Cloud computing for java/ dot net

Check following Projects ,also check if any spelling mistakes before showing to your Guide:

Data integrity and security using cloud computing

Abstract : Cloud computing has been envisioned as the de-facto solution to the rising storage costs of IT Enterprises. With the high costs of data storage devices as well as the rapid rate at which data is being generated it proves costly for enterprises or individual users to frequently update their hardware. Apart from reduction in storage costs data outsourcing to the cloud also helps in reducing the maintenance. Cloud storage moves the user's data to large data centers, which are remotely located, on which user does not have any control. However, this unique feature of the cloud poses many new security challenges which need to be clearly understood and resolved. We provide a scheme which gives a proof of data integrity in the cloud which the customer can employ to check the correctness of his data in the cloud. This proof can be agreed upon by both the cloud and the customer and can be incorporated in the Service level agreement (SLA).

Extra security using graphical password to cloud

ABSTRACT : Many security primitives are based on hard mathematical problems. Using hard AI problems for security is emerging as an exciting new paradigm, but has been under-explored. In this paper, we present a new security primitive based on hard AI problems, namely, a novel family of graphical password systems built on top of Captcha technology, which we call Captcha as graphical passwords (CaRP). CaRP is both a Captcha and a graphical password scheme. CaRP addresses a number of security problems altogether, such as online guessing attacks, relay attacks, and, if combined with dual-view technologies, shoulder-surfing attacks. Notably, a CaRP password can be found only probabilistically by automatic online guessing attacks even if the password is in the search set. CaRP also offers a novel approach to address the well-known image hotspot problem in popular graphical password systems, such as PassPoints, that often leads to weak password choices. CaRP is not a panacea, but it offers reasonable security and usability and appears to fit well with some practical applications for improving online security.

A Distributed Video Management Cloud Platform

ABSTRACT : The rapidly increasing power of personal mobile devices is providing much richer contents and social interactions to users on the move. This trend however is throttled by the limited battery lifetime of mobile devices and unstable wireless connectivity, making the highest possible quality of service experienced by mobile users not feasible. The recent cloud computing technology, with its rich resources to compensate for the limitations of mobile devices and connections, can potentially provide an ideal platform to support the desired mobile services. Tough challenges arise on how to effectively exploit cloud resources to facilitate mobile services, especially those with stringent interaction delay requirements. In this paper, we propose the design of a Cloud-based, novel Mobile social TV system.

Cloud computing single cloud to multi cloud system.

Abstract : With the rise of various cloud services, the problem of redundant data is more prominent in the cloud storage systems. How to assign a set of documents to a distributed file system, which can not only reduce storage space, but also ensure the access efficiency as much as possible, is an urgent problem which needs to be solved. Space-efficiency mainly uses data de-duplication technologies, while access-efficiency requires gathering the files with high similarity on a server. Based on the study of other data de-duplication technologies, especially the Similarity-Aware Partitioning (SAP) algorithm, this paper proposes the Frequency and Similarity-Aware Partitioning (FSAP) algorithm for cloud storage. The FSAP algorithm is a more reasonable data partitioning algorithm than the SAP algorithm. Meanwhile, this paper proposes the Space-Time Utility Maximization Model (STUMM), which is useful in balancing the relationship between space-efficiency and access-efficiency. Finally, this paper uses 100 web files downloaded from CNN for testing, and the results show that, relative to using the algorithms associated with the SAP algorithm (including the SAP-Space-Delta algorithm and the SAP-Space-Dedup algorithm), the FSAP algorithm based on STUMM reaches higher compression ratio and a more balanced distribution of data blocks.

Visual cryptography for biometric privacy in cloud.

Abstract — In this era due to unbelievable development in internet, various online attacks has been increased. From all such attacks most popular attack is phishing. This attacks are done for extracting confidential information such as banking information, passwords from unsuspecting victims for fraud purposes. Confidential data can't be directly uploaded on website since it is risky. Here in this paper data is encrypted in video and visual cryptography for login purpose in our online database system for providing more security .

Eye Retina Authentication Cloud Computing.

Abstract— Cloud computing has been envisioned as the next-generation architecture of IT enterprise. In contrast to traditional solutions, where the IT services are under proper physical, logical and personnel controls, cloud computing moves the application software and databases to the large data centers, where the management of the data and services may not be fully trustworthy. This unique attribute, however, poses many new security challenges which have not been well understood. In this article, we focus on cloud secure data storage, which has always been an important aspect of quality of service (QOS). To ensure the secure storage of u ser's data in the cloud, we propose an effective and flexible bio metric authentication using face.

Privacy Preserving and Delegated Access Control for Cloud Applications.

ABSTRACT : Data access control is an effective way to ensure the data security in the cloud. Due to data outsourcing and untrusted cloud servers, the data access control becomes a challenging issue in cloud storage systems. Ciphertext-Policy Attributebased Encryption (CP-ABE) is regarded as one of the most suitable technologies for data access control in cloud storage, because it gives data owners more direct control on access policies. However, it is difficult to directly apply existing CP-ABE schemes to data access control for cloud storage systems because of the attribute revocation problem. In this paper, we design an expressive, efficient and revocable data access control scheme for multi-authority cloud storage systems, where there are multiple authorities co-exist and each authority is able to issue attributes independently. Specifically, we propose a revocable multi-authority CP-ABE scheme, and apply it as the underlying techniques to design the data access control scheme. Our attribute revocation method can efficiently achieve both forward security and backward security. The analysis and simulation results show that our proposed data access control scheme is secure in the random oracle model and is more efficient than previous works.

Privacy aware data aggregation in mobile sensing.

ABSTRACT: The proliferation and ever-increasing capabilities of mobile devices such as smart phones give rise to a variety of mobile sensing applications. This paper studies how an untrusted aggregator in mobile sensing can periodically obtain desired statistics over the data contributed by multiple mobile users, without compromising the privacy of each user. Although there are some

existing works in this area, they either require bidirectional communications between the aggregator and mobile users in every aggregation period, or have high-computation overhead and cannot support large plaintext spaces. Also, they do not consider the Min aggregate, which is quite useful in mobile sensing. To address these problems, we propose an efficient protocol to obtain the Sum aggregate, which employs an additive homomorphic encryption and a novel key management technique to support large plaintext space. We also extend the sum aggregation protocol to obtain the Min aggregate of time-series data. To deal with dynamic joins and leaves of mobile users, we propose a scheme that utilizes the redundancy in security to reduce the communication cost for each join and leave. Evaluations show that our protocols are orders of magnitude faster than existing solutions, and it has much lower communication overhead.

Web Service QoS Prediction Based on Adaptive Dynamic Programming Using Fuzzy Neural Networks for Cloud Services.

ABSTRACT : Cloud data center management is a key problem due to the numerous and heterogeneous strategies that can be applied, ranging from the VM placement to the federation with other clouds. Performance evaluation of cloud computing infrastructures is required to predict and quantify the cost-benefit of a strategy portfolio and the corresponding quality of service (QoS) experienced byusers. Such analyses are not feasible by simulation or on-the-field experimentation, due to the great number of parameters that have to be investigated. In this paper, we present an analytical model, based on stochastic reward nets (SRNs), that is both scalable to model systems composed of thousands of resources and flexible to represent different policies and cloud-specific strategies. Several performance metrics are defined and evaluated to analyze the behavior of a cloud data center: utilization, availability, waiting time, and responsiveness. A resiliency analysis is also provided to take into account load bursts. Finally, a general approach is presented that,starting from the concept of system capacity, can help system managers to opportunely set the data center parameters under different working conditions

Online java complier in cloud computing.

Abstract : As it is a competitive world and very fast world, every thing in the universes is to be internet. In this internet world all the things are on-line. So we created software called “On-line java compiler with security editor”.

The main aim of this project we can easily to write a java program and compile it and debug in on-line. The client machine doesn't having java development kit .The client machine only connected to the server. The server having java compiler .so server executes the java code and produce the error message to the appropriate client machine.

In this project is also creating a security editor. This editor performs Encrypt and decrypts the file. Encryption and decryption process perform using RSA Algorithms. There is lot of security algorithms are there, but RSA algorithm is very efficient to encrypt and decrypt the file.

In this project is used to view all type of java API .It is very useful for writing the java program easily, for example if any error in the format of API means we can able to view API throw this modules

Frequency and Similarity-Aware Partitioning for Cloud Storage Based on space time utility maximization model.

Abstract:: With the rise of various cloud services, the problem of redundant data is more prominent in the cloud storage systems. How to assign a set of documents to a distributed file system, which can not only reduce storage space, but also ensure the access efficiency as much as possible, is an urgent problem which needs to be solved. Space-efficiency mainly uses data de-duplication technologies, while access-efficiency requires gathering the files with high similarity on a server. Based on the study of other data de-duplication technologies, especially the Similarity-Aware Partitioning (SAP) algorithm, this paper proposes the Frequency and Similarity-Aware Partitioning (FSAP) algorithm for cloud storage. The FSAP algorithm is a more reasonable data partitioning algorithm than the SAP algorithm. Meanwhile, this paper proposes the Space-Time Utility Maximization Model (STUMM), which is

useful in balancing the relationship between space-efficiency and access-efficiency. Finally, this paper uses 100 web files downloaded from CNN for testing, and the results show that, relative to using the algorithms associated with the SAP algorithm (including the SAP-Space-Delta algorithm and the SAP-Space-Dedup algorithm), the

FSAP algorithm based on STUMM reaches higher compression ratio and a more balanced distribution of data blocks

An Inline High Performance Deduplication System Used in Cloud Storage.

Abstract: : Data deduplication is an emerging and widely employed method for current storage systems. As this technology is gradually applied in inline scenarios such as with virtual machines and cloud storage systems, this study proposes a novel deduplication architecture called I-sieve. The goal of I-sieve is to realize a high performance data sieve system based on iSCSI in the cloud storage system. We also design the corresponding index and mapping tables and present a multi-level cache using a solid state drive to reduce RAM consumption and to optimize lookup performance. A prototype of I-sieve is implemented based on the open source iSCSI target, and many experiments have been conducted driven by virtual machine images and testing tools. The evaluation results show excellent deduplication and foreground performance. More importantly, I-sieve can co-exist with the existing deduplication systems as long as they support the iSCSI protocol.

Data Security in Cloud Architecture Based Elliptical Curve Cryptography.

Abstract — Secure and efficient data storage is needed in the cloud environment in modern era of information technology industry. In the present scenario the cloud verifies the authenticity of the cloud services without the knowledge of user's identity. The cloud provides massive data access directly through the internet. Centralized storage mechanism is followed here for effective accessing of data. Cloud service providers are normally acquires the software and hardware resources and the cloud consumers are avail the services through the internet access in lease basis. Cloud security is enhanced through cryptography technique applied to the cloud security to avoid vulnerability. The intractable computability is achieved in the cloud by using the public key cryptosystem. This paper proposed the approach of applying Hyper elliptic curve cryptography for data protection in the cloud with the small key size. The proposed system has the further advantage of eliminating intruder in cloud computing. Efficacy of the system is to provide the high security of the cloud data.

A Cloud Service Architecture for Analyzing Big Monitoring Data.

Abstract : Cloud monitoring is of a source of big data that are constantly produced from traces of infrastructures, platforms, and applications. Analysis of monitoring data delivers insights of the system's workload and usage pattern and ensures workloads are operating at optimum levels. The analysis process involves data query and

extraction, data analysis, and result visualization. Since the volume of monitoring data is big, these operations require a scalable and reliable architecture to extract, aggregate, and analyze data in an arbitrary range of granularity. Ultimately, the results of analysis become the knowledge of the system and should be shared and

communicated. This paper presents our cloud service architecture that explores a search cluster for data indexing and query. We develop REST APIs that the data can be accessed by different analysis modules. This architecture enables extensions to integrate with software frameworks of both batch processing (such as Hadoop) and stream processing (such as Spark) of big data. The analysis results are structured in Semantic Media Wiki pages in the context of the monitoring data source and the analysis process. This cloud architecture is empirically assessed to evaluate its responsiveness when processing a large set of data records under node failures.

Fog computing : A new concept By cisco.

ABSTRACT : Cloud computing promises to significantly change the way we use computers and access and store our personal and business information. With these new computing and communications paradigms arise new data security challenges. Existing data protection mechanisms such as encryption have failed in preventing data theft attacks, especially those perpetrated by an insider to the cloud provider. We propose a different approach for securing data in the cloud using offensive decoy technology. We monitor data access in the cloud and detect abnormal data access patterns. When unauthorized access is suspected and then verified using challenge questions, we launch a disinformation attack by returning large amounts of decoy information to the attacker. This protects against the misuse of the user's real data. Experiments conducted in a local file setting provide evidence that this approach may provide unprecedented levels of user data security in a Cloud environment.

Cloud Data Protection for the Masses.

