IBM Research - India

2015 Internship Program

Introduction

Curiosity has a way of finding answers.

Nowhere does this ring truer than at IBM Research-India, the premier software and service research facility of its kind in the Eastern hemisphere.

Since our inception in 1998 in Delhi, our single-minded mission has been to advance the state-of-the-art in information technology, through research in systems, software and services, in order to deliver innovations that not only bring tangible value to IBM's clients, but also positively impact individuals and communities, businesses and industries and society at large.

Co-located in New Delhi and in Bangalore (since August 2005), the IBM Research-India Lab are involved in a wide array of research areas, and boasts of a long list of achievements, in areas like high-performance computing, mobile-enabled emerging technologies, building enterprise resiliency, modelling natural disasters, bridging the digital divide, analytics and human language technologies…to mention just a few.

Key service innovations include Voice of Customer Analytics for contact centers, Recruiting, Staffing under Uncertainty and Business Contingency Planning technologies for workforce management, Application Assembly Optimization for streamlining global delivery, and Defect Prevention.

We owe much of our success to our culture of innovation, that thrives on global collaboration and our unmatched talent pool of world class researchers most of who are PhDs from premier global institutes like Cambridge, Carnegie Mellon, Cornell, Georgia Tech., IIT, IISc, INRIA, Stanford, University of California Berkeley, University of Maryland and Yale.

It is the curiosity displayed by these minds – the curiosity to explore uncharted frontiers and to push the innovation envelope-that is the driving force behind our achievements.


CURIOSITY
WANTED

Internship Opportunities

IBM Research – India (Bangalore & New Delhi)

We invite applications for our 2015 Internship Program in Computer Science, Mathematical Sciences and Services Science, Management & Engineering at its locations in New Delhi and Bangalore. We are seeking highly motivated graduate students, who are interested in experiencing an exciting summer of research. The selected students will have the opportunity to work closely with an outstanding research team on challenging problems that range from leading-edge exploratory work to prototyping real-world systems and applications. During the internship, the students will also have the opportunity to participate in the dynamic technical environment of the largest industrial research organization in the world and network with other students in different fields from other universities.

At a high level, our internships serve a dual purpose:

- Interns bring in fresh ideas and perspectives to the lab and help us conduct world-class research, thus creating impact.
- IBM Research provides interns an environment where they experience a world-class industrial research setting.

We measure the success of our internships in one of the following ways:

- Projects done at internships are targeted towards top conferences and journals. Most of our interns have been successful in the past in publishing their papers in top conferences.
- Projects completed during the internships become part of larger research projects at IBM.
- The projects that the students participate in are also of interest to the universities that they come from and help increase collaborations with the universities.


CURIOSITY
WANTED

Areas of Research

We offer internship positions in various research areas such as:

- **Business Analytics and Mathematical Sciences**
- Mathematical Programming and Combinatorial Optimization
- Stochastic Modeling and Optimization
- Scalable Machine Learning and Online Computation
- High Performance Computing & Analytics
- Scientific Computing & Modeling Using Coupled PDEs
- Game Theory and Mechanism Design
- **Information Management and Analytics**
- Big Data Processing
- Social Media Analytics
- Cleansing Noisy Data
- Spatio-Temporal Data
- Information Fusion @ Scale
- Information Integration and Entity Resolution
- IT Operational Analytics
- Information Trust and Security
- Network Analytics
- **Cognitive Solutions and Services**
- Natural Language Processing
- Machine Learning
- Information Retrieval and Extraction
- Text Analytics
- Natural Language Generation
- Question Answering Systems
- Dialog Systems
- Knowledge Representation and Reasoning
- **Productivity Tools and Software Engineering**
- Software Testing
- Software Development and Maintenance
- Mining Software Repositories
- Service Delivery Optimization
- Developing Cognitive Software
- **Smarter Planet Solutions**
- High-performance computing
- Weather & climate modeling
- Renewable energy
- Power systems
- Machine learning
- Optimization
- User Interface/Experience Design
- Embedded systems
- **Systems Research**
- Cloud Computing
- Distributed Systems
- Systems (Network, Storage, Servers)
- Software Defined Networking
- Data Center Networking
- Wireless Architecture

