

# Bibudh Lahiri

Résumé

Revised November 2009

## Personal Data

Address: 3125 Coover Hall Phone: (515)451-0307  
Dept. of Electrical & Computer Engineering Email: bibudh@iastate.edu  
Iowa State University Web: http://home.eng.iastate.edu/~bibudh  
Ames, IA, 50011, USA

## Research Interests

- Data stream algorithms for detecting exploit patterns from network flows
- Distributed algorithms for data aggregation in P2P and wireless sensor networks

## Education

Ph.D. Candidate Comp. Engg. Iowa State University Fall 2006 - current  
Bachelor of Engineering Comp. Sc. & Engg. Jadavpur University 1998-2002

## Professional Experience

Research Assistant ECE, ISU Fall 2006 - current Ames, IA, USA  
Assistant Systems Engineer Tata Consultancy Services Limited 2002 - 2006 Kolkata, Manchester

## Publications

### Refereed Conference Publications

- Bibudh Lahiri and Srikanta Tirthapura, “**Finding Correlated Heavy-Hitters over Data Streams**”, accepted in the *28th IEEE International Performance Computing and Communications Conference (IPCCC) 2009*
- Bibudh Lahiri and Srikanta Tirthapura, “**Computing Frequent Elements using Gossip**”, *Proc. 15th International Colloquium on Structural Information and Communication Complexity (SIROCCO) 2008*, pp. 119-130

### Invited Articles

- Bibudh Lahiri and Srikanta Tirthapura, “**Stream Sampling**”, published in the *Encyclopedia of Database Systems*, by Springer Verlag GmbH

### Papers under Review

- Bibudh Lahiri, “**A Generic Framework for Detecting Top-k Items from a Stream**”, submitted to the *ACM Student Research Competition 2010*
- Bibudh Lahiri and Srikanta Tirthapura, “**Computing Frequent Items in a Network using Gossip**”, submitted to the *Journal of Parellel and Distributed Computing (JPDC)*

## Papers under Preparation

- Bibudh Lahiri, Yang Liu, Srikanta Tirthapura and Yong Guan, “**Detecting Anomalous N-grams from Network Payloads by Asymmetric Combinatorial Group Testing**”
- Bibudh Lahiri, “**Identifying Heavy-Hitters from Data Streams with Time-Decay**”

## Talks

- “Finding Correlated Heavy-Hitters over Data Streams”, The 28th IEEE International Performance Computing and Communications Conference (IPCCC), to be held in Phoenix, AZ in December 2009
- “Computing Frequent Elements using Gossip”, The 15th International Colloquium on Structural Information and Communication Complexity (SIROCCO), Villars-sur-Ollon, Switzerland, 2008

## Honors and Achievements

Student travel grant for IMC 2009	ACM/USENIX	2009
Student travel grant for ACM PODC 2007	ACM	2007
8th in class	Bachelor of Comp. Sc. & Engg., Jadavpur University	2002
54th among around 60,000 candidates	West Bengal Engineering Entrance Examination	1998
48th among around 300,000 candidates	West Bengal Higher Secondary Examination	1998
29th among around 500,000 candidates	West Bengal Secondary Examination	1996

## Activities

- Editorial Board Member, Springer
- Student member, ACM & IEEE
- **Attending IPCCC 2009 in Phoenix, AZ in December 2009.**
- **Attended IMC 2009 at Chicago, IL in November 2009 and SIROCCO 2008 at Villars-sur-Ollon, Switzerland in June 2008.**
- Reviewer in INFOCOM '10, SIROCCO '09, ICDCN '09, LCN '08, DCOSS '08 and ICDCN '08
- **Attended PODC 2007 in Portland, Oregon in August 2007.**

## Selected Research and Development Projects

1. **Designing a Network Traffic Stream Monitoring System: ECE, ISU. Started Fall 2009 (in collaboration with Sam Ellis & Paul Bye from IBM Rochester):** Our goal is to develop a network monitoring application, utilizing the IBM Infosphere Streams middleware. The agendas are (1) to identify a set of related algorithmic problems arising in network traffic monitoring, (2) to survey the existing literature for efficient approximation algorithms designed for these problems and (3) to investigate what additional operators and aggregate queries these algorithms demand on top of Infosphere Streams.
2. **Detecting exploit patterns from network flow streams: ECE, ISU. Summer 2008 - current (with Srikanta Tirthapura, Yong Guan and Yang Liu):** Research on network traffic measurement has identified various patterns that the typical exploits on today’s Internet exhibit. The goal of this research is to devise single-pass (online) data stream algorithms for detecting these patterns from network traffic flow data, using a workspace much smaller than the size of the traffic flow, offering provable guarantees on the space and time requirements and the degree of approximation in the estimates returned.

**Roles played:** Identified a set of related problems. Currently working on the design and simulation of *random projection* sketches that can learn the traffic profiles for a site/network from training and test datasets, and can be compared using *combinatorial group testing* for detecting anomalous payloads from the test dataset.

3. **Design of gossip-based protocols for data aggregation in P2P networks: ECE, ISU. Fall 2007 - Spring 2009 (with Srikanta Tirthapura):** The objective was to design distributed *gossip*-based algorithms for computing various aggregates in large-scale P2P networks, with probabilistic guarantees on accuracy.

**Roles played:** Identified the problem, worked on the design and the probabilistic analyses of the algorithms (with Srikanta Tirthapura) and simulated the algorithms in Java

4. **A peer-to-peer file-sharing system: ECE, ISU. Fall 2007:** This application, developed using socket programming in Java, supported file sharing across multiple machines. The design was adopted from the Gnutella specification and the application was successfully tested on a platform involving multiple Linux machines.

**Roles played:** Worked on the object-oriented design of the protocol. Implemented the protocol over a network of Linux machines.

5. **Design and implementation of a distributed, lightweight sensor network directory: ECE, ISU. Fall 2006-Summer 2007 (with Srikanta Tirthapura and Bojian Xu):** The goal was to design and implement an efficient and scalable distributed sensor network directory to track mobile objects in a wireless sensor network. The distributed Arrow protocol was deployed on a fixed spanning-tree-based network of MICA2 motes.

**Roles played:** Implemented the protocol with nesC on TinyOS and CrossBow motes.

## Computer skills

**Middleware:** Java RMI, Socket API, IBM Infosphere Streams

**Embedded Systems:** MICA2 platform for sensor networks, nesC, TinyOS 1.1

**Intrusion Detection Systems:** Bro, Snort

**Programming Languages:** C, C++, Java, nesC, Unix Shell Script, Latex, PL/SQL

**Operating Systems:** Red Hat Enterprise Linux 4.0, Windows XP, TinyOS

**Database:** Oracle 9i

**J2EE frameworks:** BC4J, Struts, Maverick

**Statistical packages:** R 2.7.2

## References

Available upon request.