Abstract : Offering strong data protection to cloud users while enabling rich applications is a challenging task. We explore a new cloud platform architecture called Data Protection as a Service, which dramatically reduces the per-application development effort required to offer data protection, while still allowing rapid development and maintenance.

Ensuring Distributed Accountability For Data Sharing In The Cloud.

Abstract : Cloud computing enables highly scalable services to be easily consumed over the Internet on an as-needed basis. A major feature of the cloud services is that users' data are usually processed remotely in unknown machines that users do not own or operate. While enjoying the convenience brought by this new emerging technology, users' fears of losing control of their own data (particularly, financial and health data) can become a significant barrier to the wide adoption of cloud

services. To address this problem, here, we propose a novel highly decentralized information accountability framework to keep track of the actual usage of the users' data in the cloud. In particular, we propose an object-centered approach that enables enclosing our logging mechanism together with users' data and policies. We leverage the JAR programmable capabilities to both create a dynamic and traveling object, and to ensure that any access to users' data will trigger authentication and automated logging local to the JARs. To strengthen user's control, we also provide distributed auditing mechanisms. We provide extensive experimental studies that demonstrate the efficiency and effectiveness of the proposed approaches.

Cooperative Provable Data Possession for Integrity Verification in Multi-Cloud Storage.

ABSTRACT

Provable data possession (PDP) is a technique for ensuring the integrity of data in storage outsourcing. In this paper, we address the construction of an efficient PDP scheme for distributed cloud storage to support the scalability of service and data migration, in which we consider the existence of multiple cloud service providers to cooperatively store and maintain the clients' data. We present a cooperative PDP (CPDP) scheme based on homomorphic verifiable response and hash index hierarchy. We prove the security of our scheme based on multi-prover zero-knowledge proof system, which can satisfy completeness, knowledge soundness, and zero-knowledge properties. In addition, we articulate performance optimization mechanisms for our scheme, and in particular present an efficient method for selecting optimal parameter values to minimize the computation costs of clients and storage service providers. Our experiments show that our solution introduces lower computation and communication overheads in comparison with non-cooperative approaches.

A Secure Erasure Code-Based Cloud Storage System with Secure Data Forwarding.

ABSTRACT: A cloud storage system, consisting of a collection of storage servers, provides long-term storage services over the Internet. Storing data in a third party's cloud system causes serious concern over data confidentiality. General encryption schemes protect data confidentiality, but also limit the functionality of the storage system because a few operations are supported over encrypted data. Constructing a secure storage system that supports multiple functions is challenging when the storage system is distributed and has no central authority. We propose a threshold proxy re-encryption scheme and integrate it with a decentralized erasure code such that a secure distributed storage system is formulated. The distributed storage system not only supports secure and robust data storage and retrieval, but also lets a user forward his data in the storage servers to another user

without retrieving the data back. The main technical contribution is that the proxy re-encryption scheme supports encoding operations over encrypted messages as well as forwarding operations over encoded and encrypted messages. Our method fully integrates encrypting, encoding, and forwarding. We analyze and suggest suitable parameters for the number of copies of a message dispatched to storage servers and the number of storage servers queried by a key server. These parameters allow more flexible adjustment between the number of storage servers and robustness.

Scalable and Secure Sharing of Personal Health Records in Cloud Computing using Attribute-based Encryption.

ABSTRACT: Personal health record (PHR) is an emerging patient-centric model of health information exchange, which is often outsourced to be stored at a third party, such as cloud providers. However, there have been wide privacy concerns as personal health information could be exposed to those third party servers and to unauthorized parties. To assure the patients' control over access to their own PHRs, it is a promising method to encrypt the PHRs before outsourcing. Yet, issues such as risks of privacy exposure, scalability in key management, flexible access and efficient user revocation, have remained the most important challenges toward achieving fine-grained, cryptographically enforced data access control. In this paper, we propose a novel patient-centric framework and a suite of mechanisms for data access control to PHRs stored in semi-trusted servers. To achieve fine-grained and scalable data access control for PHRs, we leverage attribute based encryption (ABE) techniques to encrypt each patient's PHR file. Different from previous works in secure data outsourcing, we focus on the multiple data owner scenario, and divide the users in the PHR system into multiple security domains that greatly reduces the key management complexity for owners and users. A high degree of patient privacy is guaranteed simultaneously by exploiting multi-authority ABE. Our scheme also enables dynamic modification of access policies or file attributes, supports efficient on-demand user/attribute revocation and break-glass access under emergency scenarios. Extensive analytical and experimental results are presented which show the security, scalability and efficiency of our proposed scheme

Portable Organization Cloud Services .

ABSTRACT : Underneath all the hype, the essence of cloud computing is the industrialization of IT. Similar to mass production lines in other industries (such as the auto industry), cloud computing standardizes offered services and thus increases automation significantly. Consequently, enterprises are increasingly utilizing cloud technology; however, major challenges such as portability, standardization of service definitions, and security remain inadequately addressed. The ability to move cloud services and their components between providers ensures an adequate and cost-efficient IT environment and avoids vendor lock-in. Research has already addressed movability and migration

on a functional level.1,2 However, no one has yet examined cloud service portability with regard to management and operational tasks, which are a significant and increasing cost factor. One reason is the lack of an industry standard for defining composite applications and their management. Without an appropriate standardized format, ensuring compliance, trust, and security the biggest area of critique preventing the cloud's wider adoption is difficult. Dealing with these challenges in industry and research has the potential to bring cloud computing to the next level. Here, we present how the portable and standardized management of cloud services is enabled through the Topology and Orchestration Specification for Cloud Applications (TOSCA),3 a recently initiated standardization effort from OASIS. We show how TOSCA plans which capture the management aspects of cloud services in a reusable way use existing workflow technologies and research results to facilitate the portable, automated, and reusable management of cloud services throughout their life cycle.

Fog Computing: Mitigating Insider Data Theft Attacks in the Cloud.

ABSTRACT

Cloud computing promises to significantly change the way we use computers and access and store our personal and business information. With these new computing and communications paradigms arise new data security challenges. Existing data protection mechanisms such as encryption have failed in preventing data theft attacks, especially those perpetrated by an insider to the cloud provider. We propose a different approach for securing data in the cloud using offensive decoy technology. We monitor data access in the cloud and detect abnormal data access patterns. When unauthorized access is suspected and then verified using challenge questions, we launch a disinformation attack by returning large amounts of decoy information to the attacker. This protects against the misuse of the user's real data. Experiments conducted in a local file setting provide evidence that this approach may provide unprecedented levels of user data security in a Cloud environment.

A Cloud Gaming System Based on User-Level Virtualization and Its Resource Scheduling.

Abstract —Many believe the future of gaming lies in the cloud, namely Cloud Gaming, which renders an interactive gaming application in the cloud and streams the scenes as a video sequence to

the player over Internet. This paper proposes GCloud, a GPU/CPU hybrid cluster for cloud gaming based on the user-level virtualization technology. Specially, we present a performance model to analyze the server-capacity and games' resource-consumptions, which categorizes games into two types: CPU-critical and memory-io-critical. Consequently, several scheduling strategies have been proposed to improve the resource-utilization and compared with others. Simulation tests show that both of the First-Fit-like and the Best-Fit-like strategies outperform the other(s); especially they are near optimal in the batch processing mode. Other test results indicate that GCloud is efficient: An off-the-shelf PC can support five high-end video-games run at the same time. In addition, the average per-frame processing delay is 819 ms under different image-resolutions, which outperforms other similar solutions.

A Mobile Offloading Game Against Smart Attacks.

ABSTRACT : Mobile devices, such as smartphones, can ofoad applications and data to the cloud via access points or base stations to reduce energy consumption and improve user experience. However, mobile ofoading is vulnerable to smart attackers that use smart and programmable radio devices, such as universal software radio peripherals, to perform multiple types of attacks, such as spoong and jamming, based on the radio environment and ofoading transmissions. In this paper, a mobile ofoading game is investigated that consists of three players: a mobile device that chooses its ofoading rate, a smart attacker that determines its attack mode, and a security agent that decides whether or not to initiate full protection for the serving access point during the ofoading. Nash and Stackelberg equilibria of the ofoading game are derived and their existence conditions are discussed. A Q-learning-based mobile ofoading strategy is proposed for mobile devices that are unaware of system parameters, such as the channel conditions, in dynamic radio environments. Simulation results show that the proposed ofoading strategy can improve the utility of the mobile device and reduce the attack rate of smart attackers.

A Parallel Patient Treatment Time Prediction Algorithm and Its Applications in Hospital Queuing-Recommendation in a Big Data Environment.

ABSTRACT: Effective patient queue management to minimize patient wait delays and patient overcrowding is one of the major challenges faced by hospitals. Unnecessary and annoying waits for long periods result in substantial human resource and time wastage and increase the frustration endured by patients. For each patient in the queue, the total treatment time of all the patients before him is the time that he must wait. It would be convenient and preferable if the patients could receive

the most efcient treatment plan and know the predicted waiting time through a mobile application that updates in real time. Therefore, we propose a Patient Treatment Time Prediction (PTTP) algorithm to predict the waiting time for each treatment task for a patient. We use realistic patient data from various hospitals to obtain a patient treatment time model for each task. Based on this large-scale, realistic dataset, the treatment time for each patient in the current queue of each task is predicted. Based on the predicted waiting time, a Hospital Queuing-Recommendation (HQR) system is developed. HQR calculates and predicts an efcient and convenient treatment plan recommended for the patient. Because of the large-scale, realistic dataset and the requirement for real-time response, the PTTP algorithm and HQR system mandate efciency and low-latency response. We use an Apache Spark-based cloud implementation at the National Supercomputing Center in Changsha to achieve the aforementioned goals. Extensive experimentation and simulation results demonstrate the effectiveness and applicability of our proposed model to recommend an effective treatment plan for patients to minimize their wait times in hospitals.

A Tutorial on Secure Outsourcing of Large-scale Computations for Big Data.

ABSTRACT: Today's society is collecting a massive and exponentially growing amount of data that can potentially revolutionize scientic and engineering elds, and promote business innovations.With the advent of cloud computing, in order to analyze data in a cost-effective and practical way, users can outsource their computing tasks to the cloud, which offers access to vast computing resources on an on-demand and pay-per-use basis. However, since users' data contains sensitive information that needs to be kept secret for ethical, security, or legal reasons, many users are reluctant to adopt cloud computing. To this end, researchers have proposed techniques that enable users to ofoad computations to the cloud while protecting their data privacy. In this paper, we review the recent advances in the secure outsourcing of large-scale computations for a big data analysis. We rst introduce two most fundamental and common computational problems, i.e., linear algebra and optimization, and then provide an extensive review of the data privacy preserving techniques. After that, we explain how researchers have exploited the data privacy preserving techniques to construct secure outsourcing algorithms for large-scale computations.

An Anomalous Behavior Detection Model in Cloud Computing.

Abstract: This paper proposes an anomalous behavior detection model based on cloud computing. Virtual Machines (VMs) are one of the key components of cloud Infrastructure as a Service (IaaS). The security of such VMs is critical to IaaS security. Many studies have been done on cloud computing security issues, but research into VM security issues, especially regarding VM network

traffic anomalous behavior detection, remains inadequate. More and more studies show that communication among internal nodes exhibits complex patterns. Communication among VMs in cloud computing is invisible. Researchers find such issues challenging, and few solutions have been proposed—leaving cloud computing vulnerable to network attacks. This paper proposes a model that uses Software-Defined Networks (SDN) to implement traffic redirection. Our model can capture inter-VM traffic, detect known and unknown anomalous network behaviors, adopt hybrid techniques to analyze VM network behaviors, and control network systems. The experimental results indicate that the effectiveness of our approach is greater than 90%, and prove the feasibility of the model.

An Efficient Privacy-Preserving Ranked Keyword Search Method.

Abstract —Cloud data owners prefer to outsource documents in an encrypted form for the purpose of privacy preserving. Therefore it is essential to develop efficient and reliable ciphertext search techniques. One challenge is that the relationship between documents will be normally concealed in the process of encryption, which will lead to significant search accuracy performance degradation. Also the volume of data in data centers has experienced a dramatic growth. This will make it even more challenging to design ciphertext search schemes that can provide efficient and reliable online information retrieval on large volume of encrypted data. In this paper, a hierarchical clustering method is proposed to support more search semantics and also to meet the demand for fast ciphertext search within a big data environment. The proposed hierarchical approach clusters the documents based on the minimum relevance threshold, and then partitions the resulting clusters into sub-clusters until the constraint on the maximum size of cluster is reached. In the search phase, this approach can reach a linear computational complexity against an exponential size increase of document collection. In order to verify the authenticity of search results, a structure called minimum hash sub-tree is designed in this paper. Experiments have been conducted using the collection set built from the IEEE Xplore. The results show that with a sharp increase of documents in the dataset the search time of the proposed method increases linearly whereas the search time of the traditional method increases exponentially. Furthermore, the proposed method has an advantage over the traditional method in the rank privacy and relevance of retrieved documents.

