ACTEMT 2010: IEEE Globecom 2010 Workshop on Application of Communication Theory to Emerging Memory Technologies

Technical Program

Monday, December 6

9:00 AM - 10:30 AM

ACTEMT 01: Plenary Talk and Panel Discussion

Plenary Talk by Professor Jack Wolf, CMRR Endowed Chair, UC San Diego. Panel Discussion on Future of Solid State Devices

Room: Orchid C

Chairs: Rajiv Agarwal (Stanford University, USA), Marcus Marrow (Link_A_Media Devices Corporation, USA)

11:00 AM - 12:30 PM

ACTEMT 02: Performance Modelling

Room: Orchid AB

Chair: Jason Bellorado (Link-A-Media, USA)

A Quantitative Framework for Modeling and Analyzing Flash Memory Wear Leveling Algorithms

Mochan Shrestha (Wayne State University, USA); Lihao Xu (Wayne State University, USA)

Flash memory is emerging as an enabling technology that is rapidly changing the landscape of storage systems via its many desirable properties such as low power consumption and high random I/O throughput. However, due to its physical characteristics, flash cells have limited erase (program) cycles. To increase endurance and reliability, flash memory based devices employ certain wear leveling algorithms. But so far there hasn't been a rigorous mathematical tool to analyze and evaluate the effectiveness of various wear leveling algorithms. In this paper, we present a mathematical framework to model whole block wear leveling in flash-based solid state drives (SSD) based on probabilistic models of workloads, strategies and wear level states. From this we derive equations for the distribution of wear levels for SSD configurations thus enabling quantitative analysis and evaluation of wear leveling algorithms for flash memory based devices.

pp. 1898-1902

Adaptive Endurance Coding for NAND Flash

Ashish Jagmohan (IBM T. J. Watson Research Center, USA); Michele M Franceschini (IBM T.J. Watson Research Center, USA); Luis A Lastras-Montaño (IBM TJ Watson Research Center, USA); John Karidis (IBM TJ Watson Research Center, USA)

A fundamental constraint in the use of newer NAND Flash devices in the enterprise space is the low cycling endurance of such devices. As an example, the latest 2-bit MLC devices have a cycling endurance ranging from 3K to 10K program/erase cycles. Upcoming higher-density devices are expected to have even lower endurance. In this paper we propose a coding technique called Adaptive Endurance Coding (AEC) which increases the number of program/erase cycles that a Flash device can endure. The key insight leveraged by the proposed technique is the data-dependent nature of Flash cell-wear. Data-dependent wear implies that Flash chip/device lifetime can be significantly increased by converting data into bit-patterns, prior to programming, which cause minimal wear. AEC is a constrained coding technique which seeks to achieve this. Specifically, the technique can be used to generate a capacity-wear tradeoff; for compressible data, AEC can be adapted to data compressibility in order to maximize endurance gains with low system overhead costs. The technique can be implemented in the Flash device controller without requiring any hardware changes to the device itself. We present empirical results on SLC and MLC Flash chips demonstrating the improvements in retention and bit-error rate which can be obtained via this technique. We also present disk-level simulation results comparing the performance of AEC to other approaches such as lossless compression alone.

pp. 1903-1907

A closed-form expression for Write Amplification in NAND Flash

Rajiv Agarwal (Stanford University, USA); Marcus Marrow (Link_A_Media Devices Corporation, USA)

The log-structured filesystems typically used in current solid-state drive's (SSD) exhibit write amplification, whereby multiple NAND writes are required for each host write. Write amplification negatively affects the SSD endurance and write throughput. This performance loss depends on the drive over-provisioning and the garbage collection method. This paper presents a novel probabilistic model to analytically quantify the impact of over-provisioning on write amplification under a uniformly-distributed random workload and a greedy garbage collection policy. The analysis shows write amplification approximately independent of NAND block size and number of blocks in the SSD. The analysis is verified by full drive simulations.

pp. 1908-1912

An area and latency assessment for coding for memories with stuck cells

Luis A Lastras-Montaño (IBM TJ Watson Research Center, USA); Ashish Jagmohan (IBM T. J. Watson Research Center, USA); Michele M Franceschini (IBM T.J. Watson Research Center, USA)

We explore the implementation in hardware of encoders for algebraic codes for binary memories that have some cells stuck to given values. The location and stuck values of the cells are assumed to be known at encode time, but not at decode time. This type of code can be particularly useful for memories whose primary failure mechanism can be described as cells becoming stuck at a value. Recently, an algorithm for BCH-like codes for stuck cells was developed that relies on a type of polynomial interpolation in an extension field followed by a step of projection of the result of the interpolation back into the binary field. In this article we report on the area and latency of a sequential implementation of this encoding algorithm for a variety of values for the number of stuck cells. The statistics on area and latency follow from preliminary circuit synthesis results obtained from a full VHDL implementation of these examples.

