TSINGHUA SCIENCE AND TECHNOLOGY
ISSN 1007-0214 10/10 pp95-101
Volume 19, Number 1, February 2014

Mobile Internet Big Data Platform in China Unicom

Wenliang Huang, Zhen Chen*, Wenyu Dong, Hang Li, Bin Cao, and Junwei Cao

Abstract: China Unicom, the largest WCDMA 3G operator in China, meets the requirements of the historical Mobile Internet Explosion, or the surging of Mobile Internet Traffic from mobile terminals. According to the internal statistics of China Unicom, mobile user traffic has increased rapidly with a Compound Annual Growth Rate (CAGR) of 135%. Currently China Unicom monthly stores more than 2 trillion records, data volume is over 525 TB, and the highest data volume has reached a peak of 5 PB. Since October 2009, China Unicom has been developing a home-brewed big data storage and analysis platform based on the open source Hadoop Distributed File System (HDFS) as it has a long-term strategy to make full use of this Big Data. All Mobile Internet Traffic is well served using this big data platform. Currently, the writing speed has reached 1 390 000 records per second, and the record retrieval time in the table that contains trillions of records is less than 100 ms. To take advantage of this opportunity to be a Big Data Operator, China Unicom has developed new functions and has multiple innovations to solve space and time constraint challenges presented in data processing. In this paper, we will introduce our big data platform in detail. Based on this big data platform, China Unicom is building an industry ecosystem based on Mobile Internet Big Data, and considers that a telecom operator centric ecosystem can be formed that is critical to reach prosperity in the modern communications business.

Key words: big data platform; China Unicom; 3G wireless network; Hadoop Distributed File System (HDFS); mobile Internet; network forensic; data warehouse; HBase

1 Introduction

Users of the Mobile Internet[1] can access any content, anytime, and anywhere. This convenience produces a large volume of individual user network traffic on the telecom operator side, so is referred to as Mobile Traffic Deluge. According to Mary Meeker's report[2] on Mobile Internet Trends, more and more PC software is migrating to Mobile Internet devices. It is also predicted that mobile traffic will double each 14 months and that the volume of Internet traffic will quadruple between 2011 and 2016, reaching 1.3 ZB per year in 2016, as indicated by Cisco VNI[3]. China Unicom, the largest 3G operator in China, is prepared to meet this "Mobile Internet Explosion".

According to the statistics from China Unicom, who had approximately 250 million client users in 2012, mobile user traffic is increasing rapidly with a Compound Annual Growth Rate (CAGR) of 135%. Mobile Internet traffic characteristic has also

• Wenliang Huang is with China Unicom Groups, No. 21 Financial Street, Xicheng District, Beijing 100140, China.
• Zhen Chen and Junwei Cao are with Research Institute of Information Technology (RIIT) and Tsinghua National Lab for Information Science and Technology (TNList), Tsinghua University, Beijing 100084, China. E-mail: {zhenchen, jcao}@tsinghua.edu.cn.
• Wenyu Dong and Bin Cao are with Department of Computer Science and Technology, Tsinghua University, Beijing 100084, China.
• Hang Li is with Department of Computer Science and Technology, PLA Univ. of Info. & Eng., Zhengzhou 450001, China.
∗ To whom correspondence should be addressed.
Manuscript received: 2014-01-09; accepted: 2014-01-10

been investigated in Ref. [3] and a traffic prediction model based on ARMA and FARIMA has been proposed to capture the multi-fractal spectra in mobile traffic.

China Unicom's big data platform, starting from October 2009, has recorded monthly traffic of more than 2 trillion records, monthly data volume is over 525 TB, and the maximum data volume recorded has reached a peak of 5 PB. Overall writting speed has reached 1 390 000 records per second, and the recorded retrieval time in the table that contains trillions of records is less than 100 ms.

2 Related Work

Network traffic recording or archiving is always applied in network forensics, network troubleshooting, and user behavior analysis. All inbound and outbound traffic from a certain vantage point can be recorded to restore the original condition at a later time if necessary.

Regarding storage limits, only network flow data or statistics are recorded, that only contains source and destination IP addresses, ports, protocols, and timestamps.

The actual flow contents are usually neglected, which would otherwise require a huge volume of repository to accommodate. In addition, there are some legal debates in Deep Packet Inspection (DPI) of flow contents concerning user privacy issues. Sometimes this information is useful for quickly identifying phishing[4], spammers, and other types of cyber-attacks.

