

Project Ideas Brainstorming

Ioan Raicu

Computer Science Department
Illinois Institute of Technology

CS554: Data-Intensive Computing
August 28th, 2013

Developing a research proposal

- Identify a problem
- Review approaches to the problem
- Propose a novel approach to the problem
- Define, design, prototype an implementation to evaluate your approach
 - Could be a real system, simulation and/or theoretical
- Write a technical report
 - Intro, motivation, related work, proposed work, evaluation, future work, conclusion, and references
- Present your results
- Write a workshop/conference paper (optional)

Survey Results

Project	#	%
ZHT:OHT	11	8%
AWS:DQS	10	7%
AWS:DTS	10	7%
OS:DistOS	7	5%
ZHT:DMHDFS	7	5%
FusionFS:Cache	6	4%
FusionFS:DAS	6	4%
FusionFS:FFSZ	6	4%
FusionFS:FHFS	6	4%
GeMTC:MTCSim	6	4%
ZHT:Bench	6	4%
ZHT:Graph	6	4%
NET:MPNet	5	4%

Project	#	%
FusionFS:IDA	4	3%
GeMTC:GMon	4	3%
GeMTC:MIC	4	3%
MATRIX:DJLSys	4	3%
FusionFS:Lib	3	2%
FusionFS:Prov	3	2%
GeMTC:GApp	3	2%
GeMTC:Test	3	2%
MATRIX:BenchTEF	3	2%
ZHT:ZST	3	2%
FusionFS:CKPT	2	1%
MATRIX:BenchJMS	2	1%
MATRIX:DJLSim	2	1%
MATRIX:MonSim	2	1%
MATRIX:Swift/M	2	1%
ZHT:Const	2	1%
ZHT:MonSys	2	1%
GeMTC:OpenMTC	0	0%
MATRIX:HJLSim	0	0%
Other	0	0%

Project Naming Convention

- Example: FusionFS:Cache
 - FusionFS
 - Overarching project
 - Context
 - Project Mentor
 - Cache
 - Project topic

Active Projects

- **Storage**

- FusionFS: Fusion distributed File System
 - HyCache, FusionProv, IStore, RXSim
- ZHT: Zero-Hop Distributed Hash Table
 - NoVoHT

- **Computing**

- **Many-Task Computing**

- MATRIX: MAny-Task computing execution fabRIc at eXascales
 - SimMatrix
 - Falkon: Fast and Light-weight task executiON framework
 - Swift: Fast, Reliable, Loosely Coupled Parallel Computation

- **Many-Core Computing**

- GeMTC: Virtualizing GPUs to Support MTC Applications


- **Cloud Computing**

- AWS: Leveraging cloud infrastructure towards scalable distributed system services


Major Overarching Projects

- **FusionFS**: Fusion distributed File System
- **ZHT**: Zero-Hop Distributed Hash Table
- **MATRIX**: MAny-Task computing execution fabRlc at eXascales
- **GeMTC**: Virtualizing GPUs to Support MTC Applications
- **AWS**: Leveraging cloud infrastructure towards scalable distributed system services

PhD Students



- **Tonglin Li** PhD Student, 5th year (**Official TA**)
- Research Interests: *Distributed file systems, Cloud computing, Data-intensive computing, Supercomputing, and Data management*
- Project Lead: [ZHT](#)




- **Ke Wang** PhD Student, 4th year (**Official TA**)
- Research Interests: *Distributed Systems, Many-Task Computing, Data-Intensive Computing, Cloud Computing, Supercomputing, High-Performance Computing*
- Project Lead: [MATRIX](#)



- **Iman Sadooghi** PhD Student, 3rd year
- Research Interests: *Distributed Systems, Cloud Computing, Data Mining*
- Project Lead: AWS



- **Scott Krieder** PhD Student, 3rd year
- Research Interests: *Distributed Systems, Cloud Computing, Supercomputing, High Performance Computing Security, Network Security*
- Project Lead: [GeMTC](#)



- **Dongfang Zhao** PhD Student, 2nd year
- Research Interests: *Cloud Computing, Distributed File Systems, Intelligent Database Systems and Data Mining*
- Project Lead: [FusionFS](#)

Timeline

- Major topics covered: FusionFS, ZHT, MATRIX, GeMTC, and CloudKon
- 08-30-2013 (Fri 12:45PM): Project Overviews (ZHT – Tony)
- 09-03-2013 (Tue 12:45PM): Tutorial (ZHT – Tony)
- 09-03-2013 (Tue 1:50PM-3:05PM): Linux and Jarvis cluster -- Life Sciences 111, part of CS550 (optional)
- 09-04-2013 (Wed 11:25AM): Project Overviews (FusionFS – Dongfang FusionFS & Ioan Projects)
- 09-05-2013 (Thur 11:25AM): Tutorial (FusionFS – Dongfang)
- 09-06-2013 (Fri 12:45PM): Project Overviews (MATRIX – Ke)
- 09-06-2013 (Fri 2:00PM): Tutorial (MATRIX – Ke)
- 09-09-2013 (Mon 11:25AM): Project Overview (GeMTC – Scott)
- 09-09-2013 (Mon 12:45PM): Tutorials (GeMTC – Scott)
- 09-11-2013 (Wed 11:25AM): Project Overview (CloudKon – Iman)
- 09-12-2013 (Thur 12:45PM): Tutorials (AWS – Iman)
- 09-16-2013 (Mon 12:45PM): Tutorial Linux Tony

Questions

