit. The breakdown of the submissions amongst organizations, challenges, innovations, and awards is shown in Table I.

TABLE I
Breakdown of the Graph Challenge 2017 submissions amongst organizations, challenges, innovations, and awards.

| Organizations | | | | | |
|---------------|-----------|------------|--------------------|-------------------|------|
| Academic | Federal | Industry | | | |
| 20 | 9 | 7 | | | |
| Challenge | | | Static Kernel | | |
| Static | Streaming | PageRank | Triangle | K-Truss | Both |
| 17 | 3 | 1 | 10 | 3 | 4 |
| Innovation | | | | | |
| Hardware | Software | Algorithms | Viz | Systems | |
| 5 | 13 | 13 | 1 | 14 | |
| Award | | | | | |
| Champion | Finalist | Innovation | Student Innovation | Honorable Mention | |
| 5 | 13 | 13 | 1 | 14 | |

Numerous submissions implemented the triangle counting challenge in a comparable manner, resulting in over 350 distinct measurements of triangle counting execution time, T_{tri} . The number of edges, N_e , in the graph describes the overall size of the graph. The rate of edges processed in triangle counting is given by

$$\text{Rate} = N_e / T_{\text{tri}}$$

For any of these submissions, it is possible to extract two points: (1) the highest rate and its corresponding number of edges, and (2) the largest graph and its corresponding rate. These points are shown in Figure 2 and give a rough picture of the overall landscape of these submissions and highlight that the judges preferred submissions that showed high performance on large graphs. By analyzing all the data points, it is possible to obtain a detailed picture of the state of the art of this field.




Fig. 2. The largest graph and corresponding rate and the fastest graph and corresponding size for selected Graph Challenge 2017 triangle counting submissions. The one disk-based submission shows the importance of large, high-speed memories for this application.

IV. PERFORMANCE ANALYSIS

Analyzing and combining all the performance data from the submissions can be done by fitting a piecewise model to each submission and then comparing the models. For each submission, T_{tri} vs N_e is plotted on a log-log scale from which a model can be fit to the data by estimating the parameters α and β in the formula

$$T_{\text{tri}} = \alpha N_e^\beta$$

The triangle counting execution time vs number of edges and corresponding model fits are shown for the Champions in Figure 3, the Finalists in Figure 4, and the Honorable Mentions in Figure 5. The model fits illustrate the strong dependence of T_{tri} on N_e .