CNSMS[5] and TIFAflow[6] are used for traffic acquisition and aggregation for forensic analysis. CNSMS is an architecture for traffic acquisition with TIFAflow and its UTM appliance for traffic aggregation used in forensic analysis in a cloud computing based security center. TIFAflow is a software-based probe that combines TIFA[7-9] with Fastbit[10] indexing to provide granular data storage. It may be operated as an independent prober or integrated into CNSMS's UTM appliance.

Deri and Fusco[11,12] also proposed MicroCloud-based flow aggregation for fixed and mobile networks. This architecture is used to provide real-time traffic monitoring and correlation in large distributed environments. Their system is deployed in the VIVACOM (Bulgarian Telecom) mobile network and is used for monitoring the .it DNS ccTLD and a large 3G mobile network.

There are other works that use a similar platform for network data analysis, just like Lee et al.[13,14] and Qian et al.[15]

However, for any mobile network operator even only recording network flow data, the resulting data repository could easily reach the Terabyte level on a yearly basis. However, if all mobile traffic data is recorded for forensic analysis, the volume of the data could easily reach the Petabyte level. This remains a major challenge to a mobile network operator that must accommodate and index such big data for further analysis.

3 Mobile Traffic Acquisition at China Unicom

3.1 Traffic data acquisition

In China Unicom, traffic acquisition is operated at each Gn point of the GGSN in the 3G WCDMA mobile network, representing the vantage point of the mobile Internet in each province, and there are more than one hundred GGSNs used to cover all service areas. Traffic acquisition captures all the IP packets and aggregates the packets from each user properly.

The principle of the aggregation is that a user's valid behavior data should not be lost and that efficiency is required to reduce the invalid data. Then the file is produced in less than five minutes, and the volume of every file is less than 200 MB. Every file contains approximately 700 000 records. The detailed deployments of traffic probers are shown in Fig. 1.

All traffic types are resolved once the traffic is captured. The captured traffic is transmitted after being packaged using a private format that is designed

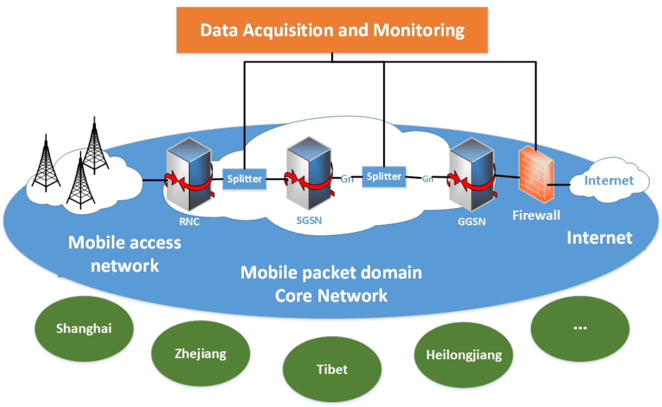

Fig. 1 Traffic acquisition in Gn point of mobile network in China Unicom.

according to China Unicom's uniform Internet records query and analysis system. The detailed format for a traffic record is shown in Table 1.

Some notes about important fields shown in Table 1 are also described as the following 6 rules:

(1) The bold field in the table needs to be captured, however, in the first stage, it is not stored, but other fields need to be captured and placed in storage.

(2) The value of traffic data packets without related information will be set to null.

(3) In the detailed record files, multiple CDRs are separated by a Carriage Return (CR) symbol and a newline symbol.

(4) To ensure that the information is immediately available for querying in 30 minutes, an intermediary log is generated every five minutes for all the protocols. The records of IM traffic (such as QQ, WeChat, Fetion, XMPP) are merged with the user login ID. The traffic records of RTSP, FTP, SIP, and other traffic types are merged with the control and data channel, and the merged record is identified with the control channel port. The traffic records of other multiple IPs and channel traffic are merged, and the merged record is identified with the first IP and port.

(5) Collect the WAP information and HTTP information that contains a complete URL field, including the "http://" and the host domain information, if there is no such information, the field must be filled with a null string.

(6) Traffic type coding is accomplished with 3 digits.