- **Mobile & Telecom Research**
  - Telecom and Mobile Analytics
  - Telecom Infrastructure and Middleware
  - Mobile Application Development Environments and Run-time
  - Mobile Device and Application Management
  - Context-aware Services
  - Machine to Machine
  - Mobile Enabled Solutions for industries

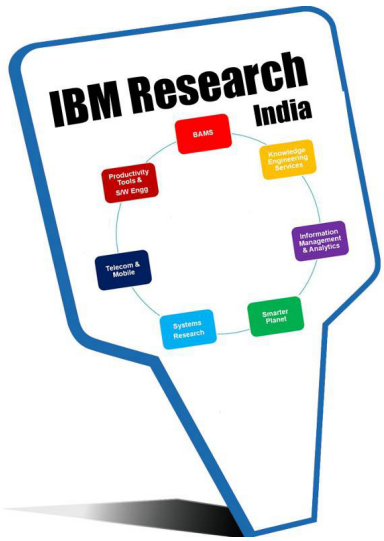

Business Analytics & Mathematical Sciences

In the Business Analytics and Mathematical Solutions (BAMS) department we conduct research in Optimization, Machine Learning and Analytics, Game Theory, and High Performance Computing. Our research work is published in highly rated conferences and journals (please see the department web-page listed below). Members of the department collaborate actively with research colleagues from academia that result in joint publications, tutorials, open collaborative projects, etc. We also lay high emphasis on applying our research and the state of the art techniques from the above fields to solve challenging real world problems: business problems faced in business settings within IBM, industrial problems faced by IBM's clients, and problems that cut across industries.

Workforce and talent analytics is one of the areas that we have been working on that is truly cross industry. This is an aspect that is important for any organization, but more so for business enterprises. In workforce and talent analytics, we work on strategic talent planning (e.g., hiring, engagement/deployment, retention/compensation, mobility), career growth and development, skills demand forecasting, capacity planning and utilization, etc. with a focus on the next-gen "smarter workforce" comprising of autonomous workers collaborating on social technology platforms. We leverage techniques from mathematical modelling, mathematical programming, machine learning, mechanism design, and discrete event simulation in our work.

We have deep domain expertise in certain industries and help improve state of the art industry practices. An example of this is the oil and gas industry. In this context of oil and gas industry, we work on (i) Challenging operational problems: shutdown management, asset operations and maintenance, and improvement of health, safety, environmental issues, and (ii) Domain specific problems: simulation and large-scale inverse problems that arise in the context of understanding the structure and production potential of oil fields such as Seismic Analysis, Full Wave Inversion, Fracture Modelling and History Matching. Technically, the work on operational domain involves large-scale mathematical programming, building predictive models from structured data coming from sensors/SCADA systems and unstructured information documented by human experts, time-series analysis, anomaly detection, and spatio-temporal analysis. The work on understanding oil fields involves solving large-scaled coupled PDEs. This requires scalable implementations of compute-intensive and/or data-intensive scientific and machine learning kernels on multi- core and many-core distributed architectures such as GPU / Linux clusters and supercomputers such as Blue Gene/P and Blue Gene/Q.

Business Analytics & Mathematical Sciences

Contd...

Techniques to infer preferences of users and relationship between users based on partial observations of their activities is an area of broad interest to the department and is applied in various domains. Some of the specific problems we have worked on are social influence analysis, active information acquisition, channel attribution, and estimating life-time value of customers. These are applied in the context of retail companies, marketing scenarios, up-sell/cross-sell in contact centers, etc. Here, we employ techniques from graphical models, online inference, non-parametric techniques, convex optimization, and mechanism design.

For further details, visit: http://researcher.ibm.com/researcher/view_project.php?id=4959

Skill Set / Level: Ph.D. or Master Students, familiar with at least a few of the following subjects either through course work or self-study: linear & nonlinear optimization, combinatorial optimization, machine learning, game theory, high performance computing, and agent-based modeling & simulation. The candidates should be able to work with one or more of following programming tools/languages/platforms: Java/C++, ILOG, SPSS/R, MATLAB, AnyLogic, MPI/OpenMP, OpenCL/CUDA, Pthreads.