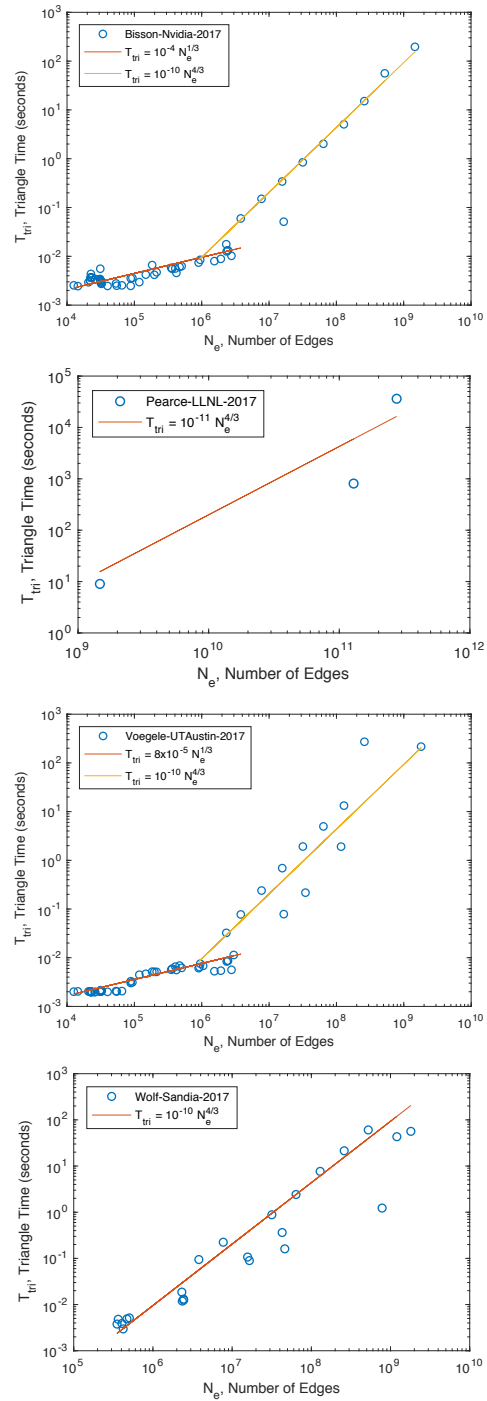


Fig. 3. Graph Challenge 2017 Champions. Triangle counting execution time vs number of edges and corresponding model fits for Bisson-Nvidia-2017 [48], Pearce-LLNL-2017 [34], Voegelé-UTAustin-2017 [24], and Wolf-Sandia-2017 [41].

V. PERFORMANCE SYNTHESIS

The model fits in Figures 3, 4, and 5 can be put in a more understandable form by transforming them to the normalized equation

$$T_{\text{tri}} = (N_e / N_1)^\beta$$




Fig. 4. Graph Challenge 2017 Finalists. Triangle counting execution time vs number of edges and corresponding model fits for Hu-GWU-2017-2017 [26], Smith-UMN-2017 [25], and Tom-UMN-2017-2017 [36].

where $N_1 = \alpha^{-1/\beta}$ is the number edges that take 1 second to process. The normalized parameters N_1 and β , along with the largest values of N_e , are shown in Table II for each submission. Submissions with larger N_e , larger N_1 , and smaller β perform best. Combined, these suggest a current state of the art performance model of

$$T_{\text{tri}} = (N_e/10^8)^{4/3}$$

and

$$\text{Rate} = \frac{10^8}{(N_e/10^8)^{1/3}}$$

The current state-of-the-art can also be seen by plotting all the model fits T_{tri} together (see Figures 6 and 7). Given the enormous diversity in processors, algorithms, and software, this relatively consistent picture of the state-of-the-art suggests that the current limitations are set by common elements across these benchmarks, such as memory bandwidth.