Analysis of Classical Encryption Techniques in Cloud Computing.

Abstract: Cloud computing has become a significant computing model in the IT industry. In this emerging model, computing resources such as software, hardware, networking, and storage can be accessed anywhere in the world on a pay-per-use basis. However, storing sensitive data on un-trusted

servers is a challenging issue for this model. To guarantee confidentiality and proper access control of outsourced sensitive data, classical encryption techniques are used. However, such access control schemes are not feasible in cloud computing because of their lack of flexibility, scalability, and fine-grained access control. Instead, Attribute-Based Encryption (ABE) techniques are used in the cloud. This paper extensively surveys all ABE schemes and creates a comparison table for the key criteria for these schemes in cloud applications.

Bayes-Based ARP Attack Detection Algorithm for Cloud Centers.

Abstract: To address the issue of internal network security, Software-Defined Network (SDN) technology has been introduced to large-scale cloud centers because it not only improves network performance but also deals with network attacks. To prevent man-in-the-middle and denial of service attacks caused by an address resolution protocol bug in an SDN-based cloud center, this study proposed a Bayes-based algorithm to calculate the probability of a host being an attacker and further presented a detection model based on the algorithm. Experiments were conducted to validate this method.

Cloud-Assisted IoT-Based SCADA Systems Security: A Review of the State of the Art and Future Challenges.

ABSTRACT : Industrial systems always prefer to reduce their operational expenses. To support such reductions, they need solutions that are capable of providing stability, fault tolerance, and exibility. One such solution for industrial systems is cyber physical system (CPS) integration with the Internet of Things (IoT) utilizing cloud computing services. These CPSs can be considered as smart industrial systems, with their most prevalent applications in smart transportation, smart grids, smart medical and eHealthcare systems, and many more. These industrial CPSs mostly utilize supervisory control and data acquisition (SCADA) systems to control and monitor their critical infrastructure (CI). For example,WebSCADA is an application used for smart medical technologies, making improved patient monitoring and more timely decisions possible. The focus of the study presented in this paper is to highlight the security challenges that the industrial SCADA systems face in an IoT-cloud environment. Classical SCADA systems are already lacking in proper security measures; however, with the integration of complex new architectures for the future Internet based on the concepts of IoT, cloud computing, mobile wireless sensor networks, and so on, there are large issues at stakes in the security and deployment of these classical systems. Therefore, the integration of these future Internet concepts needs more research effort. This paper, along with highlighting the security challenges of these CI's, also provides the existing best practices and recommendations for improving

and maintaining security. Finally, this paper briey describes future research directions to secure these critical CPSs and help the research community in identifying the research gaps in this regard.

Efficient Multicast Delivery for Data redundancy minimization over wireless data centers.

ABSTRACT: With the explosive growth of cloud-based services, large-scale data centers are widely built for housing critical computing resources to gain signicant economic benets. In data centers, the cloud services are generally accomplished by multicast-based group communications. Recently, many well-known industries, such as Microsoft, Google, and IBM, adopt high-speed wireless technologies to augment network capacity in data centers. However, those well-known multicast delivery schemes for traditional wired data centers do not consider the unique characteristics of wireless communications, which may result in unnecessary data transmissions and network congestions. Under the coexisting scenario of wired and wireless links, this paper studies multicast tree construction and maintenance problems. The objective is to minimize the total multicast trafc. We prove the problems are NP-hard and propose efcient heuristic algorithms for the two problems. Based on real traces and practical settings obtained from commercial data centers, a series of experiments are conducted, and the experimental results showthat our proposed algorithms are effective for reducing multicast data trafc. The results also provide useful insights into the design of multicast tree construction and maintenance for wireless data center networks.

Multiagent-Multiobjective Interaction Game System for Service provisoning vehicular cloud.

Abstract — The increasing number of applications based on the Internet of Things (IoT), as well as advances in wireless communication, information and communication technology, and mobile cloud computing, has allowed mobile users to access a wider range of resources when mobile. As the use of vehicular cloud computing has become more popular due to its ability to improve driver and vehicle safety, researchers and industry have a growing interest in the design and development of vehicular networks for emerging applications. Vehicle drivers can now access a variety of on demand resources en route via vehicular network service providers. The adaptation of vehicular cloud services faces many challenges, including cost, privacy and latency. The contributions of this paper are as follows: First, we propose a game theory-based framework to manage on-demand service provision in a vehicular cloud. We present three different game approaches, each of which helps drivers minimize their service costs and latency, and maximize their privacy. Secondly, we propose a Quality-of-Experience (QoE) framework for service provision in a vehicular cloud for various types of users; a simple but effective model to determine driver preferences. Third, we propose using the

Trusted Third Party (TTP) concept to represent drivers and service providers, and ensure fair game treatment. We develop and evaluate simulations of the proposed approaches under different network scenarios with respect to privacy, service cost and latency, by varying the vehicle density and driver preferences. The results show that the proposed approach outperforms conventional models, since the game theory system introduces a bounded latency of ≤ 3%, achieves service cost savings up to 65%, and preserves driver privacy by reducing revealed information by up to 47%.

Toward a Real-Time Framework in Cloudlet-Based Architecture.

Abstract: In this study, we present a framework based on a prediction model that facilitates user access to a number of services in a smart living environment. Users must be able to access all available services continuously equipped with mobile devices or smart objects without being impacted by technical constraints such as performance or memory issues, regardless of their physical location and mobility. To achieve this goal, we propose the use of cloudlet-based architecture that serves as distributed cloud resources with specific ranges of influence and a real time processing framework that tracks events and preferences of the end consumers, predicts their requirements, and recommends services to optimize resource utilization and service response time.

Towards a Virtual Domain Based Authentication on MapReduce.

ABSTRACT : This paper has proposed a novel authentication solution for the MapReduce (MR) model, a new distributed and parallel computing paradigm commonly deployed to process BigData by major IT players, such as Facebook and Yahoo. It identies a set of security, performance, and scalability requirements that are specied from a comprehensive study of a job execution process using MR and security threats and attacks in this environment. Based on the requirements, it critically analyzes the state-of-the-art authentication solutions, discovering that the authentication services currently proposed for the MR model is not adequate. This paper then presents a novel layered authentication solution for the MR model and describes the core components of this solution, which includes the virtual domain based authentication framework (VDAF). These novel ideas are signicant, because, rst, the approach embeds the characteristics of MR-in-cloud deployments into security solution designs, and this will allow the MR model be delivered as a software as a service in a public cloud environment along with our proposed authentication solution; second, VDAF supports the authentication of every interactions by any MR components involved in a job execution ow, so long as the interactions are for accessing resources of the job; third, this continuous authentication

service is provided in such a manner that the costs incurred in providing the authentication service should be as low as possible.

Validation of Pervasive Cloud Task Migration with Colored Petri Net.

Abstract: Mobile devices are resource-limited, and task migration has become an important and attractive feature of mobile clouds. To validate task migration, we propose a novel approach to the simulation of task migration in a pervasive cloud environment. Our approach is based on Colored Petri Net (CPN). In this research, we expanded the semantics of a CPN and created two task migration models with different task migration policies: one that took account of context information and one that did not. We evaluated the two models using CPN-based simulation and analyzed their task migration accessibility, integrity during the migration process, reliability, and the stability of the pervasive cloud system after task migration. The energy consumption and costs of the two models were also investigated. Our results suggest that CPN with context sensing task migration can minimize energy consumption while preserving good overall performance.

A Hybrid Cloud Approach for Secure Authorized Deduplication.

Data deduplication is one of important data compression techniques for eliminating duplicate copies of repeating data, and has been widely used in cloud storage to reduce the amount of storage space and save bandwidth. To protect the confidentiality of sensitive data while supporting deduplication, the convergent encryption technique has been proposed to encrypt the data before outsourcing. To better protect data security, this paper makes the first attempt to formally address the problem of authorized data deduplication. Different from traditional deduplication systems, the differential privileges of users are further considered in duplicate check besides the data itself. We also present several new deduplication constructions supporting authorized duplicate check in a hybrid cloud architecture. Security analysis demonstrates that our scheme is secure in terms of the definitions specified in the proposed security model. As a proof of concept, we implement a prototype of our proposed authorized duplicate check scheme and conduct testbed experiments using our prototype. We show that our proposed authorized duplicate check scheme incurs minimal overhead compared to normal operations.

AMES-Cloud: A Framework of Adaptive Mobile Video Streaming and Efficient Social Video Sharing in the Clouds.

Synopsis:

While demands on video traffic over mobile networks have been souring, the wireless link capacity cannot keep up with the traffic demand. The gap between the traffic demand and the link capacity, along with time-varying link conditions, results in poor service quality of video streaming over mobile networks such as long buffering time and intermittent disruptions. Leveraging the cloud computing technology, we propose a new mobile video streaming framework, dubbed AMES-Cloud, which has two main parts: adaptive mobile video streaming (AMoV) and efficient social video sharing (ESoV). AMoV and ESoV construct a private agent to provide video streaming services efficiently for each mobile user. For a given user, AMoV lets her private agent adaptively adjust her streaming flow with a scalable video coding technique based on the feedback of link quality. Likewise, ESoV monitors the social network interactions among mobile users, and their private agents try to prefetch video content in advance. We implement a prototype of the AMES-Cloud framework to demonstrate its performance. It is shown that the private agents in the clouds can effectively provide the adaptive streaming, and perform video sharing (i.e., prefetching) based on the social network analysis.

Cloud Computing Security From Single to Multi-Clouds.

Synopsis:

The use of cloud computing has increased rapidly in many organizations. Cloud computing provides many benefits in terms of low cost and accessibility of data. Ensuring the security of cloud computing is a major factor in the cloud computing environment, as users often store sensitive information with cloud storage providers but these providers may be untrusted. Dealing with "single cloud" providers is predicted to become less popular with customers due to risks of service availability failure and the possibility of malicious insiders in the single cloud. A movement towards "multi-clouds", or in other words, "interclouds" or "cloud-of-clouds" has emerged recently. This paper surveys recent research related to single and multi-cloud security and addresses possible solutions. It is found that the research into the use of multi-cloud providers to maintain security has received less attention from the research community than has the use of single clouds. This work aims to promote the use of multi-clouds due to its ability to reduce security risks that affect the cloud computing user.

CloudTPS Scalable Transactions for Web Applications in the Cloud.

Synopsis:

NoSQL cloud data stores provide scalability and high availability properties for web applications, but at the same time they sacrifice data consistency. However, many applications cannot afford any data inconsistency. CloudTPS is a scalable transaction manager which guarantees full ACID properties for multi-item transactions issued by web applications, even in the presence of server failures and network partitions. We implement this approach on top of the two main families of scalable data layers: Bigtable and SimpleDB. Performance evaluation on top of HBase (an open-source version of Bigtable) in our local cluster and Amazon SimpleDB in the Amazon cloud shows that our system scales linearly at least up to 40 nodes in our local cluster and 80 nodes in the Amazon cloud.

A Novel Model for Competition and Cooperation Among Cloud Providers.

Synopsis:

Having received significant attention in the industry, the cloud market is nowadays fiercely competitive with many cloud providers. On one hand, cloud providers compete against each other for both existing and new cloud users. To keep existing users and attract newcomers, it is crucial for each provider to offer an optimal price policy which maximizes the final revenue and improves the competitive advantage. The competition among providers leads to the evolution of the market and dynamic resource prices over time. On the other hand, cloud providers may cooperate with each other to improve their final revenue. Based on a service level agreement, a provider can outsource its users' resource requests to its partner to reduce the operation cost and thereby improve the final revenue. This leads to the problem of determining the cooperating parties in a cooperative environment. This paper tackles these two issues of the current cloud market. First, we solve the problem of competition among providers and propose a dynamic price policy. We employ a discrete choice model to describe the user's choice behavior based on his obtained benefit value. The choice model is used to derive the probability of a user choosing to be served by a certain provider. The competition among providers is formulated as a noncooperative stochastic game where the players are providers who act by proposing the price policy simultaneously. The game is modelled as a Markov Decision Process whose solution is a Markov Perfect Equilibrium. Then, we address the cooperation among providers by presenting a novel algorithm for determining a cooperation strategy that tells providers whether to satisfy users' resource requests locally or outsource them to a certain provider.

The algorithm yields the optimal cooperation structure from which no provider unilaterally deviates to gain more revenue. Numerical simulations are carried out to evaluate the performance of the proposed models.

CloudFTP A Case Study of Migrating Traditional Applications to the Cloud.

Synopsis:

The cloud computing is growing rapidly for it offers on-demand computing power and capacity. The power of cloud enables dynamic scalability of applications facing various business requirements. However, challenges arise when considering the large amount of existing applications. In this work we propose to move the traditional FTP service to the cloud. We implement FTP service on Windows Azure Platform along with the auto-scaling cloud feature. Based on this, we implement a benchmark to measure the performance of our Cloud FTP. This case study illustrates the potential benefits and technical issues associated with the migration of the traditional applications to the clouds.