Location: Internship positions are available at New Delhi and Bangalore.

Information Management & Analytics

The Information and Analytics group at IBM Research – India is focused on developing next-generation technologies in various areas such as database systems, cloud computing, information retrieval, distributed computing, information integration, business intelligence, data/text mining, and big data platforms. These technologies are driven by IBM Research's goal of building intelligent solutions and services to address business problems in various industrial sectors, including financial, telecommunication, retail, and healthcare, among others. Our current projects and areas of interest include:

- Spatio-Temporal Data Anaytics
- Indexing, Querying and Efficient Management of Graph Databases
- Platforms for large-scale data/text mining leveraging scale-out platforms such as Hadoop and NoSQL databases
- Information Retrieval algorithms that do joint reasoning over text sources and entity graphs
- Scalable Graph processing algorithms over platforms such as Titan and Neo4j
- Social Media Mining for applications such as Event Extraction and Visualization
- Various aspects of Data curation including discovery, integration, entity matching and fusion over structured and unstructured data sources
- Indexing and Analysis of IT operational data such as Log, config, process and devops information for Fault determination
- Theoretical aspects of Big data including, relevance, resolution, sampling and other information theoretic concepts
- Data privacy and security for Cloud DBaaS platforms
- Core database technologies

For further details, visit: http://researcher.ibm.com/researcher/view_project.php?id=4966

Skill Set/Level: Highly motivated PhD or Master Students with relevant background and strong research interest in related areas are encouraged to apply.

Location: Internship positions are available at New Delhi and Bangalore.

Cognitive Solutions and Services

The Cognitive Solutions and Services department at IBM Research - India is focused on developing the next generation of cognitive technology solutions and services to fundamentally change the way we interact with computers, people, and enterprise scale systems.
We are creating new products and service offerings based on advanced text and content analytics, machine learning, natural language processing, and the Watson technology. We are building platforms for supporting personalized education, new methods of naturally interacting with enterprise infrastructure and systems, and new methods of delivering domain specific insights to customers. We are also conducting research on the next generation of cognitive interaction technologies.
We leverage our expertise in advanced text analytics, data mining, and machine learning to maintain strong academic credentials through top-quality publications, external collaborations, and service to the research community, as well as contribute to IBM business units via well-motivated applied research.
Some of the exciting projects that we are working on include: Smarter Education, Watson for Infrastructure Management and Troubleshooting, Watson Wealth Advisor, and IBM Debating Technologies.

For further details, visit:
http://researcher.watson.ibm.com/researcher/view_group.php?id=5431

Skill Set / Level: Highly motivated Ph.D. or Master Students with relevant background and strong research interest in related areas are encouraged to apply.

Location: Internship positions are available at New Delhi and Bangalore.

Mobile and Telecom Research

There has been a tremendous growth in the penetration of the Mobile phones in the recent years. Moreover, there has been significant progress in the computing power, memory, display and other features of mobile phones. The Telecom & Mobile Department at IBM India Research Lab focuses on this exciting area of mobile computing and challenges of Telecommunication industry with the goal of creating innovative solutions and platforms. Researchers in the department collaborate extensively with other IBM business units, Telecom Service Providers, various customers looking for Mobile solutions as well as Academia.

Focused on promoting advanced telecommunications and mobile solutions and infrastructure development, the department currently conducts research in following key areas:
1. Telecom and Mobile Analytics
2. Telecom Infrastructure and Middleware
3 Mobile Application Development Environments and Run-time
4. Mobile Device and Application Management
5. Context-aware Services
6. Machine to Machine
7. Mobile Enabled Solutions for industries (such as Retail, Finance, etc.)

For further details, visit: http://researcher.ibm.com/researcher/view_project.php?id=4979

Skill Set / Level: We are seeking applications from Ph.D. or Master Students in Computer Science (or related field). Knowledge of one or more of Mobile Device Platforms, Distributed Systems, Networking, Data Mining, Programming knowledge in Java/C++ and/or platforms like iPhone, Android, etc.

Location: Internship positions are available at New Delhi and Bangalore.

