

Example Quote from Citations



Rajiv Gupta

UC Riverside

ACM Fellow, IEEE Fellow

Medusa [30] is a generalized GPU-based graph processing framework that focuses on abstractions for easy programming and scaling to multiple GPUs. CuSha primarily focuses on exploring new graph representations to allow faster

Khorasani, Farzad, Keval Vora, Rajiv Gupta, and Laxmi N. Bhuyan. "CuSha: vertex-centric graph processing on GPUs." **HPDC, 2014.**



John D. Owens

UC Davis

Medusa [2] was the pioneering mGPU graph library, taking a more general approach. It partitions the graph using Metis [11], makes replications for neighbor vertices within n hops, and updates vertex-associated values every n iterations. Their framework is limited in algorithm generality,

Pan, Yuechao, Yangzihao Wang, Yuduo Wu, Carl Yang, and John D. Owens. "Multi-GPU graph analytics." **IPDPS, 2017.**

Example Quote from Citations



Haibo Chen

Shanghai Jiao Tong Univ.

to preserve access locality for graph data. **Medusa [48]** provides users with a simple interface to write graph-parallel code on GPUs. Such techniques should be useful when extending Polymer for disk-based processing, CPU-GPU co-processing and streaming processing. However, none of them focus on leveraging NUMA character-

Zhang, Kaiyuan, Rong Chen, and Haibo Chen. "NUMA-aware graph-structured analytics." **PPoPP, 2015.**



Dick Epema

Delft University of Technology

Medusa [13] is a graph-processing framework designed to help programmers use the GPU computing power with writing only sequential code. To achieve this goal, Medusa provides a set of user-defined APIs to hide the GPU programming details. Medusa can support multiple GPUs. Medusa extends

Guo, Yong, Ana Lucia Varbanescu, Alexandru Iosup, and Dick Epema. "An empirical performance evaluation of gpu-enabled graph-processing systems." **CCGRid, 2015.**

Example Quote from Citations



Rajkumar Buyya
University of Melbourne

processing units (GPU) alongside CPU for computation (Zhang et al. 2015). Medusa (Zhong and He 2013a), for instance, was developed to make processing graphs using GPUs easier. Medusa is a programming framework that enables users to write C/C++ APIs to promote the capabilities

Heidari, Safiollah, Yogesh Simmhan, Rodrigo N. Calheiros, and Rajkumar Buyya. "Scalable graph processing frameworks: A taxonomy and open challenges." **CSUR, 2018**.



Xuanhua Shi
Huazhong University of
Science and Technology

large graphs. Medusa [37] is a general purpose GPU-based graph processing framework that provides high-level APIs for easy programming and scales to multiple GPUs. The

Zhong, Wenyong, Jianhua Sun, Hao Chen, Jun Xiao, Zhiwen Chen, Chang Cheng, and Xuanhua Shi. "Optimizing graph processing on gpus." **IPDPS, 2016**.