There is a vertical bar used as a separator between each field in a traffic record. The interval of the traffic file generation is 5 minutes by default and can be modified on demand. The size of one single file is limited to less than 200 MB. In each time interval, a traffic record file is generated and writing into the record will end when the time limit is reached or the file size limit is reached. If the size of one single file exceeds 200 MB, multiple files will be produced to guarantee that the size of the single file is below the threshold, and the additional related files are identified by appending a hexadecimal number such as [nnnnn]x.

Table 1 Detailed description of the traffic record format.

Number	Field	Remark
1	Cell phone number	Not contain a prefix such as +86, 0086, 86
2	Location area code	LAC
3	CI number	Select the first CI when a network switches
4	Terminal type	IMEI
5	Traffic type	
6	Start time	YYYY-MM-DD HH:MM:SS.1234567, Accurate to 0.1 Microsecond
7	End Time	YYYY-MM-DD HH:MM:SS.1234567, Accurate to 0.1 Microseconds
8	Duration (in seconds)	
9	Upstream traffic (in bytes)	
10	Downstream traffic (in bytes)	
11	Total traffic (in bytes)	
12	RAT type	1 represents 3G, 2 represents 2G
13	Terminal IP	
14	IP Visited	Not exist the IP information: null, Multiple IP information or the record of combined traffic: The first IP
15	Status code	
16	**User agent**	**Collect all information**
17	APN	3gwap,3gnet,uniwap,uninet,cmwap,cmnet
18	IMSI	
19	SGSN IP	The first access IP
20	GGSN IP	
21	**Content-type**	
22	Source port	
23	Destination port	
24	Record logo	0: The records unconsolidated and not split. 1: The records consolidated but not split. 2: The records unconsolidated but split. 3: The records consolidated and split.
25	Merge records	1,3: The number of combined records 0,2:null
26	URL/feature information	Business with URL/URI: the information of URL/URI Business without URL/URI: Specific information

3.2 Traffic data warehouse

The files are transmitted by FTP protocol to the twenty-four FTP servers located in Beijing. Two small provinces normally share an FTP server, while a large province normally requires two FTP servers. To reduce the bandwidth of transmission, all files are compressed by the bzip2 compression algorithm before the files are uploaded to Beijing from every province.

The warehousing program also runs on the FTP servers, and reads the files transmitted using FTP protocol. After being decompressed, the files are written into an HBase by a native Java API supported by HBase. In HBase, an online record table will be generated for each month. As the row key, the cell phone number is used to index and search the online records quickly.

4 Big Data Platform

4.1 Architecture

All functionalities of the Big Data Platform developed by China Unicom are shown in Fig. 2. The platform is based on a well-designed usage model and its capability for other purposes has been developed based on fundamental data storage, query, and analysis. Using distributed storage and a distributed HBase, Mobile Internet traffic data can be retrieved and analyzed for many purposes, such as commerce, accounting, and

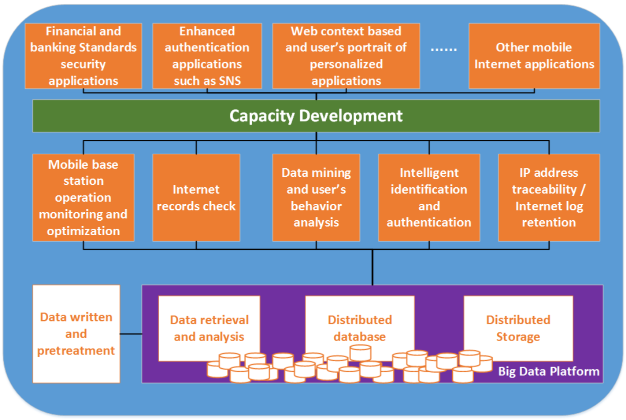

Fig. 2 Big data platform for data storage and mining in China Unicom.

national security issues.

4.2 Main features

4.2.1 Distributed storage—HDFS

The cluster contains 188 nodes or hosts. Ten hosts are used as management nodes and 159 as data nodes. The remaining 19 hosts are used for statistical data analysis and join the cluster when high loads are experienced.

There are three namenodes with high availability configuration. Each node has two Intel quad core CPUs, 96 GB memory, and 14 hard disks each 1 TB in size. The namenode and Job tracker are deployed in the same server. In addition, there is one secondary namenode and one standby namenode.

There are seven hosts used as Zookeeper and hmaster. Each host has two Intel quad core CPUs, 24 GB memory, and runs zookeeper and hmaster services for Hadoop Distributed File System (HDFS) and HBase.