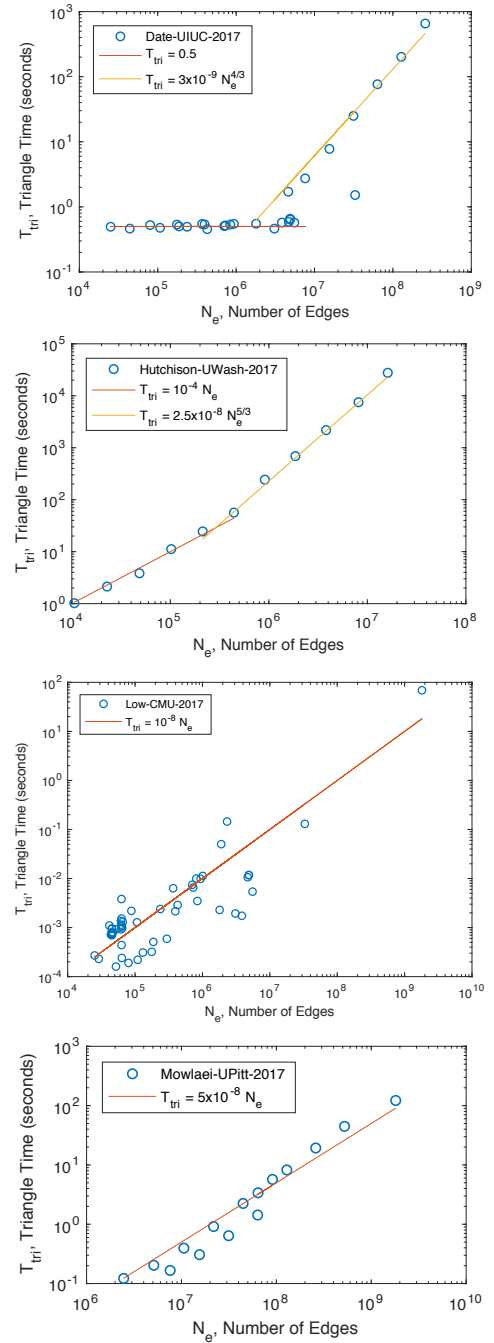


Fig. 5. Graph Challenge 2017 Honorable Mentions. Triangle counting execution time vs number of edges and corresponding model fits for Date-UIUC-2017 [49], Hutchison-UWash-2017 [40], Low-CMU-2017 [29], and Mowlaei-UPitt-2017 [31].

VI. CONCLUSION

The rapid increase in the use of graphs and has inspired new ways to measure and compare the attributes of graph analytic systems. The MIT/Amazon/IEEE Graph Challenge was created to stimulate research in graph analysis software, hardware, algorithms, and systems. The GraphChallenge.org website makes available to the world many pre-processed

TABLE II

Triangle counting time model fit coefficients for $T_{\text{tri}} = (N_e/N_1)^\beta$ for large values of N_e .

| Ref | Submission | max N_e | N_1 | β |
|------|-----------------------|----------------------|-----------------|---------|
| [48] | Bisson-Nvidia-2017 | 1.5×10^9 | 3×10^7 | 4/3 |
| [34] | Pearce-LLNL-2017 | 2.7×10^{11} | 2×10^8 | 4/3 |
| [24] | Voegelé-UTAustin-2017 | 1.8×10^9 | 3×10^7 | 4/3 |
| [41] | Wolf-Sandia-2017 | 1.8×10^9 | 3×10^7 | 4/3 |
| [26] | Hu-GWU-2017 | 3.4×10^{10} | 5×10^7 | 4/3 |
| [25] | Smith-UMN-2017 | 1.2×10^9 | 1×10^6 | 1 |
| [36] | Tom-UMN-2017 | 1.8×10^9 | 5×10^7 | 1 |
| [49] | Date-UIUC-2017 | 2.6×10^8 | 3×10^6 | 4/3 |
| [40] | Hutchison-UWash-2017 | 1.6×10^7 | 3×10^4 | 5/3 |
| [29] | Low-CMU-2017 | 1.8×10^9 | 1×10^8 | 1 |
| [31] | Mowlaei-UPitt-2017 | 1.8×10^9 | 5×10^7 | 1 |

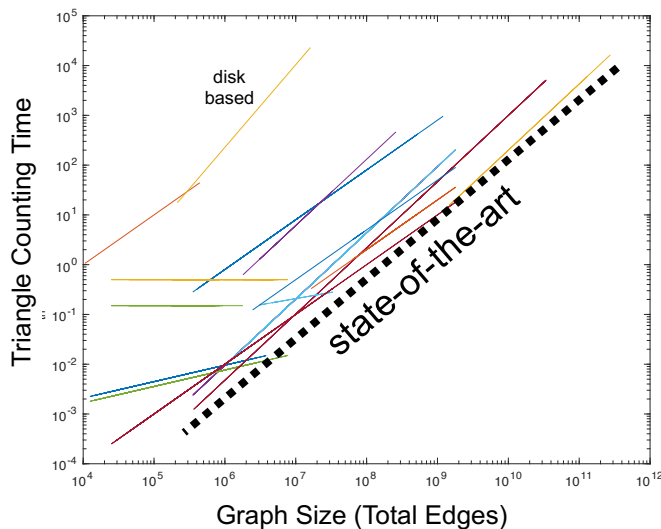


Fig. 6. Model fits of triangle execution time vs number edges for selected Graph Challenge 2017 triangle counting submissions. State-of-the-art is denoted by the black dashed line. The one disk-based submissions shows the importance of large, high-speed memories for this application.