Cloud Data Production for Masses.

Synopsis:

Offering strong data protection to cloud users while enabling rich applications is a challenging task. Researchers explore a new cloud platform architecture called Data Protection as a Service, which dramatically reduces the per-application development effort required to offer data protection, while still allowing rapid development and maintenance.

Cloud Computing for Agent-Based Urban Transportation Systems.

Synopsis:

Agent-based traffic management systems can use the autonomy, mobility, and adaptability of mobile agents to deal with dynamic traffic environments. Cloud computing can help such systems cope with the large amounts of storage and computing resources required to use traffic strategy agents and mass transport data effectively. This article reviews the history of the development of traffic control and management systems within the evolving computing paradigm and shows the state of traffic control and management systems based on mobile multiagent technology.

A Scalable and Reliable Matching Service for Content-based Publish/Subscribe Systems.

Synopsis:

Characterized by the increasing arrival rate of live content, the emergency applications pose a great challenge: how to disseminate large-scale live content to interested users in a scalable and reliable manner. The publish/subscribe (pub/sub) model is widely used for data dissemination because of its capacity of seamlessly expanding the system to massive size. However, most event matching services of existing pub/sub systems either lead to low matching throughput when matching a large number of skewed subscriptions, or interrupt dissemination when a large number of servers fail. The cloud computing provides great opportunities for the requirements of complex computing and reliable communication. In this paper, we propose SREM, a scalable and reliable event matching service for content-based pub/sub systems in cloud computing environment. To achieve low routing latency and reliable links among servers, we propose a distributed overlay SkipCloud to organize servers of SREM. Through a hybrid space partitioning technique HPartition, large-scale skewed subscriptions are mapped into multiple subspaces, which ensures high matching throughput and provides multiple candidate servers for each event. Moreover, a series of dynamics maintenance mechanisms are extensively studied. To evaluate the performance of SREM, 64 servers are deployed and millions of live content items are tested in a CloudStacktestbed. Under various parameter settings, the experimental results demonstrate that the traffic overhead of routing events in SkipCloud is at least 60 percent smaller than in Chord overlay, the matching rate in SREM is at least 3.7 times and at most 40.4 times larger than the single-dimensional partitioning technique of BlueDove. Besides, SREM enables the event loss rate to drop back to 0 in tens of seconds even if a large number of servers fail simultaneously.

A Log Based Approach to Make Digital Forensics Easier on Cloud Computing.

Synopsis:

Cloud computing is getting more and more attention from the information and communication technologies industry recently. Almost all the leading companies of the information area show their interesting and efforts on cloud computing and release services about cloud computing in succession. But if want to make it go further, we should pay more effort on security issues. Especially, the Internet environment now has become more and more unsecure. With the popularization of computers and intelligent devices, the number of crime on them has increased rapidly in last decades, and will be quicker on the cloud computing environment in future. No wall is wall in the world. We should enhance the cloud computing not only at the aspect of precaution, but also at the aspect of dealing with the security events to defend it from crime activities. In this paper, I propose a approach which using logs model to building a forensic-friendly system. Using this model we can quickly gather information from cloud computing for some kinds of forensic purpose. And this will decrease the complexity of those kinds of forensics.

Cooperative Provable Data Possession for Integrity Verification in Multi-Cloud Storage.

Synopsis:

Provable data possession (PDP) is a technique for ensuring the integrity of data in storage outsourcing. In this paper, we address the construction of an efficient PDP scheme for distributed cloud storage to support the scalability of service and data migration, in which we consider the existence of multiple cloud service providers to cooperatively store and maintain the clients' data. We present a cooperative PDP (CPDP) scheme based on homomorphic verifiable response and hash index hierarchy. We prove the security of our scheme based on multiprover zero-knowledge proof system, which can satisfy completeness, knowledge soundness, and zero-knowledge properties. In addition, we articulate performance optimization mechanisms for our scheme, and in particular present an efficient method for selecting optimal parameter values to minimize the computation costs of clients and storage service providers. Our experiments show that our solution introduces lower computation and communication overheads in comparison with noncooperative approaches.

Going Back and Forth: Efficient Multideployment and Multisnapshotting on Clouds.

Synopsis:

Infrastructure as a Service (IaaS) cloud computing has revolutionized the way we think of acquiring resources by introducing a simple change: allowing users to lease computational resources from the cloud provider’s datacenter for a short time by deploying virtual machines (VMs) on these resources.

This new model raises new challenges in the design and development of IaaS middleware. One of those challenges is the need to deploy a large number (hundreds or even thousands) of VM instances simultaneously. Once the VM instances are deployed, another challenge is to simultaneously take a snapshot of many images and transfer them to persistent storage to support management tasks, such as suspend-resume and migration. With datacenters growing rapidly and configurations becoming heterogeneous, it is important to enable efficient concurrent deployment and snapshotting that are at the same time hypervisor independent and ensure a maximum compatibility with different configurations. This paper addresses these challenges by proposing a virtual file system specifically optimized for virtual machine image storage. It is based on a lazy transfer scheme coupled with object versioning that handles snapshotting transparently in a hypervisor-independent fashion, ensuring high portability for different configurations. Large-scale experiments on hundreds of nodes demonstrate excellent performance results: speedup for concurrent VM deployments ranges from a factor of 2 up to 25, with a reduction in bandwidth utilization of as much as 90%

A Stochastic Model to Investigate Data Center Performance and QoS in IaaS Cloud Computing Systems.

Synopsis:

Cloud data center management is a key problem due to the numerous and heterogeneous strategies that can be applied, ranging from the VM placement to the federation with other clouds. Performance evaluation of cloud computing infrastructures is required to predict and quantify the cost-benefit of a strategy portfolio and the corresponding quality of service (QoS) experienced by users. Such analyses are not feasible by simulation or on-the-field experimentation, due to the great number of parameters that have to be investigated. In this paper, we present an analytical model, based on stochastic reward nets (SRNs), that is both scalable to model systems composed of thousands of resources and flexible to represent different policies and cloud-specific strategies. Several performance metrics are defined and evaluated to analyze the behavior of a cloud data center: utilization, availability, waiting time, and responsiveness. A resiliency analysis is also provided to take into account load bursts. Finally, a general approach is presented that, starting from the concept of system capacity, can help system managers to opportunely set the data center parameters under different working conditions.

Collaboration in Multicloud Computing Environments Framework and Security Issues.

Synopsis:

A proposed proxy-based multicloud computing framework allows dynamic, on-the-fly collaborations and resource sharing among cloud-based services, addressing trust, policy, and privacy issues without preestablished collaboration agreements or standardized interfaces.

Costing of Cloud Computing Services: A Total Cost of Ownership Approach.

Synopsis:

The use of Cloud Computing Services appears to offer significant cost advantages. Particularly start-up companies benefit from these advantages, since frequently they do not operate an internal IT infrastructure. But are costs associated with Cloud Computing Services really that low? We found that particular cost types and factors are frequently underestimated by practitioners. In this paper we present a Total Cost of Ownership (TCO) approach for Cloud Computing Services. We applied a multi-method approach (systematic literature review, analysis of real Cloud Computing Services, expert interview, case study) for the development and evaluation of the formal mathematical model. We found that our model fits the practical requirements and supports decision-making in Cloud Computing.

Improving Utilization of Infrastructure Clouds .Synopsis:

A key advantage of infrastructure-as-a-service (IaaS) clouds is providing users on-demand access to resources. To provide on-demand access, however, cloud providers must either significantly overprovision their infrastructure (and pay a high price for operating resources with low utilization) or reject a large proportion of user requests (in which case the access is no longer on-demand). At the same time, not all users require truly on-demand access to resources. Many applications and workflows are designed for recoverable systems where interruptions in service are expected. For instance, many scientists utilize high-throughput computing (HTC)-enabled resources, such as Condor, where jobs are dispatched to available resources and terminated when the resource is no longer available. We propose a cloud infrastructure that combines on-demand allocation of resources with opportunistic provisioning of cycles from idle cloud nodes to other processes by deploying backfill virtual machines (VMs). For demonstration and experimental evaluation, we extend the Nimbus cloud computing toolkit to deploy backfill VMs on idle cloud nodes for processing an HTC workload. Initial tests show an increase in IaaS cloud utilization from 37.5% to 100% during a portion of the evaluation trace but only 6.39% overhead cost for processing the HTC workload. We demonstrate that a shared infrastructure between IaaS cloud providers and an HTC job management system can be highly beneficial to both the IaaS cloud provider

and HTC users by increasing the utilization of the cloud infrastructure (thereby decreasing the overall cost) and contributing cycles that would otherwise be idle to processing HTC jobs.

Application-Aware Local-Global Source Deduplication for Cloud Backup Services of Personal Storage.

Synopsis:

In personal computing devices that rely on a cloud storage environment for data backup, an imminent challenge facing source deduplication for cloud backup services is the low deduplication efficiency due to a combination of the resource-intensive nature of deduplication and the limited system resources. In this paper, we present ALG-Dedupe, an Application-aware Local-Global source deduplication scheme that improves data deduplication efficiency by exploiting application awareness, and further combines local and global duplicate detection to strike a good balance between cloud storage capacity saving and deduplication time reduction. We perform experiments via prototype implementation to demonstrate that our scheme can significantly improve deduplication efficiency over the state-of-the-art methods with low system overhead, resulting in shortened backup window, increased power efficiency and reduced cost for cloud backup services of personal storage.

Efficient Security Solution for Privacy-Preserving Cloud Services.

Synopsis:

In this paper, we present a novel privacy-preserving security solution for cloud services. We deal with user anonymous access to cloud services and shared storage servers. Our solution provides registered users with anonymous access to cloud services. Our solution offers anonymous authentication. This means that users' personal attributes (age, valid registration, successful payment) can be proven without revealing users' identity. Thus, users can use services without any threat of profiling their behavior. On the other hand, if users break provider's rules, their access rights are revoked. We analyze current privacy preserving solutions for cloud services and outline our solution based on advanced cryptographic components. Our solution offers anonymous access, unlinkability and the confidentiality of transmitted data. Moreover, we implement our solution and we output the experimental results and compare the performance with related solutions.

Fog Computing: Mitigating Insider Data Theft Attacks in the Cloud.

Synopsis:

Cloud computing promises to significantly change the way we use computers and access and store our personal and business information. With these new computing and communications paradigms arise new data security challenges. Existing data protection mechanisms such as encryption have failed in preventing data theft attacks, especially those perpetrated by an insider to the cloud provider. We propose a different approach for securing data in the cloud using offensive decoy technology. We monitor data access in the cloud and detect abnormal data access patterns. When unauthorized access is suspected and then verified using challenge questions, we launch a disinformation attack by returning large amounts of decoy information to the attacker. This protects against the misuse of the user's real data. Experiments conducted in a local file setting provide evidence that this approach may provide unprecedented levels of user data security in a Cloud environment.

Secure and Practical Outsourcing of Linear Programming in Cloud Computing.

Synopsis:

Cloud computing enables customers with limited computational resources to outsource large-scale computational tasks to the cloud, where massive computational power can be easily utilized in a pay-per-use manner. However, security is the major concern that prevents the wide adoption of computation outsourcing in the cloud, especially when end-user's confidential data are processed and produced during the computation. Thus, secure outsourcing mechanisms are in great need to not only protect sensitive information by enabling computations with encrypted data, but also protect customers from malicious behaviors by validating the computation result. Such a mechanism of general secure computation outsourcing was recently shown to be feasible in theory, but to design mechanisms that are practically efficient remains a very challenging problem. Focusing on engineering computing and optimization tasks, this paper investigates secure outsourcing of widely applicable linear programming (LP) computations. In order to achieve practical efficiency, our mechanism design explicitly decomposes the LP computation outsourcing into public LP solvers running on the cloud and private LP parameters owned by the customer. The resulting flexibility allows us to explore appropriate security/efficiency tradeoff via higher-level abstraction of LP computations than the general circuit representation. In particular, by formulating private data owned by the customer for LP problem as a set of matrices and vectors, we are able to develop a set of efficient privacy-preserving problem transformation techniques, which allow customers to transform original LP problem into some random one while protecting sensitive input/output information. To validate the computation result, we further explore the fundamental duality theorem of LP computation and derive the necessary and sufficient conditions that correct result must satisfy. Such result verification mechanism is extremely efficient and incurs close-t- - o-zero

additional cost on both cloud server and customers. Extensive security analysis and experiment results show the immediate practicability of our mechanism design.

