Algorithms for Streaming Graphs PDF - herunterladen, lesen sie

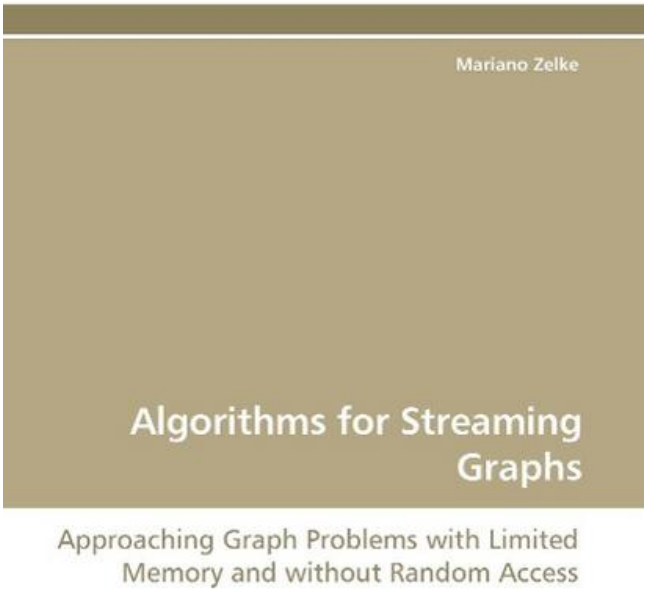

HERUNTERLADEN

LESEN

ENGLISH VERSION

DOWNLOAD

READ

Beschreibung

An algorithm solving a graph problem is usually expected to have fast random access to the input graph G and a working memory being able to store G completely. These powerful assumptions are put in question by massive graphs that exceed common working memories and that can only be stored on disks or even tapes. Here, random access is very time-consuming. To tackle massive graphs stored on external memories, the semi-streaming model has been proposed. It permits a working memory of restricted size and forbids random access to G. In contrast, the input is assumed to be a stream of edges in arbitrary order. In this book we develop algorithms in the semi-streaming model approaching different graph problems. For the problems of testing graph connectivity and bipartiteness and for the computation of a minimum spanning tree, we show how to obtain optimal running times. For the intractable problem of finding a maximum weighted matching, we present the best known approximation algorithm. Finally, we show the minimum and the maximum cut problem in a graph both to be intractable in the semi-streaming model and give algorithms that approximate respective solutions in a randomized fashion.

Figure 2: Online rectification of errors. 3.3 Consistency. Fine-grained streaming graph updates and computation on fine- grained snapshots are difficult to achieve at the same time. Graph algorithms typically require global snapshots periodically. Treating graph data as mutable state will delay updates until the computation.
Abstract— In this paper, we present a new streaming model for Graph-parallel community detection in dynamic social network using Spark GraphX tools on clouds. Two graph algorithms: SLP (streaming label propagation) and SGA. (streaming genetic algorithm), are streamlined for Graph- parallel execution in the SparkX.
Updating PageRank for Streaming Graphs. Jason Riedy. College of Computing. Georgia Institute of Technology. Atlanta, GA, USA jason.riedy@cc.gatech.edu. Abstract—Incremental graph algorithms can respond quickly to small changes in massive graphs by updating rather than recomputing analysis metrics. Here we use.
28 Jan 2010 . No algorithm truly needs to "fit into memory"--you can always page things in and out as needed. But you do want to avoid having the computation take unreasonably long--and global graph partitioning in the generic case is a NP-complete problem, which is "unreasonably long" for most problems that do not.
Deriving Streaming Graph Algorithms. From Static Definitions. David Ediger. Georgia Tech Research Institute. James P. Fairbanks. Georgia Tech Research Institute. Abstract—Increasing volumes of data and the desire for real- time query capability make the development of efficient stream- ing algorithms for data analytics.
Bad news: many graph problems cannot be solved (or even approximated) by a streaming algorithm in o(n2) space. — Example: distinguishing graphs with 0 triangles from those with 1 triangle. A bright spot: some simple properties can be solved in O(n*polylog(n)) space. xamples: bipartiteness, connectivity. These are.
CSE 522: Sublinear (and Streaming) Algorithms. Spring 2014. Lecture 17: Tolerant Testing, Graph Algorithms/Streaming. May 28, 2014. Lecturer: Paul Beame. Scribe: Paul Beame. 1 Tolerant Testing. So far in property testing we have consider testers of the form, with probability at least 2/3, deter- mine whether. • the input.
20 Nov 2017 . We introduce a novel algorithm to perform graph clustering in the edge streaming setting. In this model, the graph is presented as a sequence of edges that can be processed strictly once. Our streaming algorithm has an extremely low memory footprint as it stores only three integers per node and does not.
partitioned. In this work, we focus on studying streaming vertex-cut graph partitioning algorithms where partitioners receive a graph as a stream of vertices and edges and assign partitions to them on their arrival once and for all. Some of these algorithms maintain a state during partitioning. In some cases, the size of the state.
Section 2 presents some research contributions that are highly related to our work. Algorithms for continuous similarity computation over streaming graphs are given in Section 3.