graph data sets, graph generators, graph algorithms, prototype serial implementations in a several languages, and defined metrics for assessing performance. In 2017, Graph Challenge received 22 submissions by 111 co-authors from 36 organizations. The submissions covered new developments in hardware, software, algorithms, systems, and visualization. The comparable measurements provided by these submissions are an invaluable resource for determining the state-of-the-art in graph processing. Triangle counting was the most popular submission and produced over 350 distinct measurements of triangle execution time on a wide variety of graphs. Analysis of each of these submissions shows that execution time is a strong function of the number of edges in the graph, N_e , and is typically proportional to $N_e^{4/3}$ for large values of N_e . Combining the model fits of each of the submissions presents a picture of the current state-of-the-art of graph analysis, which is typically 10^8 edges processed per second for graphs with 10^8 edges. These results are 30 times faster than serial implementations commonly used by many graph analysts

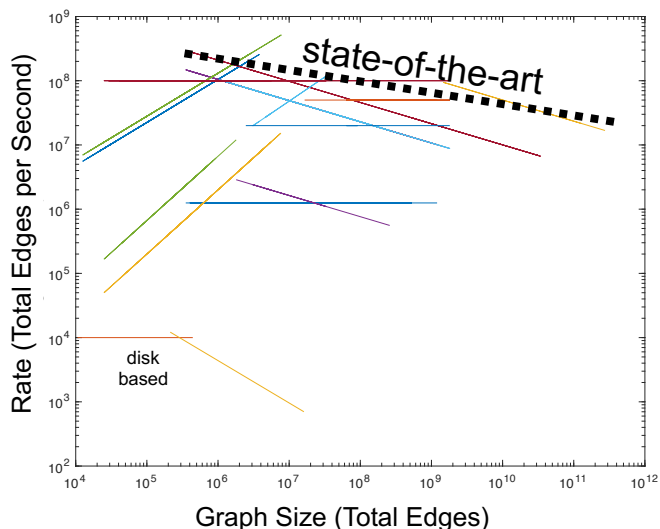


Fig. 7. Model fits of triangle execution rate vs. number edges for selected Graph Challenge 2017 triangle counting submissions. State-of-the-art is denoted by the black dashed line. The one disk-based submissions shows the importance of large, high-speed memories for this application.

and reinforces the importance of making these performance benefits available to the broader community.

Graph Challenge provides a clear picture of current graph analysis systems and underscores the need for new innovations to achieve high performance on very large graphs. In the future, Graph Challenge will examine additional data sets, algorithms, and implementations. One area of interest is sparse neural networks. Neural networks are becoming widely used in many disciplines. In some contexts, sparse neural networks are limited by their size. One approach to alleviating these size limitations is to make the representation of a neural network sparse, which may require new systems to effectively perform neural network computations.

ACKNOWLEDGMENTS

The authors wish to acknowledge the following individuals for their contributions and support: Alan Edelman, Charles Leiserson, Steve Pritchard, Michael Wright, Bob Bond, Dave Martinez, Sterling Foster, Paul Burkhardt, and Victor Roytburd.

REFERENCES

- [1] J. L. Morrison, R. Breitling, D. J. Higham, and D. R. Gilbert, “GeneRank: using search engine technology for the analysis of microarray experiments,” *BMC Bioinformatics*, vol. 6, no. 1, p. 233, 2005.
- [2] B. L. Mooney, L. R. Corrales, and A. E. Clark, “MolecularNetworks: An integrated graph theoretic and data mining tool to explore solvent organization in molecular simulation,” *Journal of Computational Chemistry*, vol. 33, no. 8, pp. 853–860, 2012.
- [3] D. Polychronopoulos, D. Sellis, and Y. Almirantis, “Conserved noncoding elements follow power-law-like distributions in several genomes as a result of genome dynamics,” *PloS one*, vol. 9, no. 5, p. e95437, 2014.
- [4] S. Dodson, D. O. Rieke, and J. Kepner, “Genetic sequence matching using D4M big data approaches,” in *High Performance Extreme Computing Conference (HPEC)*. IEEE, 2014.

