Karthik Kambatla

400 N River Rd Apt 221, West Lafayette, IN 47906 kkambatl@cs.purdue.edu +1 765 315 2742

Interests

Distributed Systems and Cloud Computing

Education

Purdue University, West Lafayette, IN
Ph.D. Computer Science (Advisor — Prof. Ananth Grama) In progress
M.S. Computer Science (**GPA 4.0/4.0**) **May 2011**

Birla Institute of Technology and Science, Pilani, India **June 2007**
M.Sc. (Hons) Mathematics (**GPA 9.14/10.0**)
B.E. (Hons) Computer Science (**GPA 9.14/10.0**)

PhD Research

Erasure-Coded Messaging Infrastructure (ECMI)

- Proposed erasure-coded message persistence for improved and efficient fault-tolerance and elasticity in message-passing based distributed execution environments.
- Implementing ECMI in Apache S4 project to demonstrate its effectiveness.

Transactional MapReduce (TransMR)

- Built *TransMR* to provide *transactional support* for *MapReduce* over key-value stores to exploit *speculative parallelism*, by enabling data-sharing across parallel computations (map/reduce tasks). [*HotCloud '11*]

Large-scale graph analysis in MapReduce

- Formulated and implemented *relaxed semantics* for MapReduce to exploit *algorithmic asynchrony* inherent in most graph algorithms. [*Cluster '10*]
- Proposed *efficient workload distribution* in large-scale graph algorithms, particularly to optimize the key underlying operation — matrix-vector product. Our mat-vec implementation demonstrates an order of magnitude performance improvement over the *Pegasus* implementation. [*Purdue Technical Report '11*]

Hadoop Provisioning

- Devised automatic provisioning of MapReduce jobs based on the application's resource consumption signature. [*HotCloud '09*]

Professional Experience

Research Associate, HP Labs, Palo Alto **Jun 2010 - Aug 2010**

- Built a high-throughput, large-scale notification system using MapReduce and Bigtable.
- Devised *data-local reduce phase* for subscription matching to avoid remote lookups and associated network congestion. [*Patent filed May 2011*]

Software Engineering Intern, Google Inc., Mountain View **May 2008 - Aug 2008**

- Parallelized internal tools that monitor video storage placement/tracking for improved performance.

Research Intern, Microsoft Research India, Bangalore **Jul 2006 – Jun 2007**

- *Model Driven Driver Development* — Built the runtime scheduler for the *Clarity* programming model.
- *Secure Protocol for Automatic Connection Establishment* — Built secure, automatic mobile pairing using 2-phase Elliptic-Curve Diffie Hellman based key agreement over trusted channels. (Demonstrated at European Tech Fest 2006 at Brussels)

Research Intern, Helsinki University of Technology, Helsinki **May 2006 - Jul 2006**

- Implemented the *Blind* protocol in HIPL (Host Identity Protocol for Linux) to eliminate potential man-in-the-middle attacks.

Patents	Processing Notifications (Pending) *Karthik Kambatla, Jun Li*	**Filed in May 2011**
Publications	TransMR: Data-Centric Programming Beyond Data Parallelism *Naresh Rapolu, Karthik Kambatla, Suresh Jagannathan, Ananth Grama*	**HotCloud '11**
	Asynchronous Algorithms in MapReduce *Karthik Kambatla, Naresh Rapolu, Suresh Jagannathan, Ananth Grama*	**Cluster '10**
	Towards Optimizing Hadoop Provisioning in the Cloud *Karthik Kambatla, Abhinav Pathak, Himabindu Pucha*	**HotCloud '09**
	Efficient Large-Scale Graph Analysis in MapReduce *Karthik Kambatla, Giorgos Kollias, Ananth Grama*	**Purdue Technical Report '11**
	Clusterken: A Reliable Object-Based Messaging Framework to Support Data Center Processing *Marc Stiegler, Jun Li, Karthik Kambatla, Alan Karp*	**HP Labs Technical Report '11**
	Revisiting I/O middleware for the Cloud [*Work-in-Progress Report]* *Karthik Kambatla, Naresh Rapolu, Jalaja Padma, Patrick Eugster, Ananth Grama*	**FAST '10**
Service	External reviewer for *(1) HotCloud 2010* and *(2) Software: Practice and Experience (SPE) 2010*	
Programming Skills	C, Java, C++, C#, Python, Perl, PL/SQL Apache Contributor for the S4 project	
Awards	Top Graduate Teaching Assistant Award, Dept. of Computer Science, Purdue University 2010	
Other Activities	Faculty Search representative, Dept. of Computer Science, Purdue University 2009 Vice-Captain, Purdue University Cricket Club 2011 Webmaster, Indian Classical Music Association at Purdue (2007-08) Chief Editor for Computer Science Association at BITS Pilani (2004-05)	
References	Available upon request	