Performance evaluation results based on real-world as well as synthetic networks are offered in Section 4. Finally, Section 5 concludes our work.
18 Dec 2017 . Request (PDF) | Clustering Streaming. | In this paper, we propose techniques for clustering large-scale "streaming" graphs where the updates to a graph are given in form of a stream of vertex or edge additions and deletions. Our algorithm handles such updates in an online and incremental manner and it.
Frequent Subgraph Discovery in Large Attributed Streaming Graphs basic idea behind it was to join two size k frequent graphs to generate size (k + 1) graph candidates, and then check the frequency of these candidates separately. The algorithm. FSG proposed by Kuramochi and Karypis (Kuramochi and Karypis, 2004).
mining streaming graph transactions for closed frequent subgraphs using a core set of closed frequent subgraphs as a compressed representation of all the closed frequent subgraphs discovered in the past. Aggarwal et al. (Aggarwal et al., 2010) propose two algorithms for finding dense subgraphs in large graphs with.
ABSTRACT. Due to a large number of applications, the problem of es- timating the number of triangles in graphs revealed as a stream of edges, and the closely related problem of estimating the graph's clustering coefficient, have received consider- able attention in the last decade. Both efficient algorithms.
I have gone through many algorithms including streaming k-means, CluStream etc and they all have their pros and cons. What is the best performing algorithm in . or too trivial a question for this community. What is the best data structure to store a time-windowed streaming graph in order to compute fast statistics over .
11 Nov 2014 . Note here in all streaming algorithms, we want to minimize space, and then update time, even through multiple passes. The accuracy of the algorithm is often defined as an (ϵ, δ)-approximation. It means the algorithm achieves an error of less than ϵ with probability $1 - \delta$. 2.1 Lower Bounds in Graph.
An algorithm solving a graph problem is usually expected to have fast random access to the input graph G and a working memory being able to store G completely. These powerful assumptions are put in question by massive graphs that exceed common working memories and that can only be stored on disks or even tapes.
So far we have studied streaming algorithms where the goal has been to obtain $(1+\epsilon)$-approximate solution in polylogarithmic space (or at least sublinear space in input size n and m) in a single pass. We now introduce graph streaming algorithms where the stream comes in the form of edges of the graph. For graph.
We formalize a potentially rich new streaming model, the semi-streaming model, that we believe is necessary for the fruitful study of efficient algorithms for solving problems on massive graphs whose edge sets cannot be stored in memory. In this model, the input graph, G = (V, E), is presented as a stream of edges (in.
As a result, it is essential to develop efficient incremental algorithms for streaming graph data. In this paper, we propose the first incremental k-core decomposition algorithms for streaming graph data. These algorithms locate a small subgraph that is guaranteed to contain the list of vertices whose maximum k-core values.
logical networks can, already today, easily exceed several hundreds of millions of nodes), so that the adjacency matrix A of the corresponding graph can become difficult to manipulate and store. We revisit network clustering problems under memory constraints. Memory limited algorithms are relevant in the streaming data.
Data stream algorithms: Finding Maximum; Counting distinct elements. Graph Stream algorithms: Insert-only streams- spanners; Sliding window- connectivity. 2. Part 1. What is a