There are 159 data nodes. Each data node has two Intel quad core CPUs, 48 GB memory, and fourteen 1 TB hard disks with a total of 12 TB disk space available for use. Each datanode runs regionserver and tasktracker services.

The total available disk space of this cluster is approximately 1.9 PB, and 1.43 PB is currently in use, or 74.1% of total capacity. By the end of September 2013, the storage used was expected to reach approximately 1.63 PB, or 85% of total capacity at the end of 2013.

4.2.2 Distributed database—HBase

The records written into Hbase exceed 7 billion each day. The peak writing rate is 145 000 records per second. HBase is configured to use LZO compression algorithm and the compression ratio reaches 28%. In particular, it has been empirically demonstrated that LZO is four to five times faster than zlib[16] in decompression, even when using zlib at the fastest compression level[17].

HBase is used to store structured data based on HDFS, which is shown in Fig. 3. Each region server in HBase contains 1500 regions.

5 Evaluation Experiments

5.1 Evaluation environments

The performance of the relational database depends on the configuration used, so we show the configuration details of both Oracle and HBase. Tables 2-5 show the test environment of Oracle, and Tables 6-8 show the test environment of HBase. We compare the performance of the Oracle and HBase by querying the record of a specified telephone number.

5.2 Evaluation results

We evaluate our optimized HBase database with Oracle database with cluster mode. Several insertion and querying experiments are conducted to compare these

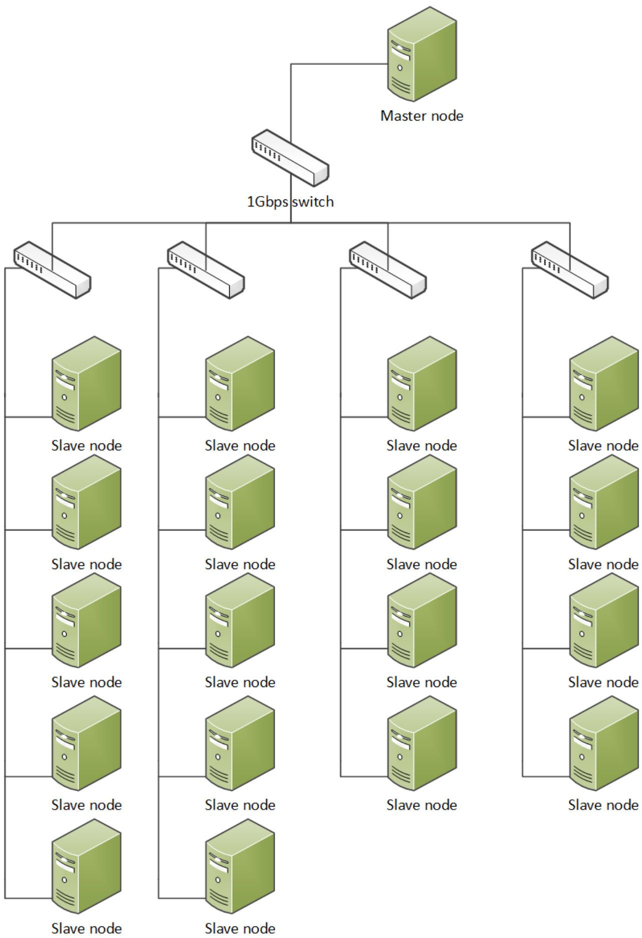

Fig. 3 Hadoop cluster for mobile Internet big data storage.

Table 2 System configuration of Oracle.

System	Windows server 2003 r2 enterprise
Processor	Intel Xeon(R) CPU X5680 3.33 GHz
Memory	Samsung DDR3 1333 2 GB×12
System type	64-bit operating system
Oracle type	Oracle 11g enterprise 64-bit r2
Hard disk 1	Seagate 1 TB 7200 r/m cache:32M SATA 2.0
Hard disk 2	Seagate 2 TB 7200 r/m cache:32M SATA 3.0
Hard disk 3	Seagate 500 GB 7200 r/m cache:32M SATA 2.0

Table 3 Configuration parameter of Oracle database.

Oracle type	Oracle 11g enterprise 64-bit r2
Management	Automatic memory management
Memory_target	18 048 GB
Memory_max_target	18 048 GB
Processes	500

Table 4 Description of the insert data in Oracle.