CLOUDQUAL: A Quality Model for Cloud Services.Synopsis:

Cloud computing is an important component of the backbone of the Internet of Things (IoT). Clouds will be required to support large numbers of interactions with varying quality requirements. Service quality will therefore be an important differentiator among cloud providers. In order to distinguish themselves from their competitors, cloud providers should offer superior services that meet customers' expectations. A quality model can be used to represent, measure, and compare the quality of the providers, such that a mutual understanding can be established among cloud stakeholders. In this paper, we take a service perspective and initiate a quality model named CLOUDQUAL for cloud services. It is a model with quality dimensions and metrics that targets general cloud services. CLOUDQUAL contains six quality dimensions, i.e., usability, availability, reliability, responsiveness, security, and elasticity, of which usability is subjective, whereas the others are objective. To demonstrate the effectiveness of CLOUDQUAL, we conduct empirical case studies on three storage clouds. Results show that CLOUDQUAL can evaluate their quality. To demonstrate its soundness, we validate CLOUDQUAL with standard criteria and show that it can differentiate service quality.

Fault Tolerance Management in Cloud Computing: A System- Level Perspective.

Synopsis:

The increasing popularity of Cloud computing as an attractive alternative to classic information processing systems has increased the importance of its correct and continuous operation even in the presence of faulty components. In this paper, we introduce an innovative, system-level, modular perspective on creating and managing fault tolerance in Clouds. We propose a comprehensive high-level approach to shading the implementation details of the fault tolerance techniques to application developers and users by means of a dedicated service layer. In particular, the service layer allows the user to specify and apply the desired level of fault tolerance, and does not require knowledge about the fault tolerance techniques that are available in the envisioned Cloud and their implementations.

Managing A Cloud for Multi-agent Systems on Ad-hoc Networks.

Synopsis:

We present a novel execution environment for multi-agent systems building on concepts from cloud computing and peer-to-peer networks. The novel environment can provide the computing power of a cloud for multi-agent systems in intermittently connected networks. We present the design and implementation of a prototype operating system for managing the environment. The operating system provides the user with a consistent view of a single machine, a single file system, and a unified programming model while providing elasticity and availability.

Decentralized Access Control with Anonymous Authentication of Data Stored in Clouds.

Synopsis:

We propose a new decentralized access control scheme for secure data storage in clouds that supports anonymous authentication. In the proposed scheme, the cloud verifies the authenticity of the series without knowing the user's identity before storing data. Our scheme also has the added feature of access control in which only valid users are able to decrypt the stored information. The scheme prevents replay attacks and supports creation, modification, and reading data stored in the cloud. We also address user revocation. Moreover, our authentication and access control scheme is decentralized and robust, unlike other access control schemes designed for clouds which are centralized. The communication, computation, and storage overheads are comparable to centralized approaches.

Framework of Data Integrity for Cross Cloud Environment Using CPDP Scheme.

Synopsis:

Provable data possession (PDP) is a technique for ensuring the integrity of data in storage outsourcing. In this paper, we address the construction of an efficient PDP scheme for distributed cloud storage to support the scalability of service and data migration, in which we consider the existence of multiple cloud service providers to cooperatively store and maintain the clients' data. We present a cooperative PDP (CPDP) scheme based on homomorphic verifiable response and hash index hierarchy. We prove the security of our scheme based on multi-prover zero-knowledge proof system, which can satisfy completeness, knowledge soundness, and zeroknowledge properties. In addition, we articulate

performance optimization mechanisms for our scheme, and in particular present an efficient method for selecting optimal parameter values to minimize the computation costs of clients and storage service providers. Our experiments show that our solution introduces lower computation and communication overheads in comparison with non-cooperative approaches.

Monitoring and Detecting Abnormal Behavior in Mobile Cloud Infrastructure

Synopsis:

Recently, several mobile services are changing to cloud-based mobile services with richer communications and higher flexibility. We present a new mobile cloud infrastructure that combines mobile devices and cloud services. This new infrastructure provides virtual mobile instances through cloud computing. To commercialize new services with this infrastructure, service providers should be aware of security issues. In this paper, we first define new mobile cloud services through mobile cloud infrastructure and discuss possible security threats through the use of several service scenarios. Then, we propose a methodology and architecture for detecting abnormal behavior through the monitoring of both host and network data. To validate our methodology, we injected malicious programs into our mobile cloud test bed and used a machine learning algorithm to detect the abnormal behavior that arose from these programs.

Expressive, Efficient, and Revocable Data Access Control for Multi-Authority Cloud Storage

Synopsis:

Data access control is an effective way to ensure the data security in the cloud. Due to data outsourcing and untrusted cloud servers, the data access control becomes a challenging issue in cloud storage systems. Ciphertext-Policy Attribute-based Encryption (CP-ABE) is regarded as one of the most suitable technologies for data access control in cloud storage, because it gives data owners more direct control on access policies. However, it is difficult to directly apply existing CP-ABE schemes to data access control for cloud storage systems because of the attribute revocation problem. In this paper, we design an expressive,

efficient and revocable data access control scheme for multi-authority cloud storage systems, where there are multiple authorities co-exist and each authority is able to issue attributes independently. Specifically, we propose a revocable multi-authority CP-ABE scheme, and apply it as the underlying techniques to design the data access control scheme. Our attribute revocation method can efficiently achieve both forward security and backward security. The analysis and simulation results show that our proposed data access control scheme is secure in the random oracle model and is more efficient than previous works.

Mining Contracts for Business Events and Temporal Constraints in Service.

Synopsis:

Contracts are legally binding descriptions of business service engagements. In particular, we consider business events as elements of a service engagement. Business events such as purchase, delivery, bill payment, and bank interest accrual not only correspond to essential processes but are also inherently temporally constrained. Identifying and understanding the events and their temporal relationships can help a business partner determine what to deliver and what to expect from others as it participates in the service engagement specified by a contract. However, contracts are expressed in unstructured text and their insights are buried therein. Our contributions are threefold. We develop a novel approach employing a hybrid of surface patterns, parsing, and classification to extract 1) business events and 2) their temporal constraints from contract text. We use topic modeling to 3) automatically organize the event terms into clusters. An evaluation on a real-life contract dataset demonstrates the viability and promise of our hybrid approach, yielding an F-measure of 0.89 in event extraction and 0.90 in temporal constraints extraction. The topic model yields event term clusters with an average match of 85 percent between two independent human annotations and an expert-assigned set of class labels for the clusters.

MORPHOSYS: Efficient Colocation of QoS-Constrained Workloads in the Cloud.

Synopsis:

In hosting environments such as IaaS clouds, desirable application performance is usually guaranteed through the use of Service Level Agreements (SLAs), which specify minimal fractions of resource capacities that must be allocated for unencumbered use for proper operation. Arbitrary colocation of applications with different SLAs on a single host may result in inefficient utilization of the host's resources. In this paper, we propose that periodic resource allocation and consumption models -- often used to characterize real-time

workloads -- be used for a more granular expression of SLAs. Our proposed SLA model has the salient feature that it exposes flexibilities that enable the infrastructure provider to safely transform SLAs from one form to another for the purpose of achieving more efficient colocation. Towards that goal, we present MORPHOSYS: a framework for a service that allows the manipulation of SLAs to enable efficient colocation of arbitrary workloads in a dynamic setting. We present results from extensive trace-driven simulations of colocated Video-on-Demand servers in a cloud setting. These results show that potentially-significant reduction in wasted resources (by as much as 60%) are possible using MORPHOSYS.

Key-Aggregate Cryptosystem for Scalable Data Sharing in Cloud Storage.

Synopsis:

Data sharing is an important functionality in cloud storage. In this paper, we show how to securely, efficiently, and flexibly share data with others in cloud storage. We describe new public-key cryptosystems that produce constant-size ciphertexts such that efficient delegation of decryption rights for any set of ciphertexts are possible. The novelty is that one can aggregate any set of secret keys and make them as compact as a single key, but encompassing the power of all the keys being aggregated. In other words, the secret key holder can release a constant-size aggregate key for flexible choices of ciphertext set in cloud storage, but the other encrypted files outside the set remain confidential. This compact aggregate key can be conveniently sent to others or be stored in a smart card with very limited secure storage. We provide formal security analysis of our schemes in the standard model. We also describe other application of our schemes. In particular, our schemes give the first public-key patient-controlled encryption for flexible hierarchy, which was yet to be known.

Outsourcing Privacy-Preserving Social Networks to a Cloud.

Synopsis:

In the real world, companies would publish social networks to a third party, e.g., a cloud service provider, for marketing reasons. Preserving privacy when publishing social network data becomes an important issue. In this paper, we identify a novel type of privacy attack, termed 1*-neighborhood attack. We assume that an attacker has knowledge about the degrees of a target's one-hop neighbors, in addition to the target's 1-neighborhood graph, which consists of the one-hop neighbors of the target and the relationships among these neighbors. With this information, an attacker may re-identify the target from a k-anonymity

social network with a probability higher than 1/k, where any node's 1-neighborhood graph is isomorphic with k - 1 other nodes' graphs. To resist the 1*-neighborhood attack, we define a key privacy property, probability indistinguishability, for an outsourced social network, and propose a heuristic indistinguishable group anonymization (HIGA) scheme to generate an anonymized social network with this privacy property. The empirical study indicates that the anonymized social networks can still be used to answer aggregate queries with high accuracy.

Reliable Re-Encryption in Unreliable Clouds.

Synopsis:

A key approach to secure cloud computing is for the data owner to store encrypted data in the cloud, and issue decryption keys to authorized users. Then, when a user is revoked, the data owner will issue re-encryption commands to the cloud to re-encrypt the data, to prevent the revoked user from decrypting the data, and to generate new decryption keys to valid users, so that they can continue to access the data. However, since a cloud computing environment is comprised of many cloud servers, such commands may not be received and executed by all of the cloud servers due to unreliable network communications. In this paper, we solve this problem by proposing a time-based re-encryption scheme, which enables the cloud servers to automatically re-encrypt data based on their internal clocks. Our solution is built on top of a new encryption scheme, attribute-based encryption, to allow fine-grain access control, and does not require perfect clock synchronization for correctness.

On the Security of a Public Auditing Mechanism for Shared Cloud Data Service.

Synopsis:

Recently, a public auditing protocol for shared data called Panda (IEEE Transactions on Services Computing, doi: 10.1109/TSC.2013.2295611) was proposed to ensure the correctness of the outsourced data. A distinctive feature of Panda is the support of data sharing and user revocation. Unfortunately, in this letter, we show that Panda is insecure in the sense that a cloud server can hide data loss without being detected. Specifically, we show that even some stored file blocks have been lost, the server is able to generate a valid proof by replacing a pair of lost data block and its signature with another block and signature pair. We also provide a solution to the problem while preserving all the desirable features of the original protocol.

Privacy-Assured Outsourcing of Image Reconstruction Service in Cloud.

Synopsis:

Large-scale image data sets are being exponentially generated today. Along with such data explosion is the fast-growing trend to outsource the image management systems to the cloud for its abundant computing resources and benefits. How to protect the sensitive data while enabling outsourced image services, however, becomes a major concern. To address these challenges, we propose outsourced image recovery service (OIRS), a novel outsourced image recovery service architecture, which exploits different domain technologies and takes security, efficiency, and design complexity into consideration from the very beginning of the service flow. Specifically, we choose to design OIRS under the compressed sensing framework, which is known for its simplicity of unifying the traditional sampling and compression for image acquisition. Data owners only need to outsource compressed image samples to cloud for reduced storage overhead. In addition, in OIRS, data users can harness the cloud to securely reconstruct images without revealing information from either the compressed image samples or the underlying image content. We start with the OIRS design for sparse data, which is the typical application scenario for compressed sensing, and then show its natural extension to the general data for meaningful tradeoffs between efficiency and accuracy. We thoroughly analyze the privacy-protection of OIRS and conduct extensive experiments to demonstrate the system effectiveness and efficiency. For completeness, we also discuss the expected performance speedup of OIRS through hardware built-in system design.

Oruta Privacy-Preserving Public Auditing for Shared Data in the Cloud.

Synopsis:

With cloud data services, it is commonplace for data to be not only stored in the cloud, but also shared across multiple users. Unfortunately, the integrity of cloud data is subject to skepticism due to the existence of hardware/software failures and human errors. Several mechanisms have been designed to allow both data owners and public verifiers to efficiently audit cloud data integrity without retrieving the entire data from the cloud server. However, public auditing on the integrity of shared data with these existing mechanisms will inevitably reveal confidential information-identity privacy-to public verifiers. In this paper, we propose a novel privacy-preserving mechanism that supports public auditing on shared data stored in

the cloud. In particular, we exploit ring signatures to compute verification metadata needed to audit the correctness of shared data. With our mechanism, the identity of the signer on each block in shared data is kept private from public verifiers, who are able to efficiently verify shared data integrity without retrieving the entire file. In addition, our mechanism is able to perform multiple auditing tasks simultaneously instead of verifying them one by one. Our experimental results demonstrate the effectiveness and efficiency of our mechanism when auditing shared data integrity.

Privacy-Preserving Public Auditing for Secure Cloud Storage.