Productivity Tools & Software Engineering

Productivity Tools and Software Engineering

The research agenda of the Productivity Tools and Software Engineering group at IBM Research- India focuses on improving software product and service quality throughout the life-cycle and developing new techniques for mobile and cloud enablement of software and services. Our agenda is driven specifically by the needs of the software services industry. Our research uses many core techniques, such as program analysis, text analysis, and data mining. In addition to developing new technologies that have a significant business impact, we strive for creating intellectual property and make broader scientific impact by publishing our research in top conferences. Below we describe our current research themes.

Software Testing

Our research agenda in software testing focuses on developing techniques and tools that bring automation and rigor to the tasks that are performed manually in testing services, often in an ad-hoc manner, and are prone to human lapses. Our research covers a broad spectrum of topics, including test automation, test data generation, test suite reduction, test repair, and regression testing.

Software Development and Maintenance

Our group is developing scalable automated techniques for API extraction from legacy code, fault detection, localization, explanation and repair. We develop new techniques for migration to mobile and cloud platforms and create productivity tools to improve design of cognitive, database-driven, cloud-based applications. Our methods employ sophisticated constraint-solving and theorem-proving techniques coupled with efficient program-analysis techniques, such slicing and differencing, along with effective information visualization and summarization for reducing cognitive overload.

Mining Software Repositories

Repositories, such as version management systems and bug management systems, contain a wealth of information on how an application evolves over time. This information, if mined appropriately, can help project teams gain insights for supporting maintenance of the software, improving design/ reuse, and enabling effective transitioning of new people into a project. Our research agenda in this area focuses on using analytics on code and other artifacts generated during SDLC, and visualization to enable decision making, based on knowledge derived from software repositories.

Service Delivery Optimization

The goal is to improve the efficiency of service delivery processes through intelligent assignment, scheduling and routing of work for management of practitioner utilization, and application of simulation and optimization techniques to determine optimal staffing in the delivery organization. A specific context that motivates this research is the factory model of shared delivery, where instead of having customer-specific teams, a common pool of practitioners in a given specialization area is responsible for meeting the service requests of multiple customers in that area. Building automation tooling capabilities that assist practitioners in their day-to-day work is also an area of interest. Ticket lifecycle management, application configuration discovery,patch management and server health-check monitoring

Productivity Tools & Software Engineering

Contd...

are all human-intensive processes which are made easier and more efficient by tools and technologies that allow auto-discovery, auto-configuration and auto-remediation. We are also exploring the social network of practitioners and evolution of skills in large-scale global delivery centers following such a model. Information visualization is also another area of work for aiding decision making by service delivery managers.

Developing Cognitive Software

We are gradually moving towards an era of cognitive software development, where humans can interact with apps in more 'natural' ways and apps are able to exploit their human and physical contexts more skillfully. Creating such cognitive apps is a complex endeavor involving natural interaction, reactive mobile interfaces, carefully-orchestrated distributed computation pipelines and seamless cloud-based deployment. Our goal is to investigate new abstractions, domain-specific languages and distributed computation paradigms to enable effortless development of robust next-generation cognitive apps.

For further details, visit: http://researcher.ibm.com/researcher/view_project.php?id=4955

Skill Set / Level: We are seeking applications from Ph.D. or Master Students in Computer Science (or related field). Candidates with knowledge of one or more in empirical software engineering, compilers, program analysis, web technologies, machine learning, simulation & optimization, mining and analysis of software engineering artifacts are welcome to apply.

Location: Internship positions are available at New Delhi and Bangalore.

Smarter Planet Solutions

A significant initiative has been taken at IBM to make the world better instrumented, highly interconnected and more intelligent. We believe the recently announced, new era of "Cognitive Computing" will accelerate realization of this goal. Several key research directions aligned with this initiative are being pursued at the Smarter Planet Solutions department of IBM Research in India. For example, leveraging the power of High Performance Computing, we are creating ultra high-resolution weather models to provide highly customized weather prediction and integrating them with hydrology models to provide early warning on potential floods. High-resolution weather models combined with long-term climate models are being utilized to help farmers better plan their farm activities (sowing, irrigation, pesticide application, etc.) and also identify the right crop for their geographical region. We are working on combining weather modeling with renewable energy generation to come-up with optimal wind-farm and solar-farm design and scheduling algorithms at power generation plants. We are working on various machine learning and optimization algorithms using the data collected from power generation plants, transmission and distribution networks to help companies predict and manage demand peaks, identify leaks, predict black-outs and estimate connectivity models to utilize their resources most efficiently. We are also innovating on "smarter plugs" which can sense grid status, appliance characteristics and user preferences and analyze to flatten the demand curve and help stabilize both grid and appliance operations. See the website given below for more details.