Attribute num.	26
Avg. size	370 byte
Phone number (Not repeated)	2 128 795

Table 5 Table information of Oracle.

	Insert line num.	Size of the table (GB)	Size of the index (GB)
Table 1	About 600 million	190.3	7.002
Table 2	About 1200 million	380.6	13.890
Table 3	About 2300 million	760.7	27.680
Table 4	About 2900 million	1031.3	37.430

Table 6 Cluster configuration of HBase.

OS	CentOS 6.2
Processor	Intel(R) Xeon(R) CPU E5-2620 0 @ 2.00 GHz
Memory	48 GB
System Type	64-bit opreating system
HBase Version	0.94.1-Intel
Hadoop Version	1.0.3-Intel
Nodes	5

Table 7 Table attribute of HBase.

Avg. size	370 byte
Rigion num.	200
Compressing solution	SNAPPY

Table 8 Table information of HBase.

	Data line num.	Size of the table (GB)
Table 1	About 600 million	24
Table 2	About 1200 million	47
Table 3	About 2500 million	99
Table 4	About 3100 million	123

two solutions. The throughput of insertion and querying with these two solutions are also shown in Figs. 4-7.

First, we compare the insertion rate of traffic records in our home-brewed HBase with Oracle database. Figure 4 shows the insertion rate of traffic records for Oracle database with the volume of traffic

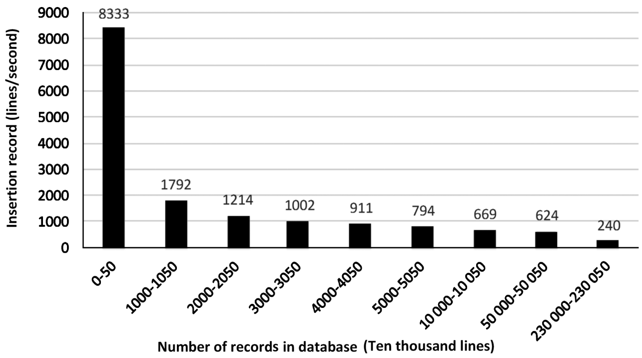

Fig. 4 The insertion rate of traffic records in Oracle database for comparison.

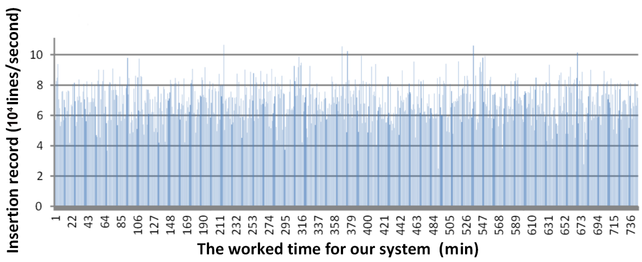

Fig. 5 The insertion speed of records in our system.

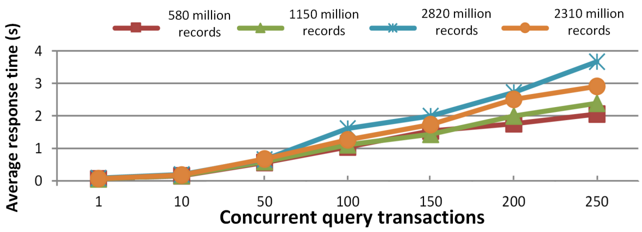

Fig. 6 Data query experiments in concurrent query transactions with Oracle database.

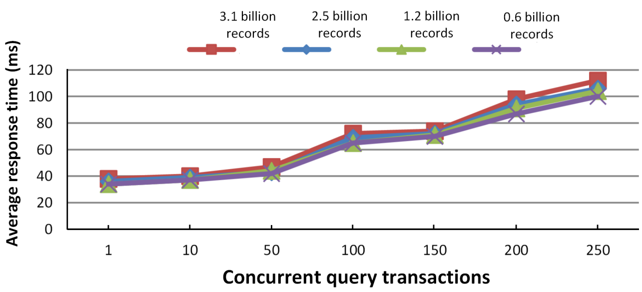

Fig. 7 Data query experiments in concurrent query transactions with our system.

records already inserted. It is obvious that the insertion rate will decrease dramatically (by about 4 times) after 500 000 traffic records are inserted, and the more records inserted, the slower the insertion rate becomes. Compared with Oracle database, our HBase shows very consistent performance, and the peak insertion rate reaches approximately 100 000 records per second as shown in Fig. 5.