streaming algorithm? 3. Data streams. A data stream is a (very big) amount of data that cannot be stored. Most algorithms usually require for the.
Graph Streaming Models. In the updatable edge-streaming model of graphs, the stream is viewed as a sequence of tuples of the form (u,v,+) or (u,v,−), corresponding, respectively, to the insertion or the deletion of the edge (u,v). In the updatable model, once an edge (u,v) is inserted, it remains current in the graph until a.
29 Apr 2014 . game of communication between Alice and Bob in which both parties attempt to solve the s–t reachability problem of a 3-layer graph using small memory. The solution to this problem helps us construct a streaming algorithm for topological sorting. Unfortunately, using the notion of additive combinatorics.
28 Jun 2012 . Streaming Relevance. • Computer communication networks link entities. • Represent relationships with a graph. Cyber traffic/activity is a stream through time. • Stream are huge. – Humans cannot keep up. – Gap will only increase. • It's time to develop fundamental streaming graph algorithms to partially.
The aim of this paper is to find a general streaming algorithm, which will carry out well-compressed results relying only on structures of graphs. An algorithm named as STT is introduced. Characteristics of STT algorithm are as follows. – It's able to compress graph streams without occupying extra storage. – It achieves both.
Intractability of Min- and Max-Cut in Streaming Graphs. Mariano Zelke. Institut für Informatik, Goethe-Universität, Frankfurt am Main, Germany. Abstract. We show that the exact computation of a minimum or a maximum cut of a given graph G is out of reach for any one-pass streaming algorithm, that is, for any algorithm that.
General Terms: Algorithms, Design, Experimentation, Performance. Additional Key Words and Phrases: Network sampling, social network analysis, graph streams, relational classification. ACM Reference Format: Ahmed, N.K., Neville, J., and Kompella, R. 2013. Network Sampling: From Static to Streaming Graphs.
As the computer science community increases its development of algorithms and codes for large-scale graph problems, no canonical graph representation has yet to emerge. Without a standard graph representation, algorithms that are implemented for one framework may require substantial programming efforts to port to a.
After attending this course, the participants will be familiar with the state of the art of streaming algorithms/complexities and will have an overview of research . Kannan, A. Mcgregor, S. Suri, and J. Zhang; [Bas08] Streaming Algorithm for Graph Spanners - Single Pass and Constant Processing Time per Edge by S. Baswana.
ABSTRACT. Over the last decade, there has been considerable in- terest in designingalgorithms for processing massive graphs in the data stream model. Theoriginal moti- vation was two-fold: a) in many applications, the dy- namic graphs that arise are too large to be stored in the main memory of asingle machine andb).
8 Dec 2017 . We introduce a novel algorithm to perform graph clustering in the edge streaming setting. In this model, the graph is presented as a sequence of edges that can be processed strictly once. Our streaming algorithm has an extremely low memory footprint as it stores only three integers per node and does not.
This paper discusses a method to derive algorithms for streaming graph analysis from static formulations Combining tuned graph algorithms building blocks with an appropriate functional language, a graph query planner should be able to correctly implement most static and streaming versions of an algorithm from a single.
graph is rapidly changing and received as a real-time stream of nodes and arcs. In this paper we will review current graph reachability algorithms and focus on how they can be adapted to