Synopsis:

Using cloud storage, users can remotely store their data and enjoy the on-demand high-quality applications and services from a shared pool of configurable computing resources, without the burden of local data storage and maintenance. However, the fact that users no longer have physical possession of the outsourced data makes the data integrity protection in cloud computing a formidable task, especially for users with constrained computing resources. Moreover, users should be able to just use the cloud storage as if it is local, without worrying about the need to verify its integrity. Thus, enabling public auditability for cloud storage is of critical importance so that users can resort to a third-party auditor (TPA) to check the integrity of outsourced data and be worry free. To securely introduce an effective TPA, the auditing process should bring in no new vulnerabilities toward user data privacy, and introduce no additional online burden to user. In this paper, we propose a secure cloud storage system supporting privacy-preserving public auditing. We further extend our result to enable the TPA to perform audits for multiple users simultaneously and efficiently. Extensive security and performance analysis show the proposed schemes are provably secure and highly efficient. Our preliminary experiment conducted on Amazon EC2 instance further demonstrates the fast performance of the design.

Performance and cost evaluation of an adaptive encryption architecture for cloud databases.

Synopsis:

The cloud database as a service is a novel paradigm that can support several Internet-based applications, but its adoption requires the solution of information confidentiality problems. We propose a novel architecture for adaptive encryption of public cloud

databases that offers an interesting alternative to the tradeoff between the required data confidentiality level and the flexibility of the cloud database structures at design time. We demonstrate the feasibility and performance of the proposed solution through a software prototype. Moreover, we propose an original cost model that is oriented to the evaluation of cloud database services in plain and encrypted instances and that takes into account the variability of cloud prices and tenant workloads during a medium-term period.

Supporting the Migration of Applications to the Cloud through a Decision Support System.

Synopsis:

The motivation for this work is the necessity to be able to select an appropriate Cloud service provider offering for the migration of existing applications, based on cost minimization. While service providers offer pricing information publicly, and online tools allow for the calculation of cost for various Cloud offerings, the selection of which offering fits better the application requirements is left to application developers. For this purpose, this work proposes a migration decision support system that incorporates both offering matching and cost calculation, combining features from various approaches in the State of the Art. The proposed approach is then evaluated against existing tools.

Secure kNN Query Processing in Untrusted Cloud Environments.

Synopsis:

Mobile devices with geo-positioning capabilities (e.g., GPS) enable users to access information that is relevant to their present location. Users are interested in querying about points of interest (POI) in their physical proximity, such as restaurants, cafes, ongoing events, etc. Entities specialized in various areas of interest (e.g., certain niche directions in arts, entertainment, travel) gather large amounts of geo-tagged data that appeal to subscribed users. Such data may be sensitive due to their contents. Furthermore, keeping such information up-to-date and relevant to the users is not an easy task, so the owners of such data sets will make the data accessible only to paying customers. Users send their current location as the query parameter, and wish to receive as result the nearest POIs, i.e., nearest-neighbors (NNs). But typical data owners do not have the technical means to support processing queries on a large scale, so they outsource data storage and querying to a cloud service provider. Many such cloud providers exist who offer powerful storage and computational infrastructures at low cost. However, cloud providers are not fully trusted, and typically behave in an honest-but-curious fashion. Specifically, they follow the protocol to answer queries correctly, but they also collect the locations of the POIs and the subscribers for other purposes. Leakage of POI locations can lead to privacy breaches as well as financial losses to the data owners, for whom the POI data set is an important source of revenue. Disclosure of user locations leads to privacy violations and may deter subscribers from using the service altogether. In this paper, we propose a family of techniques that allow processing of NN queries in an untrusted outsourced environment, while at the same time protecting both the POI and querying users' positions. Our techniques rely on mutable order preserving encoding (mOPE), the only secure order-preserving encryption method known to-date. We also provide performance optimizations to decrease the computational cost inherent to processing on encrypted data, and we consider the case of incrementally updating data sets. We present an extensive performance evaluation of our techniques to illustrate their viability in practice.

Shared Authority Based Privacy-preserving Authentication Protocol in Cloud Computing.

Synopsis:

Cloud computing is an emerging data interactive paradigm to realize users' data remotely stored in an online cloud server. Cloud services provide great conveniences for the users to enjoy the on-demand cloud applications without considering the local infrastructure limitations. During the data accessing, different users may be in a collaborative relationship, and thus data sharing becomes significant to achieve productive benefits. The existing security solutions mainly focus on the authentication to realize that a user's privative data cannot be illegally accessed, but neglect a subtle privacy issue during a user challenging the cloud server to request other users for data sharing. The challenged access request itself may reveal the user's privacy no matter whether or not it can obtain the data access permissions. In this paper, we propose a shared authority based privacy-preserving authentication protocol (SAPA) to address above privacy issue for cloud storage. In the SAPA, 1) shared access authority is achieved by anonymous access request matching mechanism with security and privacy considerations (e.g., authentication, data anonymity, user privacy, and forward security); 2) attribute based access control is adopted to realize that the user can only access its own data fields; 3) proxy re-encryption is applied to provide data sharing among the multiple users. Meanwhile, universal composability (UC) model is established to prove that the SAPA theoretically has the design correctness. It indicates that the proposed protocol is attractive for multi-user collaborative cloud applications.

Winds of Change From Vendor Lock-In to the Meta Cloud.

Synopsis:

The emergence of yet more cloud offerings from a multitude of service providers calls for a meta cloud to smoothen the edges of the jagged cloud landscape. This meta cloud could solve the vendor lock-in problems that current public and hybrid cloud users face.

A Mechanism Design Approach to Resource Procurement in Cloud Computing.

Synopsis:

We present a cloud resource procurement approach which not only automates the selection of an appropriate cloud vendor but also implements dynamic pricing. Three possible

mechanisms are suggested for cloud resource procurement: cloud-dominant strategy incentive compatible (C-DSIC), cloud-Bayesian incentive compatible (C-BIC), and cloud optimal (C-OPT). C-DSIC is dominant strategy incentive compatible, based on the VCG mechanism, and is a low-bid Vickrey auction. C-BIC is Bayesian incentive compatible, which achieves budget balance. C-BIC does not satisfy individual rationality. In C-DSIC and C-BIC, the cloud vendor who charges the lowest cost per unit QoS is declared the winner. In C-OPT, the cloud vendor with the least virtual cost is declared the winner. C-OPT overcomes the limitations of both C-DSIC and C-BIC. C-OPT is not only Bayesian incentive compatible, but also individually rational. Our experiments indicate that the resource procurement cost decreases with increase in number of cloud vendors irrespective of the mechanisms. We also propose a procurement module for a cloud broker which can implement C-DSIC, C-BIC, or C--OPT to perform resource procurement in a cloud computing context. A cloud broker with such a procurement module enables users to automate the choice of a cloud vendor among many with diverse offerings, and is also an essential first step toward implementing dynamic pricing in the cloud.

Cloud based emails boundaries and vulnerabilities .

Synopsis:

Since there is significant increase in adoption of cloud computing, securing users emails is also a growing concern. This paper reviews the boundaries, privacy, vulnerabilities, varying legislations of cloud based emails, and how to mitigate such to provide reliable and secure cloud based email services for users and organizations. We propose a new framework to improve security of cloud based email messages: Intelligent Cloud Based Email Encryption and Decryption System (ICLEEDS). The goal is to encrypt content of email mail messages from users' mail box before being sent. The intelligent machine learning encryption system helps to protect users against email interception, re-construction, phishing attacks, relaying of previous messages, spoofing, eavesdropping and provide high level of privacy.

Automatic Reconfiguration for Large-Scale Reliable Storage Systems.

Synopsis:

Byzantine-fault-tolerant replication enhances the availability and reliability of Internet services that store critical state and preserve it despite attacks or software errors. However, existing Byzantine-fault-tolerant storage systems either assume a static set of replicas, or have limitations in how they handle reconfigurations (e.g., in terms of the scalability of the solutions or the consistency levels they provide). This can be problematic in long-lived, large-scale systems where system membership is likely to change during the system lifetime. In this paper, we present a complete solution for dynamically changing system membership in a large-scale Byzantine-fault-tolerant system. We present a service that tracks system membership and periodically notifies other system nodes of membership changes. The membership service runs mostly automatically, to avoid human configuration errors; is itself Byzantine-fault-tolerant and reconfigurable; and provides applications with a sequence of consistent views of the system membership. We demonstrate the utility of this membership service by using it in a novel distributed hash table called dBQS that provides atomic semantics even across changes in replica sets. dBQS is interesting in its own right because its storage algorithms extend existing Byzantine quorum protocols to handle changes in the replica set, and because it differs from previous DHTs by providing Byzantine fault tolerance and offering strong semantics. We implemented the membership service and dBQS. Our results show that the approach works well, in practice: the membership service is able to manage a large system and the cost to change the system membership is low.

A Medical Image Archive Solution in the Cloud.

Synopsis:

Growing long-term cost of managing an onsite medical imaging archive has been a subject which the health care industry struggles with. Based on the current trend, it is estimated that over 1 billion diagnostic imaging procedures will be performed in the United States during year 2014, generating about 100 Petabytes of data. The high volume of medical images is leading to scalability and maintenance issues with healthcare providers' onsite picture archiving and communication system (PACS) and network. Cloud computing promises lower cost, high scalability, availability and disaster recoverability which can be a natural solution some of the problems we faced for long-term medical image archive. A prototype system was implemented to study such as solution on one of the industry leading cloud computing platform, Microsoft Windows Azure. It includes a Digital Imaging and Communications in Medicine (DICOM) server which handles standard store/query/retrieve requests; a DICOM image indexer that parses the metadata and store them in a SQL Azure database; and a web UI for searching and viewing archived images based on patient and image attributes. The comprehensive tools and functionality of Windows Azure made it an ideal platform to develop and deploy this kind of service oriented applications.

Optimal Multicast capacity and delay tradeoff in manet.

Synopsis:

In this paper, we give a global perspective of multicast capacity and delay analysis in Mobile Ad Hoc Networks (MANETs). Specifically, we consider four node mobility models: (1) two-dimensional i.i.d. mobility, (2) two-dimensional hybrid random walk, (3) one-dimensional i.i.d. mobility, and (4) one-dimensional hybrid random walk. Two mobility time-scales are investigated in this paper: (i) fast mobility where node mobility is at the same time-scale as data transmissions and (ii) slow mobility where node mobility is assumed to occur at a much slower time-scale than data transmissions. Given a delay constraint D, we first characterize the optimal multicast capacity for each of the eight types of mobility models, and then we develop a scheme that can achieve a capacity-delay tradeoff close to the upper bound up to a logarithmic factor. In addition, we also study heterogeneous networks with infrastructure support.

A Scientometric Analysis of Cloud Computing Literature.

Synopsis:

The popularity and rapid development of cloud computing in recent years has led to a huge amount of publications containing the achieved knowledge of this area of research. Due to the interdisciplinary nature and high relevance of cloud computing research, it becomes increasingly difficult or even impossible to understand the overall structure and development of this field without analytical approaches. While evaluating science has a long tradition in many fields, we identify a lack of a comprehensive scientometric study in the area of cloud computing. Based on a large bibliographic data base, this study applies scientometric means to empirically study the evolution and state of cloud computing research with a view from above the clouds. By this, we provide extensive insights into publication patterns, research impact and research productivity. Furthermore, we explore the interplay of related subtopics by analyzing keyword clusters. The results of this study provide a better understanding of patterns, trends and other important factors as a basis for directing research activities, sharing knowledge and collaborating in the area of cloud computing research

Scalable and Secure Sharing of Personal Health Records in Cloud Computing using Attribute-based Encryption.

Synopsis:

Personal health record (PHR) is an emerging patient-centric model of health information exchange, which is often outsourced to be stored at a third party, such as cloud providers. However, there have been wide privacy concerns as personal health information could be exposed to those third party servers and to unauthorized parties. To assure the patients' control over access to their own PHRs, it is a promising method to encrypt the PHRs before outsourcing. Yet, issues such as risks of privacy exposure, scalability in key management, flexible access, and efficient user revocation, have remained the most important challenges toward achieving fine-grained, cryptographically enforced data access control. In this paper, we propose a novel patient-centric framework and a suite of mechanisms for data access control to PHRs stored in semitrusted servers. To achieve fine-grained and scalable data access control for PHRs, we leverage attribute-based encryption (ABE) techniques to encrypt each patient's PHR file. Different from previous works in secure data outsourcing, we focus on the multiple data owner scenario, and divide the users in the PHR system into multiple security domains that greatly reduces the key management complexity for owners and users. A high degree of patient privacy is guaranteed simultaneously by exploiting multiauthority ABE. Our scheme also enables dynamic modification of access policies or file attributes, supports efficient on-demand user/attribute revocation and break-glass access under emergency scenarios. Extensive analytical and experimental results are presented which show the security, scalability, and efficiency of our proposed scheme.

Automatic Protocol Blocker for Privacy-Preserving Public Auditing in Cloud Computing.