Second, we compare the query rate of traffic records in our home-brewed database with Oracle database. As shown in Fig. 6 for Oracle database, the higher the number of concurrent query transactions conducted, the slower the average response time becomes. The impact of the size of records in the database has a deleterious effect on the query performance. However, for our optimized HBase system, the latency of most responses is in milliseconds, and the impact of the records already in the database is quite low compared with Oracle database, which is shown in Fig. 7.

Our work is optimized based on the open source nature of HBase, while Oracle database is a proprietary one where we cannot optimize the code to speed up transactions in the traffic records repository. The performance gain is not achieved without pain, because the optimization method is also applied in dynamic multi-dimensional load balancing, fast disk access and batching writing, data block allocation and access, incremental full-text indexing, multiple distributed statistical model and hierarchical aggregation SQL execution engine, etc. Due to space constraints, we will provide more details in the future.

6 Conclusions

Mobile Internet represents a major opportunity to transform a telecom operator to become a Big Data operator. To achieve this, some obstacles need to be overcome. China Unicom takes the lead to embrace the Mobile Internet Explosion, and builds a big data platform to solve the challenges of data acquisition, data analysis, and data value-added services. Compared with the proprietary solution, the open source solution adopted by China Unicom offers us more advantages to optimize data storage, speed up database transactions, and achieve better performance. With this achievement, a telecom operator centric ecosystem that is critical for a telecom business to prosper in today's mobile data society is formed.

Acknowledgements

This work was supported in part by the National Key Basic Research and Development (973) Program of China (Nos. 2013CB228206 and 2012CB315801), the National Natural Science Foundation of China (Nos. 61233016 and 61140320). This work was also supported by the Intel Research Council under the title of "Security Vulnerability Analysis Based on Cloud Platform with Intel IA Architecture".

References

[1] J. Manyika, M. Chui, B. Brown, J. Bughin, R. Dobbs, C. Roxburgh, and A. H. Byers, Big data: The next frontier for innovation, competition, and productivity, Technical report, McKinsey Global Institute, 2011.

[2] M. Meeker and L. Wu, Mary Meeker's 2013 Internet Trends report, http://www.kpcb.com/insights/2013-internet-trends, 2013.

[3] Cisco Inc., Visual networking index: Forecast and methodology, 2010-2015, http://apo.org.au/research/cisco-visual-networking-index-forecast-and-methodology-2010-2015, 2011.

[4] T. Li, F. Y. Han, S. Ding, and Z. Chen, LARX: Large-scale anti-phishing by retrospective data-exploring based on a cloud computing platform, in *Proc. 20th International Conference on Computer Communications and Networks (ICCCN)*, Hawaii, USA, 2011, pp. 1-5.

[5] Z. Chen, F. Y. Han, J. W. Cao, X. Jiang, and S. Chen, Cloud computing-based forensic analysis for collaborative network security management system, *Tsinghua Science and Technology*, vol. 18, no. 1, pp. 40-50, 2013.

[6] Z. Chen, L. Y. Ruan, J. W. Cao, Y. Yu, and X. Jiang, TIFAflow: Enhancing traffic archiving system with flow granularity for forensic analysis in network security, *Tsinghua Science and Technology*, vol. 18, no. 4, pp. 406-417, 2013.

[7] S. Kornexl, V. Paxson, H. Dreger, A. Feldmann, and R. Sommer, Building a time machine for efficient recording and retrieval of high-volume network traffic, in *Proc. 5th ACM SIGCOMM Conference on Internet Measurement*, 2005.

[8] G. Maier, R. Sommer, H. Dreger, A. Feldmann, V. Paxson, and F. Schneider, Enriching network security analysis with time travel, in *Proc. ACM SIGCOMM 2008 Conference on Data Communication*, New York, USA, 2008, pp. 183-194.

[9] J. Li, S. Ding, M. Xu, F. Y. Han, X. Guan, and Z. Chen, TIFA: Enabling real-time querying and storage of massive stream data, in *Proc. International Conference on Networking and Distributed Computing (ICNDC)*, Osaka, Japan, 2011, pp. 61-64.