the streaming setting. We will also outline a new algorithm for answering reachability queries on huge, rapidly changing graphs.
A fast streaming spanner algorithm for incrementally constructing sparse roadmaps. Weifu Wang. Devin Balkcom. Amit Chakrabarti. Abstract— Sampling-based probabilistic roadmap algorithms such as PRM and PRM* have been shown to be effective at solving certain motion planning problems, but the large graphs.
Streaming problems and techniques. – Estimating number of distinct elements in a stream. – Other quantities: Norms, moments, heavy hitters… – What else ? Geometry, graphs, text,… • Streaming and sparse approximations. – Connections (compressive sensing, coding theory). – New developments. • 1-2 proofs, 3 2 open.
Streaming Balanced Graph Partitioning Algorithms for Random. Graphs. Isabelle Stanton. Google Inc stanton@google.com. Abstract. With recent advances in storage technology, it is now possible to store the vast amounts of data generated by cloud computing applications. The sheer size of 'big data' motivates the need for.
Abstract. Considerable effort has been devoted to the development of streaming algorithms for analyzing mas- sive graphs. Unfortunately, many results have been negative, establishing that a wide variety of problems require Ω(n2) space to solve. One of the few bright spots has been the development of semi-streaming.
incremental algorithms for streaming graph data. In this paper, we propose the first incremental k-core decomposition algorithms for streaming graph data. These algorithms locate a small subgraph that is guaranteed to contain the list of vertices whose maximum k-core values have to be updated, and efficiently process this.
18 Jan 2010 . Abstract. Streaming is an important paradigm for handling massive graphs that are too large to fit in the main memory. In the streaming computational model, algorithms are restricted to use much less space than they would need to store the input. Furthermore, the input is accessed in a sequential fashion,.
Abstract: Analysis of social networks is challenging due to the rapid changes of its members and their relationships. For many cases it impractical to recompute the metric of interest, therefore, streaming algorithms are used to reduce the total runtime following modifications to the graph. Centrality is often used for determining.
24 Oct 2017 . Computer programs that perform these kinds of on-the-go calculations are called streaming algorithms. Because data comes at .. The important feature of an expander graph is that instead of merely connecting each point with its adjoining blocks, you connect each two-digit block with multiple other blocks.
Abstract. We consider the problem of estimating the size of a maximum matching when the edges are revealed in a streaming fashion. When the input graph is planar, we present a simple and elegant streaming algorithm that with high probability estimates the size of a maximum matching within a constant factor using.
earlier approximate algorithm. Index Terms— social network analysis, streaming data, graph analysis, parallel processing. 1. INTRODUCTION. Applications ranging from business intelligence and finance to computational biology and computer security are generating data at a massive rate. Social networks such as those.
Algorithms, Performance. Keywords triangle counting; streaming algorithm; parallel algorithm; parallel cache oblivious (PCO); massive graphs. 1. INTRODUCTION. The number of triangles in a graph is an important metric in so- cial network analysis [30, 21], identifying thematic structures of networks [10], spam and fraud.
README.md. Gelly Streaming. An experiemental API for single-pass graph streaming analytics on Apache Flink. A Graph Streaming Model. We implement a light-weight

distributed graph streaming model for online processing of graph statistics, improving aggregates, approximations, one-pass algorithms and graph.
finedbyastreamofdata. For example, the streamcould consist of the edges of the graph. Algorithms in this model must process the input stream in theorder it ar- . the semi-streaming model [27,52]. In this model the. datastreamalgorithmispermittedO(npolylogn)space where n is the number of nodes in the graph. Thisis.
13 Nov 2008 . Algorithms for Streaming Graphs. DISSERTATION zur Erlangung des akademischen Grades doctor rerum naturalium. (Dr. rer. nat.) im Fach Informatik eingereicht an der. Mathematisch-Naturwissenschaftlichen Fakultät II. Humboldt-Universität zu Berlin von. Herr Dipl.-Inf. Mariano Zelke geboren am.
Abstract. We formalize a potentially rich new streaming model, the semi-streaming model, that we believe is necessary for the fruitful study of efficient algorithms for solving problems on massive graphs whose edge sets cannot be stored in memory. In this model, the input graph, G = (V,E), is presented as a stream of edges.
Mathematics]: Graph Theory—Graph algorithms. General Terms: Algorithms, Theory. Additional Key Words and Phrases: Graph spanners, streaming algorithms. ACM Reference Format: Elkin, M. 2011. Streaming and fully dynamic centralized algorithms for constructing and maintaining sparse spanners. ACM Trans. Algor.
(Insert-Only Streams) We present a one-pass algorithm that takes O(α log n) space and approximates the size of the maximum matching in graphs with arboricity α within a factor of O(α). This improves significantly upon the state-of-the-art ˜O(αn2/3)-space streaming algorithms, and is the first poly-logarithmic space.
whole graph. The partitioning algorithm we present manages to create partitions from streaming graphs with low memory usage, but can also adapt partitions overtime based on different application needs such as minimizing cross-partition edges, balancing load across partitions, elastically adapting partitions based on a.
Streaming Algorithm for Graph Spanners - Single Pass and Constant. Processing Time per Edge. Surender Baswana ∗. Department of Computer Science & Engineering. Indian Institute of Technology, Kanpur - 208016, INDIA. Email : sbaswana@cse.iitk.ac.in. Keywords : Analysis of algorithms, On-line algo- rithms.
contains 500 million objects and 3.5 billion relation- ships [4]. In the massive streaming data analytics model, we view the graph as an infinite stream of edge insertions, deletions, and updates. Keeping complex analytics up to date at these high rates is a challenge that requires new algorithms that exploit opportunities for.
19 May 2014 . Outline. Motivation: Graph Algorithms for Analysis. Graphs and Streaming Data. STING/STINGER Analysis Framework. Building Blocks for Streaming Graph Data. PageRank. Triangle Counting. Agglomerative Communities. Observations. Riedy, Bader— STINGER: Building Blocks. 19 May 2014. 3 / 24.
Available online 28 October 2010. Communicated by F.Y.L. Chin. Keywords: Graph algorithms. Streaming algorithms. Intractability. Min-cut. Max-cut. We show that the exact computation of a minimum or a maximum cut of a given graph G is out of reach for any one-pass streaming algorithm, that is, for any algorithm that runs.
There has been a great deal of recent interest in the streaming model of computation where algorithms are restricted to a single pass over the data and have significantly less internal memory available than would be required to store the entire stream of data. Because of the inherent difficulty of solving graph problems in the.
Abstract. Estimating the size of the maximum matching is a canonical problem in graph analysis, and one that has attracted extensive study over a range of different computational