pp. 1913-1917

ACTEMT 03: Coding for Memories I

Room: Orchid C

Chair: Henry D Pfister (Texas A&M University, USA)

Error Characterization and Coding Schemes for Flash Memories

Eitan Yaakobi (University of California, San Diego, USA); Paul H. Siegel (University of California, San Diego, USA); Steven Swanson (University of California, San Diego, USA); Jack Wolf (UCSD, USA); Laura M Grupp (University of California, San Diego, USA); Jing Ma (National University of Singapore, Singapore)

In this work, we use an extensive empirical database of errors induced by write, read, and erase operations to develop a comprehensive understanding of the error behavior of flash memories. Error characterization of MLC and SLC flash is given on the block, page, and bit level. Based on our error characterization in MLC flash, we propose an error-correcting scheme which outperforms the conventional BCH code. We compare several schemes which use an MLC flash block as an SLC block. Finally, an implementation of two-write WOM-codes in SLC flash is given as well as the BER for the first and second write.

pp. 1918-1922

Rewriting Codes for Flash Memories Based Upon Lattices, and an Example Using the E8 Lattice

Brian Michael Kurkoski (University of Electro-Communications, Japan)

A rewriting code construction for flash memories based upon lattices is described, where the values stored in flash cells correspond to lattice points. This construction encodes information to lattice points in such a way that data can be written to the memory multiple times without decreasing the cell values. The construction partitions the flash memory's cubic signal space into blocks, which aids with encoding. The minimum number of writes is approximately linear in one of the code parameters. Using the E8 lattice as an example, the average number of writes can be increased by introducing randomization in the encoding.

pp. 1923-1927

Rank Modulation with Multiplicity

Anxiao Andrew Jiang (Texas A&M University, USA); Yue Wang (Texas A&M University, USA)

Rank modulation is a scheme that uses the relative order of cell levels to represent data. Its applications include flash memories, phase-change memories, etc. An extension of rank modulation is studied in this paper, where multiple cells can have the same rank. We focus on the rewriting of data based on this new scheme, and study its basic properties.

pp. 1928-1932

Not Just for Errors: Codes for Fast and Secure Flash Storage

Yuval Cassuto (EPFL, Switzerland)

Error-correcting codes are normally employed in storage devices to guarantee the integrity of data in the presence of errors. This paper presents two schemes where error-correcting codes are used for entirely different purposes. In the first part of the paper, a new coding paradigm is proposed to improve the write performance of multi-level flash devices. By slightly relaxing the accuracy of cell programming, significant speed up can be achieved. The resulting write inaccuracies are then corrected by codes that are tailored for the appropriately restricted error model. In the second part, new low-complexity codes are proposed to protect the security of sensitive data in the presence of imperfect physical erasure process. Codes that have optimal encoding and decoding complexities are constructed to allow fast storing and retrieval of secret data, and guarantee unconditional security of data against an adversary with access to parts of the secret that

failed to erase.
pp. 1933-1937

2:00 PM - 3:30 PM

ACTEMT 04: Modelling for Storage Media

Room: Orchid AB
Chair: Ashish Jagmohan (IBM T. J. Watson Research Center, USA)

Challenges and opportunities for information theory-based design of Phase Change Memories

Andrea Marinoni (University of Pavia, Italy); Alessandro Cabrini (University of Pavia, Italy); Eugenio Costamagna (University of Pavia, Italy); Guido Torelli (University of Pavia, Italy); Paolo Gamba (Università degli Studi di Pavia, Italy)

Phase Change Memory (PCM) definitely represents one of the most promising technologies among the non-volatile memories to be used in the next decade. Even though the noise resilience of PCMs looks to be very strong, an information theory-based design may improve the error-rate performance of the PCM reading process. In this paper, we introduce a channel model for the information recovery in PCMs and discuss the related mutual information. Experimental results are provided in order to highlight relationships between information recovery performance, electrical parameters of the memory cell, and the sensitivity of the reading architecture. This information theory-based point of view represents a base for an effective optimization of the error-rate performance during the PCM reading process and opens the path to very interesting new research lines.
pp. 1938-1942

Exploration on Sub-nanosecond Spin Torque Random Access Memory

Xiaobin Wang (Seagate Technology, USA)

a key attribute that spin torque random access memory (SPRAM) offers as advanced nonvolatile memory is the fast read and write performance. In this paper, based upon model and experiment data, we study SPRAM switching speed under switching variation constrains. We analyze factors limiting SPRAM switching speed and examine possibilities that lead SPRAM speed down to sub-nanosecond region. We explore paths that are different from direct scaling down approaches. Particularly we quantify magnetization relaxation and temperature effects on SPRAM switching speed. We suggest that linear magnetization reversal mechanism at elevated temperature provides a path that leads to sub-nanosecond SPRAM with well controlled switching time variation.
pp. 1943-1947