- [5] S. Dodson, D. O. Rieke, J. Kepner, N. Chiu, and A. Shcherbina, "Rapid sequence identification of potential pathogens using techniques from sparse linear algebra," in *Symposium on Technologies for Homeland Security*. IEEE, 2015.
- [6] N. Gouda, Y. Shiwa, M. Akashi, H. Yoshikawa, K. Kasahara, and M. Furusawa, "Distribution of human single-nucleotide polymorphisms is approximated by the power law and represents a fractal structure," *Genes to Cells*, vol. 21, no. 5, pp. 396–407, 2016.
- [7] A. Fornito, "Graph theoretic analysis of human brain networks," *fMRI Techniques and Protocols*, pp. 283–314, 2016.
- [8] S. Brin and L. Page, "The anatomy of a large-scale hypertextual web search engine," *Computer Networks and ISDN Systems*, vol. 30, no. 1, pp. 107–117, 1998.
- [9] M. Faloutsos, P. Faloutsos, and C. Faloutsos, "On power-law relationships of the internet topology," in *ACM SIGCOMM Computer Communication Review*, vol. 29.4. ACM, 1999, pp. 251–262.
- [10] G. Yan, G. Tsekenis, B. Barzel, J.-J. Slotine, Y.-Y. Liu, and A.-L. Barabási, "Spectrum of controlling and observing complex networks," *Nature Physics*, vol. 11, no. 9, pp. 779–786, 2015.
- [11] R. Fontugne, P. Abry, K. Fukuda, D. Veitch, K. Cho, P. Borgnat, and H. Wendt, "Scaling in internet traffic: a 14 year and 3 day longitudinal study, with multiscale analyses and random projections," *IEEE/ACM Transactions on Networking*, 2017.
- [12] M. Zuckerberg, "Facebook and computer science," Harvard University CS50 guest lecture, Dec. 7 2005.
- [13] H. Kwak, C. Lee, H. Park, and S. Moon, "What is Twitter, a social network or a news media?" in *Proceedings of the 19th International Conference on World Wide Web*. ACM, 2010, pp. 591–600.
- [14] S. Shao, X. Huang, H. E. Stanley, and S. Havlin, "Percolation of localized attack on complex networks," *New Journal of Physics*, vol. 17, no. 2, p. 023049, 2015.
- [15] S. Yu, G. Gu, A. Barnawi, S. Guo, and I. Stojmenovic, "Malware propagation in large-scale networks," *IEEE Transactions on Knowledge and Data Engineering*, vol. 27, no. 1, pp. 170–179, 2015.
- [16] H. Lee, C. Ekanadham, and A. Y. Ng, "Sparse deep belief net model for visual area v2," in *Advances in neural information processing systems*, 2008, pp. 873–880.
- [17] M. Ranzato, Y.-l. Boureau, and Y. L. Cun, "Sparse feature learning for deep belief networks," in *Advances in neural information processing systems*, 2008, pp. 1185–1192.
- [18] X. Glorot, A. Bordes, and Y. Bengio, "Deep sparse rectifier neural networks," in *Aistats*, vol. 15, no. 106, 2011, p. 275.
- [19] D. Yu, F. Seide, G. Li, and L. Deng, "Exploiting sparseness in deep neural networks for large vocabulary speech recognition," in *Acoustics, Speech and Signal Processing (ICASSP), 2012 IEEE International Conference on*. IEEE, 2012, pp. 4409–4412.
- [20] J. Kepner, M. Kumar, J. Moreira, P. Pattnaik, M. Serrano, and H. Tufo, "Enabling massive deep neural networks with the GraphBLAS," in *High Performance Extreme Computing Conference (HPEC)*. IEEE, 2017.