models. We present improved streaming algorithms for approximating the size of maximum matching with sparse. (bounded arboricity) graphs.
ABSTRACT. This paper presents a distributed, streaming graph parti- tioner, Graph Streaming Partitioner (GraSP), which makes partition decisions as each vertex is read from memory, simulating an online algorithm that must process nodes as they arrive. GraSP is a lightweight high-performance comput- ing (HPC) library.
Streaming Algorithms for Estimating the Matching Size in Planar Graphs and Beyond. Hossein Esfandiari *†. Mohammad T Hajiaghayi *†. Vahid Liaghat *†. Morteza Monemizadeh ‡†. Krzysztof Onak §. Abstract. We consider the problem of estimating the size of a maximum matching when the edges are revealed in a.
Abstract. We present three streaming algorithms that $(\epsilon, \delta)-$ approx- imate 1 the number of triangles in graphs. Similar to the previous al- gorithms [3], the space usage of presented algorithms are inversely pro- portional to the number of triangles while, for some families of graphs, the space usage is improved. We also.
8 Nov 2016 . Google details a graph streaming algorithm for constant runtime over large graphs of varying complexity space and predictor outputs.
11 Apr 2017 . In order to handle this problem, we propose an association-oriented streaming graph partitioning method named Assc. This approach first computes the rank values of vertices with a hybrid approximate PageRank algorithm. After splitting these vertices with an adapted variant affinity propagation algorithm,.
6 May 2009 . Algorithms for streaming graphs. Zelke, Mariano. Mathematisch-Naturwissenschaftliche Fakultät II. Für einen Algorithmus zum Lösen eines Graphenproblems wird üblicherweise angenommen, dieser sei mit wahlfreiem Zugriff (random access) auf den Eingabegraphen G ausgestattet, als auch mit einem.
14 May 2014 . Streaming Algorithms graphs), e.g. how well connected the graph is. Although such statistics is not precise for most cases, this approximate information is usually used to speedup subsequent processes. Thees examples motivate the study of streaming algorithms, in which algorithms have access to the.
23. 2.7.2 Regular Bipartite Multigraphs and Edge Coloring . . . . . 24. 2.8 An $\Omega(nd)$ Lower Bound for Deterministic Algorithms . . . . . 25. 2.9 An $\Omega(n \log n)$ High Probability Lower Bound For Dense Graphs . . . . . 28. 3 Matching covers and streaming. 31. 3.1 Preliminaries .
Algorithms for Streaming Graphs: Approaching Graph Problems with Limited Memory and without Random Access - Buy Algorithms for Streaming Graphs: Approaching Graph Problems with Limited Memory and without Random Access by zelke, mariano|author only for Rs. 6099 at Flipkart.com. Only Genuine Products.
In computer science, streaming algorithms are algorithms for processing data streams in which the input is presented as a sequence of items and can be examined in only a few passes (typically just one). In most models, these algorithms have access to limited memory (generally logarithmic in the size of and/or the.
22 Oct 2017 . For most algorithms we provide two procedures, one that writes results back to the graph as node-properties and reports statistics. And another (named algo.<name>.stream) that returns a stream of data, e.g. node-ids and computed values. For large graphs the streaming procedure might return millions or.
Seamlessly work with both graphs and collections. GraphX unifies ETL, exploratory analysis, and iterative graph computation within a single system. You can view the same data as both graphs and collections, transform and join graphs with RDDs efficiently, and write custom iterative graph algorithms using the Pregel API.

