Prateek Sharma

prateek3.14@gmail.com
www.cse.iitb.ac.in/~prateeks

Research Interests

- Operating Systems, Virtualization, Storage, Language-runtimes.

Education

- **M.Tech Computer Science** — July 2009 - July 2012
  Indian Institute of Technology - Bombay — *GPA : 8.99/10*
- **M.Sc.(Tech.) Information Systems** — July 2005 - June 2009
  Birla Institute of Technology and Science - Pilani — *GPA : 7.14/10*

Experience

- **EMC Data Storage Systems** — Aug 2012 - Present
  Software Engineer, R&D Team — *Bangalore*
  - Work on Cloud, Big-Data analytics solutions involving text mining and social-networks.
- **Dept. of Computer Science Engineering - IIT Bombay** — July 2009- July 2012
  Research Assistant — *Mumbai*
  - **System administrator** incharge of maintaining critical computing facilities(30 servers, 200 desktops) and services(storage, authentication, email,backups).
- **IBM India Software Lab** — July 2008 - Dec. 2008
  Intern. Implemented Eclipse plugins for business report generation. — *Pune*
- **The Tata Power Company** — Mumbai
  Summer Intern — *May 2007 - July 2007*

Publications

- Singleton : Systemwide Page-Deduplication in Virtual Environments. Prateek Sharma, Purushottam Kulkarni. *HPDC 2012*.
- Per-file page cache for Linux. Prateek Sharma, Purushottam Kulkarni. *Technical Report*
- Software Testing using meta-heuristic techniques.Praveen Srivastava, Vinod Ramachandran, Manish Kumar, Prateek Sharma. *IEEE TENCON 2008*

Theses

- **Page-cache Management in Virtual Environments** — May 2011 - July 2012
  M.Tech. Thesis — *Advisor: Prof. Purushottam Kulkarni*
  Investigated memory-management and I/O performance problems in hypervisors and developed solutions for KVM (Kernel Virtual Machine). Performed extensive profiling to detect fundamental performance and design issues and developed these solutions:
  - **Inter-VM page deduplication:** Reduced the overhead of page deduplication by a factor of 10 by exploiting spatial locality, better data-structures and CPU features.
  - **Exclusive-caching:** Developed the first black-box exclusive page cache mechanism to eliminate duplicate content across guest and hypervisor.
  - **Page cache redesign:** Implemented a per-file page cache in Linux which reduces page cache size yet provides higher hit-ratio and improved scalability. Cache size and eviction algorithm can be adjusted on a per-file basis. Motivating example: hypervisors can run specialized second-level caching algorithms for virtual machines.

Projects

- **Memory Fingerprinting in QEMU**
  Course: Virtualization & Cloud Computing *Prof. Purushottam Kulkarni*
  - Developed a memory-fingerprinting tool for KVM/QEMU, allowing page-deduplication effectiveness to be accurately measured.
  - Empirically analyzed effectiveness of inter-VM page sharing for various workloads.
- **rCloud Deployment** July 2011 - June 2012
  - rCloud is IIT Bombay's private research cloud currently deployed on 8 physical servers.
  - Setup and maintained the hardware, hypervisors(KVM), networking, and cloud-management software (CloudStack)
- **DrawCAD** Jul 2010 - Dec 2010
  R & D Project *Prof. Abhiram Ranade*
  - DrawCAD is a sketch based CAD tool developed in JAVA.[50K lines]
  - Implemented several new features and improved the stroke/pattern recognition accuracy.
- **Binary translator for GCC**
  Course:Design and Implementation of GCC *Prof. Uday Khedker*
  - Proposed and prototyped a novel way of optimizing assembly code for a processor family by using GCC infrastructure(Machine Descriptions and RTL).
  - Use case: Most binaries are compiled for lowest common denominator (i386). Using the binary translator, it can be optmized for Core2,Opteron etc. at link-time.
- **Map-Reduce-Merge for Clusters**
  Course: Distributed Systems *Prof. Umesh Bellur*
  - Implemented an MPI(Message Passing Interface) based Map-Reduce framework from scratch.
- **Reed-Muller Error Correcting Codes**
  Seminar *Prof. Ajit Diwan*
  - Studied combinatorial properties and various decoding algorithms.
- **Intelligent Linux Installer** Jan 2007 - Mar 2007
  2nd prize in Software design contest in BITS-Pilani's annual technical festival
  - Developed a completely hands-off disk partitioner for Linux - which ensured safe partitioning and no data-loss. Incorporated various heuristics for optimizing free disk-space.

Skills & Expertise

- **Languages:** C, Java, Lisp, Python, Haskell, JavaScript, Go, R, Bash.
- **Linux:** Kernel development, Performance-optimization, System Administration.
- **Tools:** MPI, DTrace, Puppet, ZFS, Lucene, LLVM, GCC(internals), V8 JS-engine.

Contact Information

Email: prateek3.14@gmail.com
Webpage:`http://www.cse.iitb.ac.in/~prateeks`
Phone: (+91)9920092815
Address: 2012, Maruthi Apartments, 2nd Cross - Wind Tunnel Road, Bangalore, India