Synopsis:

Cloud Computing is nothing but specific style of computing where everything from computing power to infrastructure, business apps are provided “as a service”. In cloud, shared resources, softwares and information is provided as a metered service over the network. By data outsourcing, users can be relieved from the burden of local data storage and maintenance. However, the fact that users no longer have physical possession of the possibly large size of outsourced data makes the data integrity protection in Cloud Computing a very challenging and potentially formidable task, especially for users with constrained computing resources and capabilities. Thus, enabling public auditability for cloud data storage security is of critical importance so that users can resort to an external audit party to check the integrity of outsourced data when needed. In particular, we consider the task of allowing a third party auditor (TPA), on behalf of the cloud client, to verify the integrity of the dynamic data stored in the

cloud. In this paper we are extending the previous system by using automatic blocker for privacy preserving public auditing for data storage security in cloud computing. We utilize the public key based homomorphic authenticator and uniquely integrate it with random mask technique and automatic blocker. In particular, to achieve efficient data dynamics, we improve the existing proof of storage models by manipulating the block tag authentication. Thus, TPA eliminates the involvement of the client through the auditing of whether his Data stored in the Cloud are indeed intact, which can be important in achieving economies of scale For Cloud Computing

A Secured Cost-effective Multi-Cloud Storage in Cloud Computing.

Synopsis:

The end of this decade is marked by a paradigm shift of the industrial information technology towards a pay-per-use service business model known as cloud computing. Cloud data storage redefines the security issues targeted on customer's outsourced data (data that is not stored/retrieved from the costumers own servers). In this work we observed that, from a customer's point of view, relying upon a solo SP for his outsourced data is not very promising. In addition, providing better privacy as well as ensuring data availability, can be achieved by dividing the user's data block into data pieces and distributing them among the available SPs in such a way that no less than a threshold number of SPs can take part in successful retrieval of the whole data block. In this paper, we propose a secured cost-effective multi-cloud storage (SCMCS) model in cloud computing which holds an economical distribution of data among the available SPs in the market, to provide customers with data availability as well as secure storage. Our results show that, our proposed model provides a better decision for customers according to their available budgets.

A Social Compute Cloud: Allocating and Sharing Infrastructure Resources via Social Networks.

Synopsis:

Social network platforms have rapidly changed the way that people communicate and interact. They have enabled the establishment of, and participation in, digital communities as well as the representation, documentation and exploration of social relationships. We

believe that as `apps' become more sophisticated, it will become easier for users to share their own services, resources and data via social networks. To substantiate this, we present a social compute cloud where the provisioning of cloud infrastructure occurs through “friend” relationships. In a social compute cloud, resource owners offer virtualized containers on their personal computer(s) or smart device(s) to their social network. However, as users may have complex preference structures concerning with whom they do or do not wish to share their resources, we investigate, via simulation, how resources can be effectively allocated within a social community offering resources on a best effort basis. In the assessment of social resource allocation, we consider welfare, allocation fairness, and algorithmic runtime. The key findings of this work illustrate how social networks can be leveraged in the construction of cloud computing infrastructures and how resources can be allocated in the presence of user sharing preferences.

Exploiting Rateless Codes in Cloud Storage Systems.

Synopsis:

Block-level cloud storage (BLCS) offers to users and applications the access to persistent block storage devices (virtual disks) that can be directly accessed and used as if they were raw physical disks. In this paper we devise ENIGMA, an architecture for the back-end of BLCS systems able to provide adequate levels of access and transfer performance, availability, integrity, and confidentiality, for the data it stores. ENIGMA exploits LT ratelesscodes to store fragments of sectors on storage nodes organized in clusters. We quantitatively evaluate how the various ENIGMA system parameters affect the performance, availability, integrity, and confidentiality of virtual disks. These evaluations are carried out by using both analytical modeling (for availability, integrity, and confidentiality) and discrete event simulation (for performance), and by considering a set of realistic operational scenarios. Our results indicate that it is possible to simultaneously achieve all the objectives set forth for BLCS systems by using ENIGMA, and that a careful choice of the various system parameters is crucial to achieve a good compromise among them. Moreover, they also show that LT coding-based BLCS systems outperform traditional BLCS systems in all the aspects mentioned before.

Preserving Integrity of Data and Public Auditing for Data Storage Security in Cloud Computing.

Synopsis:

Cloud Computing is the long dreamed vision of computing as a utility, where users can remotely store their data into the cloud so as to enjoy the on-demand high quality applications and services from a shared pool of configurable computing resources. By data

outsourcing, users can be relieved from the burden of local data storage and maintenance. However, the fact that users no longer have physical possession of the possibly large size of outsourced data makes the data integrity protection in Cloud Computing a very challenging and potentially formidable task, especially for users with constrained computing resources and capabilities. Thus, enabling public auditability for cloud data storage security is of critical importance so that users can resort to an external audit party to check the integrity of outsourced data when needed. To securely introduce an effective third party auditor (TPA), the following two fundamental requirements have to be met: 1) TPA should be able to efficiently audit the cloud data storage without demanding the local copy of data, and introduce no additional on-line burden to the cloud user; 2) The third party auditing process should bring in no new vulnerabilities towards user data privacy. In this paper, we utilize and uniquely combine the public key based homomorphic authenticator with random masking to achieve the privacy-preserving public cloud data auditing system, which meets all above requirements. To support efficient handling of multiple auditing tasks, we further explore the technique of bilinear aggregate signature to extend our main result into a multi-user setting, where TPA can perform multiple auditing tasks simultaneously. Extensive security and performance analysis shows the proposed schemes are provably secure and highly efficient.

Data Integrity Proofs in Cloud Storage.

Synopsis:

Cloud computing has been envisioned as the de-facto solution to the rising storage costs of IT Enterprises. With the high costs of data storage devices as well as the rapid rate at which data is being generated it proves costly for enterprises or individual users to frequently update their hardware. Apart from reduction in storage costs data outsourcing to the cloud also helps in reducing the maintenance. Cloud storage moves the user's data to large data centers, which are remotely located, on which user does not have any control. However, this unique feature of the cloud poses many new security challenges which need to be clearly understood and resolved. One of the important concerns that need to be addressed is to assure the customer of the integrity i.e. correctness of his data in the cloud. As the data is physically not accessible to the user the cloud should provide a way for the user to check if the integrity of his data is maintained or is compromised. In this paper we provide a scheme which gives a proof of data integrity in the cloud which the customer can employ to check the correctness of his data in the cloud. This proof can be agreed upon by both the cloud and the customer and can be incorporated in the Service level agreement (SLA). This scheme ensures that the storage at the client side is minimal which will be beneficial for thin clients.

Achieving Effective Cloud Search Services: Multi-keyword Ranked Search over Encrypted Cloud Data Supporting Synonym Query.

Synopsis:

In recent years, consumer-centric cloud computing paradigm has emerged as the development of smart electronic devices combined with the emerging cloud computing technologies. A variety of cloud services are delivered to the consumers with the premise that an effective and efficient cloud search service is achieved. For consumers, they want to find the most relevant products or data, which is highly desirable in the "pay-as-you use" cloud computing paradigm. As sensitive data (such as photo albums, emails, personal health records, financial records, etc.) are encrypted before outsourcing to cloud, traditional keyword search techniques are useless. Meanwhile, existing search approaches over encrypted cloud data support only exact or fuzzy keyword search, but not semantics-based multi-keyword ranked search. Therefore, how to enable an effective searchable system with support of ranked search remains a very challenging problem. This paper proposes an effective approach to solve the problem of multi-keyword ranked search over encrypted cloud data supporting synonym queries. The main contribution of this paper is summarized in two aspects: multi-keyword ranked search to achieve more accurate search results and synonym-based search to support synonym queries. Extensive experiments on real-world dataset were performed to validate the approach, showing that the proposed solution is very effective and efficient for multikeyword ranked searching in a cloud environment.

Fuzzy Authorization for Cloud Storage.

Synopsis:

By leveraging and modifying ciphertext-policy attribute based encryption (CP-ABE) and OAuth, we propose a new authorization scheme, called fuzzy authorization, to facilitate an application registered with one cloud party to access data residing in another cloud party. The new proposed scheme enables the fuzziness of authorization to enhance the scalability and flexibility of file sharing by taking advantage of the one-to-one correspondence between linear secret-sharing scheme (LSSS) and generalized Reed Solomon (GRS) code. Furthermore, by conducting attribute distance checking and distance adjustment, operations like sending attribute sets and satisfying an access tree are eliminated. In addition, the automatic revocation is realized with update of TimeSlot attribute when data owner modifies the data. The security of the fuzzy authorization is proved under the d-BDHE assumption. In order to measure and estimate the performance of our scheme, we have implemented the protocol flow of fuzzy authorization with OMNET++ 4.2.2 and realized the cryptographic part

with pairing-based cryptography (PBC) library. Experimental results show that fuzzy authorization can achieve fuzziness of authorization among heterogeneous clouds with security and efficiency.

Privacy Preserving Delegated Access Control in Public Clouds .

Synopsis:

Current approaches to enforce fine-grained access control on confidential data hosted in the cloud are based on fine-grained encryption of the data. Under such approaches, data owners are in charge of encrypting the data before uploading them on the cloud and re-encrypting the data whenever user credentials change. Data owners thus incur high communication and computation costs. A better approach should delegate the enforcement of fine-grained access control to the cloud, so to minimize the overhead at the data owners, while assuring data confidentiality from the cloud. We propose an approach, based on two layers of encryption, that addresses such requirement. Under our approach, the data owner performs a coarse-grained encryption, whereas the cloud performs a fine-grained encryption on top of the owner encrypted data. A challenging issue is how to decompose access control policies (ACPs) such that the two layer encryption can be performed. We show that this problem is NP-complete and propose novel optimization algorithms. We utilize an efficient group key management scheme that supports expressive ACPs. Our system assures the confidentiality of the data and preserves the privacy of users from the cloud while delegating most of the access control enforcement to the cloud.

Dynamic Audit Services for Integrity Verification of Outsourced Storages in Clouds.

Synopsis:

In this paper, we propose a dynamic audit service for verifying the integrity of an untrusted and outsourced storage. Our audit service is constructed based on the techniques, fragment structure, random sampling, and index-hash table, supporting provable updates to outsourced data and timely anomaly detection. In addition, we propose a method based on probabilistic query and periodic verification for improving the performance of audit services. Our experimental results not only validate the effectiveness of our approaches, but also show our audit system verifies the integrity with lower computation overhead and requiring less extra storage for audit metadata.

An Efficient Certificateless Encryption for Secure Data Sharing in Public Clouds.

Synopsis:

We propose a mediated certificateless encryption scheme without pairing operations for securely sharing sensitive information in public clouds. Mediated certificateless public key encryption (mCL-PKE) solves the key escrow problem in identity based encryption and certificate revocation problem in public key cryptography. However, existing mCL-PKE schemes are either inefficient because of the use of expensive pairing operations or vulnerable against partial decryption attacks. In order to address the performance and security issues, in this paper, we first propose a mCL-PKE scheme without using pairing operations. We apply our mCL-PKE scheme to construct a practical solution to the problem of sharing sensitive information in public clouds. The cloud is employed as a secure storage as well as a key generation center. In our system, the data owner encrypts the sensitive data using the cloud generated users' public keys based on its access control policies and uploads the encrypted data to the cloud. Upon successful authorization, the cloud partially decrypts the encrypted data for the users. The users subsequently fully decrypt the partially decrypted data using their private keys. The confidentiality of the content and the keys is preserved with respect to the cloud, because the cloud cannot fully decrypt the information. We also propose an extension to the above approach to improve the efficiency of encryption at the data owner. We implement our mCL-PKE scheme and the overall cloud based system, and evaluate its security and performance. Our results show that our schemes are efficient and practical.

Identity-Based Distributed Provable Data Possession in Multi- Cloud Storage.

Synopsis:

Remote data integrity checking is of crucial importance in cloud storage. It can make the clients verify whether their outsourced data is kept intact without downloading the whole data. In some application scenarios, the clients have to store their data on multicloud servers. At the same time, the integrity checking protocol must be efficient in order to save the verifier's cost. From the two points, we propose a novel remote data integrity checking model: ID-DPDP (identity-based distributed provable data possession) in multicloud storage. The formal system model and security model are given. Based on the bilinear pairings, a concrete ID-DPDP protocol is designed. The proposed ID-DPDP protocol is provably secure under the hardness assumption of the standard CDH (computational Diffie-Hellman) problem. In addition to the structural advantage of elimination of certificate

management, our ID-DPDP protocol is also efficient and flexible. Based on the client's authorization, the proposed ID-DPDP protocol can realize private verification, delegated verification, and public verification.

Remote Display Solutions for Mobile Cloud Computing.

Synopsis:

Proposed optimization techniques address the major challenges that varying wireless channel conditions, short battery lifetime, and interaction latency pose for the remote display of cloud applications on mobile devices.

Automatic Scaling of Internet Applications for Cloud Computing Services.

