A new record of graph enumeration enabled by parallel processing

Zhipeng Xu
Sun Yat-sen University, Stony Brook University
Guangzhou, Guangdong, P.R. China, Stony Brook, NY USA
xuzhp9@mail2.sysu.edu.cn

Xiaolong Huang
Stony Brook University
Stony Brook, NY, USA
xiaolong.huang@stonybrook.edu

Yidan Zhang
Stony Brook University
Stony Brook, NY, USA
yidan.zhang@stonybrook.edu

Yuefan Deng
Stony Brook University
Stony Brook, NY, USA
yuefan.deng@stonybrook.edu

Analysis of regular graphs for their properties, including eigenspectra and automorphisms, is a fertile field for discovering and applications in algebraic graph theory [2]. In the area of unlabelled regular graph counting, for 3-regular graphs of order n, Robinson and Wormald [9] presented all counting results for $n \leq 40$, while pointing out that enumeration for unlabeled k-regular graphs with $k > 3$ is an unsolved problem. Meringer [4] proposed a practical method to construct regular graphs without pairwise isomorphism checking but with a fast test for canonicity. Kimberley [8] contributed many results (A068934) to these databases by a package called GENREG developed by Meringer [4]. In addition to its challenges of pure mathematics, the graph counting problem is the root of topics in reliability, artificial intelligence, reasoning, statistical physics [10], life sciences, chemistry [5] and even the search for the origins of life [6].

Method & Results: GENREG, efficient for small-scale clusters due to its feature of task partition, approaches a hard wall of speedup for fine-grained partitioning on large-scale clusters, caused mainly by load imbalance. For obtaining larger graphs, we extend GENREG for distributed clusters by using the message passing interface (MPI). Using the parallel GENREG we developed, we have obtained the following results:

(1) a series of optimal base graphs that can be used as interconnection networks (see [12]),
(2) the exact counts of 4-regular graphs of order 23 that is 429,668,180,677,439 by using the three supercomputer clusters located in the US, China, and Ecuador (see [11]).

Among our results, the first [12] has been applied to forming in interconnection networks [1] for studying benchmark relationship of the graph ASPLs to network performance latencies. The second expands n from 22 to 23, for the first time, in the sequence A006820 [7] of OEIS, which is the number of connected 4-regular graphs of order n. Kimberley [7] used GENREG to enumerate the 4-regular graphs for up to the order 22 in 2011 [3]. This record for $n = 23$ remained unchallenged until our enumeration for $n = 23$, enabled by our parallel computing implementation to advance it a step.

Our work to obtain the new enumeration for $n = 23$ is estimated to cost nearly 100 core-years. We ganged three supercomputers, the SeaWulf at Stony Brook University can process 178,000 graphs per second per core, while the Tianhe-1 with Intel Xeon X5670 processors and the IBM Quinde 1 with Power8 processors can process 113,000 and 56,469 graphs per second per core, respectively.

Conclusions: Using three supercomputers, we broke a record set in 2011, in the enumeration of non-isomorphic regular graphs by expanding the sequence of A006820 in Online Encyclopedia of Integer Sequences (OEIS), to achieve the number for 4-regular graphs of order 23 as 429,668,180,677,439, while discovering a series of optimal base graphs that can be used as interconnection networks for parallel computers.

REFERENCES

[1] Yuefan Deng, Meng Guo, Alexandre F. Ramos, Xiaolong Huang, Zhipeng Xu, and Weifeng Liu. 2019. Optimal Low-Latency Network Topologies for Cluster Performance Enhancement. *arxiv* (2019). arXiv:cs.NI/http://arxiv.org/abs/1904.00513v1
[2] Chris Godsil and Gordon Royle. 2001. *Algebraic Graph Theory*. Springer New York. https://doi.org/10.1007/978-1-4613-0163-9
[3] F. Larrión, M.A. Pizaña, and R. Villarroel-Flores. 2016. On Self-clique Shoal Graphs. *Discrete Appl. Math.* 205, C (May 2016), 86–100. https://doi.org/10.1016/j.dam.2016.01.013
[4] Markus Meringer. 1999. Fast generation of regular graphs and construction of cages. *Journal of Graph Theory* 30, 2 (feb 1999), 137–146. https://doi.org/10.1002/(sici)1097-0118(199902)30:2<137::aid-jgt7>3.0.co;2-g
[5] Markus Meringer. 2010. Structure Enumeration and Sampling. In *Handbook of Chemoinformatics Algorithms*. Chapman and Hall/CRC, 233–267. https://doi.org/10.1201/9781420082999-c8
[6] Markus Meringer and H. James Cleaves. 2017. Exploring astrobiology using in silico molecular structure generation. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 375, 2109 (nov 2017), 20160344. https://doi.org/10.1098/rsta.2016.0344
[7] OEIS Foundation Inc. 2019. A006820 in the on-line encyclopaedia of integer sequences. http://oeis.org/A006820
[8] OEIS Foundation Inc. 2019. A068934 in the on-line encyclopaedia of integer sequences. http://oeis.org/wiki/User:Jason_Kimberley/A068934
[9] R. W. Robinson and N. C. Wormald. 1983. Numbers of cubic graphs. *Journal of Graph Theory* 7, 4 (1983), 463–467. https://doi.org/10.1002/jgt.3190070412
[10] Salil P. Vadhan. 2001. The Complexity of Counting in Sparse, Regular, and Planar Graphs. *SIAM J. Comput.* 31, 2 (jan 2001), 398–427. https://doi.org/10.1137/s0097539797321602
[11] Zhipeng Xu, Xiaolong Huang, Fabian Jimenez, and Yuefan Deng. 2019. A new record of enumeration of regular graphs by parallel processing. (2019). arXiv:cs.DM/1907.12455v2
[12] Yidan Zhang, Xiaolong Huang, Zhipeng Xu, and Yuefan Deng. 2019. A Structured Table of Graphs with Symmetries and Other Special Properties. (2019). arXiv:cs.DM/http://arxiv.org/abs/1910.13539v3