The Cyber Engineering Research Institute (CERI) of Sandia National Laboratories, partnering with The Center for Discrete Mathematics and Theoretical Computer Science (DIMACS) will host an invitation-only workshop on streaming graph algorithms.
Abstract—While the algorithms for streaming graph parti- tioning are proved promising, they fall short of creating timely partitions when applied on large graphs. For example, it takes. 415 seconds for a state-of-the-art partitioner to work on a social network graph with 117 millions edges. We introduce an efficient platform for.
30 Apr 2016 . Abstract: In this paper we study graph problems in dynamic streaming model, where the input is defined by a sequence of edge insertions and deletions. As many natural problems require space, where is the number of vertices, existing works mainly focused on designing space algorithms. Although.
Reductions in Streaming Algorithms, with. an Application to Counting Triangles in. Graphs. Ziv Bar-Yossef. Computer Science Division. U.C. Berkeley. http://www.cs.berkeley.edu/ zivi. Ravi Kumar. D. Sivakumar. IBM Almaden Research Center. 1.
This thesis studies algorithms for graph problems in a data streaming setting. In its classical form, the Streaming computational model assumes that the input data are presented as a sequential read-only stream, that has to be processed in one or more passes, using a memory that is small compared to the length of the.
We extend previous work on such annotation models by considering a number of graph streaming problems. Without annotations, streaming algorithms for graph problems generally require significant memory; we show that for many standard problems, including all graph problems that can be expressed with totally.
are few existing scalable algorithms for monitor- ing complex global quantities like decomposition into community structure. Using our framework. STING, we present the first known parallel algo- rithm specifically for monitoring communities in this massive, streaming, graph-structured data. Our algorithm performs.
Streaming is an important paradigm for handling massive graphs that are too large to fit in the main memory. In the streaming computational model, algorithms are restricted to use much less space than they would need to store the input. Furthermore, the input is accessed in a sequential fashion, therefore, can be viewed as.
streaming model, such that the bound is $O(n \cdot \text{polylog } n)$. The essay describes algorithms for ap- proximations of the unweighted and weighted matching problem and gives a $o(\log^{1-\epsilon} n)$ lower bound for approximations of the diameter. Finally, some results for further graph problems are discussed. 1. Introduction .
Algorithms for Big Data: Streaming and Sublinear Time Algorithms. Moran Feldman . Algorithm. Trivial Algorithm. Query the distance between every pair of points, and return the maximum one. Time complexity: $O(P^2)$. More Involved Algorithm. Fix an arbitrary point .. Streaming Algorithms for Graph Problems. Streaming for.
12 Aug 2012 . Graph Partitioning, Streaming Algorithms, Distributed Graphs,. Experimental Evaluation. Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies.
What? This class will give you a biased sample of techniques for scalable data anslysis. Target audience are students interested in algorithms, statistics, machine learning, data mining and related areas. Who? Grigory Yaroslavtsev. When? Fall 2015, MW 10:30 – 12:00; Where? Towne 307. Need permission? Please, send.
large-scale “streaming” graphs where the updates to a graph are given in form of a stream of vertex or edge additions and deletions. Our algorithm handles such updates in an online and