- [21] T. H. Cormen, C. E. Leiserson, R. L. Rivest, and C. Stein, *Introduction to Algorithms*. Cambridge: MIT Press, 2009.
- [22] B. A. Miller, N. Arcolano, M. S. Beard, J. Kepner, M. C. Schmidt, N. T. Bliss, and P. J. Wolfe, "A scalable signal processing architecture for massive graph analysis," in *Acoustics, Speech and Signal Processing (ICASSP), 2012 IEEE International Conference on*. IEEE, 2012, pp. 5329–5332.
- [23] A. Buluç, G. Ballard, J. Demmel, J. Gilbert, L. Grigori, B. Lippshitz, A. Lugowski, O. Schwartz, E. Solomonik, and S. Toledo, "Communication-avoiding linear-algebraic primitives for graph analytics," in *International Parallel and Distributed Processing Symposium Workshops (IPDPSW)*. IEEE, 2014.
- [24] C. Voegelé, Y.-S. Lu, S. Pai, and K. Pingali, "Parallel triangle counting and k-truss identification using graph-centric methods," in *High Performance Extreme Computing Conference (HPEC)*. IEEE, 2017.
- [25] S. Smith, X. Liu, N. K. Ahmed, A. S. Tom, F. Petrini, and G. Karypis, "Truss decomposition on shared-memory parallel systems," in *High Performance Extreme Computing Conference (HPEC)*. IEEE, 2017.
- [26] Y. Hu, P. Kumar, G. Swope, and H. H. Huang, "Trix: Triangle counting at extreme scale," in *High Performance Extreme Computing Conference (HPEC)*. IEEE, 2017.
- [27] T. La Fond, G. Sanders, C. Klymko *et al.*, "An ensemble framework for detecting community changes in dynamic networks," in *High Performance Extreme Computing Conference (HPEC)*. IEEE, 2017.
- [28] D. Zhuzhunashvili and A. Knyazev, "Preconditioned spectral clustering for stochastic block partition streaming graph challenge (preliminary version at arxiv)," in *High Performance Extreme Computing Conference (HPEC)*. IEEE, 2017.
- [29] T. M. Low, V. N. Rao, M. Lee, D. Popovici, F. Franchetti, and S. McMillan, "First look: Linear algebra-based triangle counting without matrix multiplication," in *High Performance Extreme Computing Conference (HPEC)*. IEEE, 2017.
- [30] A. J. Uppal, G. Swope, and H. H. Huang, "Scalable stochastic block partition," in *High Performance Extreme Computing Conference (HPEC)*. IEEE, 2017.
- [31] S. Mowlaei, "Triangle counting via vectorized set intersection," in *High Performance Extreme Computing Conference (HPEC)*. IEEE, 2017.
- [32] A. Buluç and J. R. Gilbert, "The combinatorial BLAS: Design, implementation, and applications," *The International Journal of High Performance Computing Applications*, vol. 25, no. 4, pp. 496–509, 2011.
- [33] J. Kepner, W. Arcand, W. Bergeron, N. Bliss, R. Bond, C. Byun, G. Condon, K. Gregson, M. Hubbell, J. Kurz, A. McCabe, P. Michaleas, A. Prout, A. Reuther, A. Rosa, and C. Yee, "Dynamic Distributed Dimensional Data Model (D4M) database and computation system," in *2012 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2012, pp. 5349–5352.