Synopsis:

Many Internet applications can benefit from an automatic scaling property where their resource usage can be scaled up and down automatically by the cloud service provider. We present a system that provides automatic scaling for Internet applications in the cloud environment. We encapsulate each application instance inside a virtual machine (VM) and use virtualization technology to provide fault isolation. We model it as the Class Constrained Bin Packing (CCBP) problem where each server is a bin and each class represents an application. The class constraint reflects the practical limit on the number of applications a server can run simultaneously. We develop an efficient semi-online color set algorithm that achieves good demand satisfaction ratio and saves energy by reducing the number of servers used when the load is low. Experiment results demonstrate that our system can improve the throughput by 180% over an open source implementation of Amazon EC2 and restore the normal QoS five times as fast during flash crowds. Large scale simulations demonstrate that our algorithm is extremely scalable: the decision time remains under 4 s for a system with 10 000 servers and 10 000 applications. This is an order of magnitude improvement over traditional application placement algorithms in enterprise environments.

Panda: Public Auditing for Shared Data with Efficient User Revocation in the Cloud.

Synopsis:

With data storage and sharing services in the cloud, users can easily modify and share data as a group. To ensure shared data integrity can be verified publicly, users in the group need to compute signatures on all the blocks in shared data. Different blocks in shared data are generally signed by different users due to data modifications performed by different users. For security reasons, once a user is revoked from the group, the blocks which were previously signed by this revoked user must be re-signed by an existing user. The straightforward method, which allows an existing user to download the corresponding part of shared data and re-sign it during user revocation, is inefficient due to the large size of shared data in the cloud. In this paper, we propose a novel public auditing mechanism for the integrity of shared data with efficient user revocation in mind. By utilizing the idea of proxy re-signatures, we allow the cloud to re-sign blocks on behalf of existing users during user revocation, so that existing users do not need to download and re-sign blocks by themselves. In addition, a public verifier is always able to audit the integrity of shared data without retrieving the entire data from the cloud, even if some part of shared data has been re-signed by the cloud. Moreover, our mechanism is able to support batch auditing by verifying multiple auditing tasks simultaneously. Experimental results show that our mechanism can significantly improve the efficiency of user revocation.

Scalable and Secure Sharing of Personal Health Records in Cloud Computing using Attribute-based Encryption.

Synopsis:

Personal health record (PHR) is an emerging patient-centric model of health information exchange, which is often outsourced to be stored at a third party, such as cloud providers. However, there have been wide privacy concerns as personal health information could be exposed to those third party servers and to unauthorized parties. To assure the patients' control over access to their own PHRs, it is a promising method to encrypt the PHRs before outsourcing. Yet, issues such as risks of privacy exposure, scalability in key management, flexible access, and efficient user revocation, have remained the most important challenges toward achieving fine-grained, cryptographically enforced data access control. In this paper, we propose a novel patient-centric framework and a suite of mechanisms for data access control to PHRs stored in semitrusted servers. To achieve fine-grained and scalable data access control for PHRs, we leverage attribute-based encryption (ABE) techniques to encrypt each patient's PHR file. Different from previous works in secure data outsourcing,

we focus on the multiple data owner scenario, and divide the users in the PHR system into multiple security domains that greatly reduces the key management complexity for owners and users. A high degree of patient privacy is guaranteed simultaneously by exploiting multiauthority ABE. Our scheme also enables dynamic modification of access policies or file attributes, supports efficient on-demand user/attribute revocation and break-glass access under emergency scenarios. Extensive analytical and experimental results are presented which show the security, scalability, and efficiency of our proposed scheme.

Planning vs dynamic control: Resource allocation in corporate clouds.

Synopsis:

Nowadays corporate data centers leverage virtualization technology to cut operational and management costs. Virtualization allows splitting and assigning physical servers to virtual machines (VM) that run particular business applications. This has led to a new stream in the capacity planning literature dealing with the problem of assigning VMs with volatile demands to physical servers in a static way such that energy costs are minimized. Live migration technology allows for dynamic resource allocation, where a controller responds to overload or underload on a server during runtime and reallocates VMs in order to maximize energy efficiency. Dynamic resource allocation is often seen as the most efficient means to allocate hardware resources in a data center. Unfortunately, there is hardly any experimental evidence for this claim. In this paper, we provide the results of an extensive experimental analysis of both capacity management approaches on a data center infrastructure. We show that with typical workloads of transactional business applications dynamic resource allocation does not increase energy efficiency over the static allocation of VMs to servers and can even come at a cost, because migrations lead to overheads and service disruptions.

Compatibility-Aware Cloud Service Composition under Fuzzy Preferences of Users.

Synopsis:

When a single Cloud service (i.e., a software image and a virtual machine), on its own, cannot satisfy all the user requirements, a composition of Cloud services is required. Cloud service composition, which includes several tasks such as discovery, compatibility checking, selection, and deployment, is a complex process and users find it difficult to select the best one among the hundreds, if not thousands, of possible compositions available. Service composition in Cloud raises even new challenges caused by diversity of users with different expertise requiring their applications to be deployed across difference

geographical locations with distinct legal constraints. The main difficulty lies in selecting a combination of virtual appliances (software images) and infrastructure services that are compatible and satisfy a user with vague preferences. Therefore, we present a framework and algorithms which simplify Cloud service composition for unskilled users. We develop an ontology-based approach to analyze Cloud service compatibility by applying reasoning on the expert knowledge. In addition, to minimize effort of users in expressing their preferences, we apply combination of evolutionary algorithms and fuzzy logic for composition optimization. This lets users express their needs in linguistics terms which brings a great comfort to them compared to systems that force users to assign exact weights for all preferences.

Towards Secure and Dependable Storage Services in Cloud Computing.

Synopsis:

Cloud storage enables users to remotely store their data and enjoy the on-demand high quality cloud applications without the burden of local hardware and software management. Though the benefits are clear, such a service is also relinquishing users' physical possession of their outsourced data, which inevitably poses new security risks toward the correctness of the data in cloud. In order to address this new problem and further achieve a secure and dependable cloud storage service, we propose in this paper a flexible distributed storage integrity auditing mechanism, utilizing the homomorphic token and distributed erasure-coded data. The proposed design allows users to audit the cloud storage with very lightweight communication and computation cost. The auditing result not only ensures strong cloud storage correctness guarantee, but also simultaneously achieves fast data error localization, i.e., the identification of misbehaving server. Considering the cloud data are dynamic in nature, the proposed design further supports secure and efficient dynamic operations on outsourced data, including block modification, deletion, and append. Analysis shows the proposed scheme is highly efficient and resilient against Byzantine failure, malicious data modification attack, and even server colluding attacks.

Privacy-Preserving Multi-keyword Ranked Search over Encrypted Cloud Data.

Synopsis:

With the advent of cloud computing, data owners are motivated to outsource their complex data management systems from local sites to the commercial public cloud for great flexibility and economic savings. But for protecting data privacy, sensitive data have to be encrypted before outsourcing, which obsoletes traditional data utilization based on plaintext

keyword search. Thus, enabling an encrypted cloud data search service is of paramount importance. Considering the large number of data users and documents in the cloud, it is necessary to allow multiple keywords in the search request and return documents in the order of their relevance to these keywords. Related works on searchable encryption focus on single keyword search or Boolean keyword search, and rarely sort the search results. In this paper, for the first time, we define and solve the challenging problem of privacy-preserving multi-keyword ranked search over encrypted data in cloud computing (MRSE). We establish a set of strict privacy requirements for such a secure cloud data utilization system. Among various multi-keyword semantics, we choose the efficient similarity measure of "coordinate matching," i.e., as many matches as possible, to capture the relevance of data documents to the search query. We further use "inner product similarity" to quantitatively evaluate such similarity measure. We first propose a basic idea for the MRSE based on secure inner product computation, and then give two significantly improved MRSE schemes to achieve various stringent privacy requirements in two different threat models. To improve search experience of the data search service, we further extend these two schemes to support more search semantics. Thorough analysis investigating privacy and efficiency guarantees of proposed schemes is given. Experiments on the real-world data set further show proposed schemes indeed introduce low overhead on computation and communication.

Consistency as a Service: Auditing Cloud Consistency.

Synopsis:

Cloud storage services have become commercially popular due to their overwhelming advantages. To provide ubiquitous always-on access, a cloud service provider (CSP) maintains multiple replicas for each piece of data on geographically distributed servers. A key problem of using the replication technique in clouds is that it is very expensive to achieve strong consistency on a worldwide scale. In this paper, we first present a novel consistency as a service (CaaS) model, which consists of a large data cloud and multiple small audit clouds. In the CaaS model, a data cloud is maintained by a CSP, and a group of users that constitute an audit cloud can verify whether the data cloud provides the promised level of consistency or not. We propose a two-level auditing architecture, which only requires a loosely synchronized clock in the audit cloud. Then, we design algorithms to quantify the severity of violations with two metrics: the commonality of violations, and the staleness of the value of a read. Finally, we devise a heuristic auditing strategy (HAS) to reveal as many violations as possible. Extensive experiments were performed using a combination of simulations and real cloud deployments to validate HAS.

Secure kNN Query Processing in Untrusted Cloud Environments

Synopsis:

Mobile devices with geo-positioning capabilities (e.g., GPS) enable users to access information that is relevant to their present location. Users are interested in querying about points of interest (POI) in their physical proximity, such as restaurants, cafes, ongoing events, etc. Entities specialized in various areas of interest (e.g., certain niche directions in arts, entertainment, travel) gather large amounts of geo-tagged data that appeal to subscribed users. Such data may be sensitive due to their contents. Furthermore, keeping such information up-to-date and relevant to the users is not an easy task, so the owners of such data sets will make the data accessible only to paying customers. Users send their current location as the query parameter, and wish to receive as result the nearest POIs, i.e., nearest-neighbors (NNs). But typical data owners do not have the technical means to support processing queries on a large scale, so they outsource data storage and querying to a cloud service provider. Many such cloud providers exist who offer powerful storage and computational infrastructures at low cost. However, cloud providers are not fully trusted, and typically behave in an honest-but-curious fashion. Specifically, they follow the protocol to answer queries correctly, but they also collect the locations of the POIs and the subscribers for other purposes. Leakage of POI locations can lead to privacy breaches as well as financial losses to the data owners, for whom the POI data set is an important source of revenue. Disclosure of user locations leads to privacy violations and may deter subscribers from using the service altogether. In this paper, we propose a family of techniques that allow processing of NN queries in an untrusted outsourced environment, while at the same time protecting both the POI and querying users' positions. Our techniques rely on mutable order preserving encoding (mOPE), the only secure order-preserving encryption method known to-date. W- also provide performance optimizations to decrease the computational cost inherent to processing on encrypted data, and we consider the case of incrementally

updating data sets. We present an extensive performance evaluation of our techniques to illustrate their viability in practice.

Securing the cloud storage audit service: defending against frame and collude attacks of third party auditor.

Synopsis:

Cloud computing has been envisioned as the next generation architecture of the IT enterprise, but there exist many security problems. A significant problem encountered in the context of cloud storage is whether there exists some potential vulnerabilities towards cloud storage system after introducing third parties. Public verification enables a third party auditor (TPA), on behalf of users who lack the resources and expertise, to verify the integrity of the stored data. Many existing auditing schemes always assume TPA is reliable and independent. This work studies the problem what if certain TPAs are semi-trusted or even potentially malicious in some situations. Actually, the authors consider the task of allowing such a TPA to involve in the audit scheme. They propose a feedback-based audit scheme via which users are relaxed from interacting with cloud service provider (CSP) and can check the integrity of stored data by themselves instead of TPA yet. Specifically, TPA generates the feedback through processing the proof from CSP and returns it to user which is yet unforgeable to TPA and checked exclusively by user. Through detailed security and performance analysis, the author's scheme is shown to be more secure and lightweight.

Strategy-proof Pricing for Cloud Service Composition

Synopsis:

The on-demand provisions of cloud services create a service market, where users can dynamically select services based on such attractive criteria as price and quality. An intuitive model of a service market is a reverse auction. In the first price auction, however, a service that is cheaper and provides better quality is not necessarily selected. This causes undesirable outcomes both for users and service providers. A possible solution is the Vickrey-Clarke-Groves (VCG) mechanism, where the dominant strategy for a service provider is to report the true cost of his service. In spite of this desirable property, implementing the VCG mechanism for service composition suffers from computational cost. The calculation of payments to service providers based on the VCG mechanism requires iterative service selection, even though each service selection can be NP-hard. Approximation algorithms cannot be applied because approximate solutions do not assure the desirable property of the VCG mechanism. Thus, we model VCG payments for service markets and propose a dynamic programming (DP)-based algorithm for service selection and VCG payment calculation. Our proposed algorithm solves service selection in quasi-

polynomial time and gives an exact solution. Moreover, we extend it and focus on the iterative service selection process for VCG payment calculation to improve its performance. Our series of experiments show that our proposed algorithm solves practical scale service composition.