incremental manner and it can be easily parallelized. Several previous graph clustering algorithms fall short of han- dling massive and.
15, 10/20, Streaming Graph Algorithms: Graph Connectivity and Spanners in the Insert-Only Model. Graph Connectivity in the Turnstile Model via L_0 Sampling. For an intro to graph streams and connectivity/spanners in the insert-only model, see Andrew's Lecture 2.1 slides. For Graph Connectivity in the turnstile model,.
In this lecture, we continue our discussion of graph streaming algorithms. In particular, we establish some space lower bounds for approximating the all-pairs shortest paths problem. We define cut and spectral sparsifiers and go over an algorithm to compute a spectral sparsifier of a graph G in a streaming environment.
Since the latter tends to be relational in nature, graphs are a natural abstraction. This motivates future research into efficient algorithms for fundamental graph problems in a high-volume, streaming environment. Algorithms designed using current theoretical models for streaming graph algorithms are not directly suitable for.
14 Jul 2017 . A New Algorithm Model for Massive-Scale Streaming Graph Analysis. 1. A New Algorithm Model for Massive-Scale Streaming Graph Analysis E. Jason Riedy, Chunxing Yin, and David A. Bader Georgia Institute of Technology SIAM Workshop on Network Science, 14 July 2017; 2. Outline Motivation and.
analysis algorithms to compute an initial metric on the graph and then a final metric on the graph after all updates. The underlying assumption is that the time window is large and the network changes substantially so that the entire metric must be recomputed. In the massive streaming data analytics model, algorithms react.
26 Apr 2015 . Abstract. We give the first treatment of the classic independent set problem in graphs and hy- pergraphs in the streaming setting. The objective is to find space-efficient algorithms that output independent sets that are “combinatorially optimal”, that is, with size guarantee in terms of the degree sequence.
cluster with real-world web and social network graphs. The results demonstrate that the proposed approach achieves significant improvement compared with the state-of-the-art solutions. Index Terms—Graph Partitioning, Streaming Algorithms, Heterogeneous Environment, BSP Model. ♢. 1 INTRODUCTION. IN recent years.
29 Oct 2017 . This best-in-class streaming algorithm works by remembering just enough of what it's seen to tell you what it's seen most frequently. .. The important feature of an expander graph is that instead of merely connecting each point with its adjoining blocks, you connect each two-digit block with multiple other.
13 May 2014 . 7. K. J. Ahn, S. Guha, and A. McGregor. Spectral sparsification of dynamic graph streams. In International Workshop on Approximation Algorithms for Combinatorial Optimization Problems, 2013. 8. M. Badoiu, A. Sidiropoulos, and V. Vaikuntanathan. Computing s-t min-cuts in a semi-streaming model.
7 Oct 2015 . This time, the main topics were around the distributed stateful streaming and graph analysis with Apache Flink. We are fortunate enough to have . All of the examples of stateful processing and algorithms mentioned before rely on one common pattern, which is a stateful operator. In the talk, the stateful.
algorithms are practical, effective and efficient in this world, and how graph processing parallel or distributed applica- tions can benefit, if at all, from streaming graph partition- ing. Keywords. Streaming, Graph Partitioning, Edge-Cuts, Vertex-Cuts. 1. INTRODUCTION. Traditional graph partitioning methods, such as METIS [7].
incremental algorithms for streaming graph data. In this paper, we propose the first incremental k-core decomposition algorithms for streaming graph data. These algorithms

locate a small subgraph that is guaranteed to contain the list of vertices whose maximum k-core values have to be updated, and efficiently process this.
15 Mar 2016 . Graph mining is an important research area with a plethora of prac- tical applications. . of core decomposition in graphs, the associated algorithms for its efficient .. Baseline algorithm. – An O(|E|) algorithm for k-core decomposition. – Streaming k-core decomposition. – Distributed k-core decomposition.
15 Feb 2012 . Outline. Motivation. Technical. Why analyze data streams? Overall streaming approach. Clustering coefficients. Connected components. Common aspects and questions. Session. SIAM PP 2012—Scalable Algorithms for Analysis of Massive, Streaming Graphs—Jason Riedy. 2/29.
size. In our work, we provide compression, streaming, and parallel algorithms for three important graph analytics problems: centrality computation, dense subgraph discovery and community detection. In addition, we introduce new dense subgraph discovery algorithms to better model the cohesion in real-world networks.
Abstract. In this paper we study graph problems in dynamic streaming model, where the input is defined by a sequence of edge insertions and deletions. As many natural problems require Ω(n) space, where n is the number of vertices, existing works mainly focused on designing ˜O(n) space algorithms. Although sublinear in.
sidered problems related to following paths in directed graphs and connectivity. Most of the work on graph streams has occurred in the last decade and focuses on the semi-streaming model [27, 52]. In this model the data stream algorithm is permitted O(n polylog n) space where n is the number of nodes in the graph. This is.