- [34] R. Pearce, "Triangle counting for scale-free graphs at scale in distributed memory," in *High Performance Extreme Computing Conference (HPEC)*. IEEE, 2017.
- [35] M. Halappanavar, H. Lu, A. Kalyanaraman, and A. Tumeo, "Scalable static and dynamic community detection using grappolo," in *High Performance Extreme Computing Conference (HPEC)*. IEEE, 2017.
- [36] A. S. Tom, N. Sundaram, N. K. Ahmed, S. Smith, S. Eyerman, M. Kodyath, I. Hur, F. Petrini, and G. Karypis, "Exploring optimizations on shared-memory platforms for parallel triangle counting algorithms," in *High Performance Extreme Computing Conference (HPEC)*. IEEE, 2017.
- [37] O. Green, J. Fox, E. Kim, F. Busato, N. Bombieri, K. Lakhota, S. Zhou, S. Singapura, H. Zeng, R. Kannan *et al.*, "Quickly finding a truss in a haystack," in *High Performance Extreme Computing Conference (HPEC)*. IEEE, 2017.
- [38] H. Kabir and K. Madduri, "Parallel k-truss decomposition on multicore systems," in *High Performance Extreme Computing Conference (HPEC)*. IEEE, 2017.
- [39] S. Zhou, K. Lakhota, S. G. Singapura, H. Zeng, R. Kannan, V. K. Prasanna, J. Fox, E. Kim, O. Green, and D. A. Bader, "Design and implementation of parallel pagerank on multicore platforms," in *High Performance Extreme Computing Conference (HPEC)*. IEEE, 2017.
- [40] D. Hutchison, "Distributed triangle counting in the graphulo matrix math library," in *High Performance Extreme Computing Conference (HPEC)*. IEEE, 2017.
- [41] M. M. Wolf, M. Deveci, J. W. Berry, S. D. Hammond, and S. Rajamanickam, "Fast linear algebra-based triangle counting with kokkoskernels," in *High Performance Extreme Computing Conference (HPEC)*. IEEE, 2017.
- [42] T. Mattson, D. Bader, J. Berry, A. Buluç, J. Dongarra, C. Faloutsos, J. Feo, J. Gilbert, J. Gonzalez, B. Hendrickson, J. Kepner, C. Leiserson, A. Lumsdaine, D. Padua, S. Poole, S. Reinhardt, M. Stonebraker, S. Wallach, and A. Yoo, "Standards for graph algorithm primitives," in *High Performance Extreme Computing Conference (HPEC)*. IEEE, 2013.
- [43] J. Kepner, D. Bader, A. Buluç, J. Gilbert, T. Mattson, and H. Meyerhenke, "Graphs, matrices, and the graphblas: Seven good reasons," *Procedia Computer Science*, vol. 51, pp. 2453–2462, 2015.
- [44] J. Kepner, P. Aaltonen, D. Bader, A. Buluç, F. Franchetti, J. Gilbert, D. Hutchison, M. Kumar, A. Lumsdaine, H. Meyerhenke *et al.*, "Mathematical foundations of the graphblas," in *High Performance Extreme Computing Conference (HPEC)*. IEEE, 2016.
- [45] A. Buluç, T. Mattson, S. McMillan, J. Moreira, and C. Yang, "Design of the graphblas api for c," in *Parallel and Distributed Processing Symposium Workshops (IPDPSW), 2017 IEEE International*. IEEE, 2017, pp. 643–652.
- [46] T. Davis, "Suitesparse:graphblas," in *High Performance Extreme Computing Conference (HPEC)*. IEEE, 2017.
- [47] W. S. Song, V. Gleyzer, A. Lomakin, and J. Kepner, "Novel graph processor architecture, prototype system, and results," in *High Performance Extreme Computing Conference (HPEC)*. IEEE, 2016.