http://researcher.ibm.com/researcher/view_project.php?id=4990

Skill Set/Level: We are inviting applications for summer internship from highly motivated MS and Ph. D students with a strong academic record, working in the following areas.

- High-performance computing
- Weather & climate modeling
- Renewable energy
- Power systems
- Machine learning
- Optimization
- User Interface/Experience Design
- Embedded systems

Location: Internship positions are available at New Delhi and Bangalore.

Systems Research

The goal of system research department is to design and build next generation systems with higher levels of elasticity, flexibility, performance, and cost efficiency based on a holistic view of compute, network, storage, & software. Our current focus areas include enterprise cloud and data center networking. High-level pre-requisites and current projects are listed below.

Distributed Systems

The Distributed Systems group works in the area of cloud computing to make it more consumable in an enterprise setting. To that end we are developing new techniques for efficiently migrating traditional enterprise client systems to managed cloud environment, and for automating system management in enterprise clouds.

Pre-requisites:

- Server Virtualization Basics, Strong knowledge of operating systems and computer architectures
- Algorithms and machine learning knowledge (preferred)
- Strong system building skills in C/C++/Python/Java, shell scripting

Data Center Networking

The Data Center Networking group is exploring ways to prevent network from becoming a performance and manageability bottleneck in large-scale cloud data centers. Some of the topics that we are currently investigating include: software defined networking (SDN) - unique value it can bring to an enterprise network and challenges to its adoption; and opportunities for co-optimization of computation, storage & networking using network-aware virtual machine management.

Pre-requisites:

- IP communication network fundamentals
- Solid programming skills in any one high level language (Java preferred)
- Familiarity with network simulation and modeling tools (desired)
- Familiarity/expertise in python (a plus)
- Familiarity/expertise in Openstack Networking – Neutron (a plus)

Systems Research

Contd..

Smarter Wireless, Network Analytics

With the increasing compute and networking capabilities of general purpose processors in the cloud, the wireless/IT research group explores applying traditional IT optimizations and analytics to wireless infrastructures. Our research mainly focuses on (i) developing next generation wireless architectures on cloud platforms and (ii) building smarter wireless networks, specifically cellular networks by applying spatio-temporal analytics and optimizations. Interest in systems research towards building solutions in cellular networks in a large team with emphasis on applied research in real network deployments.

Pre-requisites:

- Wireless and IP communication network fundamentals
- Basic machine learning, data mining and their application to wireless and networking.
- Familiarity with network simulation and modeling tools
- Solid programming skills in Java or C++

For further details, visit: http://researcher.ibm.com/researcher/view_project.php?id=4948

Skill Set/Level: We are inviting applications for summer internship from highly motivated MS and Ph. D students with strong academic record and pre-requisites given above.

Location: Internship positions are available at New Delhi and Bangalore.

How to apply

Eligibility Criteria: Graduate students (preferably PhD) in Computer Science, Mathematical Science and Services Science, Management & Engineering and related.

Stipend and Relocation: The interns will be paid a stipend that is competitive with what other leading companies pay in India. IBM Research - India will also cover expenses related to travel to the work location and provide logistical support for accommodation. In the case of PhD interns, additional monetary support will also be provided for accommodation.

Location: New Delhi and Bangalore.

Application portal: Application along with your latest CV can be submitted at https://university-relations.in/wps/portal/internship

Contact: If you have any queries, reach us at urirl@in.ibm.com

Application Deadline: You must submit your internship application on or before **January 31, 2015** for consideration in the 2015 program. Short-listed candidates will be contacted for interviews in February 2015.