[10] K. Wu, S. Ahern, E. W. Bethel, J. Chen, H. Childs, E. Cormier-Michel, C. Geddes, J. Gu, H. Hagen, B. Hamann, W. Koegler, J. Lauret, J. Meredith, P. Messmer, E. Otoo, V. Perevoztchikov, A. Poskanzer, Prabhat, O. Rbel, A. Shoshani, A. Sim, K. Stockinger,

G. Weber, and W. M. Zhang, FastBit: Interactively searching massive data, *Journal of Physics: Conference Series*, vol. 180, no. 1, 2009.

[11] L. Deri and F. Fusco, MicroCloud-based network traffic monitoring, in *IFIP/IEEE International Symposium on Integrated Network Management (IM)*, 2013.

[12] L. Deri and F. Fusco, Real-time MicroCloud-based flow aggregation for fixed and mobile networks, in *9th IEEE International Wireless Communications and Mobile Computing Conference (IWCMC)*, 2013, pp. 96-101.

[13] Y. Lee and Y. Lee, Toward scalable internet traffic measurement and analysis with Hadoop, *ACM SIGCOMM Computer Communication Review*, vol. 43, no. 1, pp. 5-13, 2012.

[14] Y. Lee, W. Kang, and H. Son, An internet traffic analysis method with mapreduce, in *Network Operations and Management Symposium Workshops (NOMS Wksps)*, 2010, pp. 357-361.

[15] F. Qian, Z. Wang, Y. Gao, J. Huang, A. Gerber, Z. Mao, S. Sen, and O. Spatscheck, Periodic transfers in mobile applications: Network-wide origin, impact, and optimization, in *Proc. 21st International Conference on World Wide Web*, Lyon, France, 2012, pp. 51-60.

[16] J. L. Gailly and M. Adler, The ZLIB library, http://www.zlib.org/, 2013.

[17] M. F. Oberhumer, Lzo documentation, http://www.oberhumer.com/opensource/lzo/lzodoc.php, 2011.

Wenliang Huang is an associate dean of China Unicom Groups Labs. He received his MEng and PhD degrees from Zhejiang University in 1990 and 2008, respectively. Dr. Huang takes charge of Big Data Platform for all China Unicom. He has applied 20 patents and granted 4 patents related with Big Data Platform in all China Unicom groups. He has leaded several national programs, such as 973 plan, 863 program, NSFC, etc. He has been awarded national award for science and technology progress in 2013. His research interests include network architecture, mobile network, and big data.

Zhen Chen is an associate professor in Research Institute of Information Technology in Tsinghua University. He received his BEng and PhD degrees from Xidian University in 1998 and 2004, respectively. He once worked as postdoctoral researcher in Network Institute of Department of Computer Science and Technology in Tsinghua University during 2004 to 2006. He was also a visiting scholar in UC Berkeley ICSI in 2006. He joined Research Institute of Information Technology in Tsinghua University since 2006. His research interests include network architecture, computer security, and data analysis. He has published around 80 academic papers.

Wenyu Dong is currently a master student of Department of Computer Science and Technology at Tsinghua University. He got his BEng degree from PLA Univ. of Info. & Eng. in 2009. Currently his main research interests focus on computer network security and mobile safety.

Hang Li is currently a master student of Department of Computer Science and Technology at PLA Univ. of Info. & Eng. She got her BEng degree from PLA Univ. of Info. & Eng. in 2013. Her main research interest focuses on computer network security.

Bin Cao is currently a master student of Department of Computer Science and Technology at Tsinghua University. He got his BEng degree from PLA Univ. of Info. & Eng. in 2005. His research interests include computer network security and mobile safety.

Junwei Cao is currently a professor and deputy director of Research Institute of Information Technology, Tsinghua University, China. He is also the director of Open Platform and Technology Division, Tsinghua National Laboratory for Information Science and Technology. His research is focused on advanced computing technology and applications. Before joining Tsinghua in 2006, Junwei Cao was a research scientist of Massachusetts Institute of Technology, USA. Before that he worked as a research staff member of NEC Europe Ltd., Germany. Junwei Cao got his PhD degree in computer science from University of Warwick, UK in 2001. He got his MEng and BEng degrees from Tsinghua University in 1998 and 1996, respectively. Junwei Cao has published over 130 academic papers and books, cited by international researchers for over 3000 times. Junwei Cao is the senior member of the IEEE Computer Society and the member of the ACM and CCF.