- [48] M. Bisson and M. Fatica, “Static graph challenge on gpu,” in *High Performance Extreme Computing Conference (HPEC)*. IEEE, 2017.
- [49] K. Date, K. Feng, R. Nagi, J. Xiong, N. S. Kim, and W.-M. Hwu, “Collaborative (cpu+ gpu) algorithms for triangle counting and truss decomposition on the minsky architecture: Static graph challenge: Subgraph isomorphism,” in *High Performance Extreme Computing Conference (HPEC)*. IEEE, 2017.
- [50] E. P. DeBenedictis, J. Cook, S. Srikanth, and T. M. Conte, “Superstrider associative array architecture,” in *High Performance Extreme Computing Conference (HPEC)*. IEEE, 2017.
- [51] S. Manne, B. Chin, and S. K. Reinhardt, “If you build it, will they come?” *IEEE Micro*, vol. 37, no. 6, pp. 6–12, 2017.
- [52] P. M. Kogge, “Graph analytics: Complexity, scalability, and architectures,” in *Parallel and Distributed Processing Symposium Workshops (IPDPSW), 2017 IEEE International*. IEEE, 2017, pp. 1039–1047.
- [53] R. Gioiosa, A. Tumeo, J. Yin, T. Warfel, D. Haglin, and S. Betelu, “Exploring datavortex systems for irregular applications,” in *Parallel and Distributed Processing Symposium (IPDPS), 2017 IEEE International*. IEEE, 2017, pp. 409–418.
- [54] C. Seshadhri, T. G. Kolda, and A. Pinar, “Community structure and scale-free collections of erdős-rényi graphs,” *Physical Review E*, vol. 85, no. 5, p. 056109, 2012.
- [55] J. Kepner and J. Gilbert, *Graph algorithms in the language of linear algebra*. SIAM, 2011.
- [56] J. Kepner, “Perfect power law graphs: Generation, sampling, construction and fitting,” in *SIAM Annual Meeting*, 2012.
- [57] V. Gadepally and J. Kepner, “Using a power law distribution to describe big data,” in *High Performance Extreme Computing Conference (HPEC), 2015 IEEE*. IEEE, 2015, pp. 1–5.
- [58] P. Dreher, C. Byun, C. Hill, V. Gadepally, B. Kuzmaul, and J. Kepner, “Pagerank pipeline benchmark: Proposal for a holistic system benchmark for big-data platforms,” in *Parallel and Distributed Processing Symposium Workshops, 2016 IEEE International*. IEEE, 2016, pp. 929–937.
- [59] S. Samsi, V. Gadepally, M. Hurley, M. Jones, E. Kao, S. Mohindra, P. Monticciolo, A. Reuther, S. Smith, W. Song, D. Staheli, and J. Kepner, “Static graph challenge: Subgraph isomorphism,” in *High Performance Extreme Computing Conference (HPEC)*. IEEE, 2017.
- [60] E. Kao, V. Gadepally, M. Hurley, M. Jones, J. Kepner, S. Mohindra, P. Monticciolo, A. Reuther, S. Samsi, W. Song, D. Staheli, and S. Smith, “Streaming Graph Challenge - Stochastic Block Partition,” in *High Performance Extreme Computing Conference (HPEC)*. IEEE, 2017.
- [61] A. Pavan, K. Tangwongsan, S. Tirthapura, and K.-L. Wu, “Counting and sampling triangles from a graph stream,” *Proc. VLDB Endow.*, vol. 6, no. 14, pp. 1870–1881, Sep. 2013.
- [62] M. M. Wolf, J. W. Berry, and D. T. Stark, “A task-based linear algebra building blocks approach for scalable graph analytics,” in *2015 IEEE High Performance Extreme Computing Conference (HPEC)*, Sept 2015, pp. 1–6.
- [63] P. Burkhardt, “Graphing trillions of triangles,” *Information Visualization*, vol. 0, no. 0, p. 1473871616666393, 2016.
- [64] A. Azad, A. Buluç, and J. Gilbert, “Parallel triangle counting and enumeration using matrix algebra,” in *Proceedings of the 2015 IEEE International Parallel and Distributed Processing Symposium Workshop*, ser. IPDPSW '15. Washington, DC, USA: IEEE Computer Society, 2015, pp. 804–811.
- [65] J. Cohen, “Graph twiddling in a mapreduce world,” *Computing in Science and Engg.*, vol. 11, no. 4, pp. 29–41, Jul. 2009.
- [66] L. Wang, Y. Wang, C. Yang, and J. D. Owens, “A comparative study on exact triangle counting algorithms on the gpu,” in *Proceedings of the ACM Workshop on High Performance Graph Processing*, ser. HPGP '16. New York, NY, USA: ACM, 2016, pp. 1–8.