Understanding of switching phenomena in unipolar NiO-based RRAM

Hyung Dong Lee (Stanford University, USA); Yoshio Nishi (Stanford University, USA)

The electronic band structure of cation and anion vacancies in NiO is assessed using density functional theory in conjunction with the local density approximation and employing on site Coulomb corrections within the LDA+U method. The calculated formation energies identify the stability of charged vacancy states consistent with experimental reports. We propose a microscopic model for the formation and rupture of an electrically active filament in NiO targeted to explain the unipolar switching phenomenon observed in resistive change memory devices. The filament formation and rupture process are linked to the migration of oxygen in the oxide coupled to the oxidation/reduction of nickel atoms. The calculated migration barrier consistent with experimental activation barrier for retention can be used for making a retention model and evaluating the retention time.
pp. 1948-1951

The inner workings of phase change memory: lessons from prototype PCM devices

Geoffrey W Burr (IBM Almaden Research Center, USA); Alvaro Padilla (IBM Almaden Research Center, USA); Michele M Franceschini (IBM T.J. Watson Research Center, USA); Bryan Jackson (IBM Almaden Research Center, USA); Diego Dupouy (IBM Almaden Research Center, USA); Charles T. Rettner (IBM Almaden Research Center, USA); Kailash Gopalakrishnan (IBM Almaden Research Center, USA); Rohit Shenoy (IBM Almaden Research Center, USA); John Karidis (IBM TJ Watson Research Center, USA)

We describe some observations into the inner workings of phase-change memory devices, obtained by fabrication, electrical characterization, failure analysis and modeling of prototype phase change memory devices over the past few years. Experiments involving the RESET and SET operations, the speed of SET operations, the impact of voltage polarity, and the drift of RESET resistances after programming have been performed. Our prototype devices include PCM "pore" devices down to 20nm in diameter, PCM "bridge" devices down to 20nm in width, and novel "parallel cell" devices designed to trade off resistance contrast for lower resistance drift. Devices have been fabricated with the standard GST:225 as well as several other material variants. Simple resistors integrated in series with these devices allow programming with pulses down to <10ns pulse-widths. Highly accurate current measurements can be triggered as soon as 1 millisecond after programming, so that even five-minute-long measurements can capture drift over five orders of magnitude in time. Our customized finite-difference PCM simulator is capable of handling large and arbitrary 3-D structures, and can be matched against fast electrical SET and RESET experiments, slow thin-film crystallization experiments, or optical pulse experiments. Our results suggest that many of the idiosyncrasies of real PCM devices, such as "telegraph"-like noise, the dependence of SET resistance on pulse duration, and the variability of both repeated programming events and drift measurements even on the same device, can be traced to the important yet unappreciated role of poly-crystalline grains and grain boundaries within the PCM device.
pp. 1952-1956

ACTEMT 05: Coding for Memories II

Room: Orchid C
Chair: Luis A Lastras-Montaño (IBM TJ Watson Research Center & IBM Corporation, USA)

A Multibit-Per-Cell Memory Model and Nonbinary LDPC Codes

Seungjune Jeon (Carnegie Mellon University, USA); Euiseok Hwang (Carnegie Mellon University, USA); B. V. K. Vijaya Kumar (Carnegie Mellon University, USA); Michael K. Cheng (Jet Propulsion Laboratory, USA)

Protecting nonvolatile memory systems in harsh radiation environments encountered in space missions is important and error correcting schemes can extend the lifetime of those memory systems. For example, recent research has shown that LDPC codes can extend the lifetime of nonvolatile memory under space radiation environment more than Bose-Chaudhuri-Hocquenghem (BCH) or Reed-Solomon (RS) codes at fixed codeword error rates. However, conventional memory models assume that bit errors are independent, but multibit errors were reported in satellite experiments. Moreover, memory feature sizes are shrinking and multibit-per-cell structures are becoming standard so radiation will increasingly lead to multibit errors. For these reasons, we can expect that the bit errors in memory systems will be correlated. In this work, we develop a mathematical multibit-per-cell memory model under a radiation environment. In this memory model, bit errors are correlated and the probability of errors depends on radiation parameters and time. For correlated bit errors, nonbinary codes can be more effective than binary codes. We will demonstrate that nonbinary LDPC code can outperform conventional BCH and RS codes in a correlated multibit error environment.
pp. 1957-1961

Agile Encoder Architectures for Strength-Adaptive Long BCH Codes

Raghunath Cherukuri (CodePhy Inc., USA)

Long Bose-Chaudhuri-Hocquenghem (BCH) codes are the choice of error correction codes for FLASH memory applications. Quite often, for FLASH memory applications, along with the need for high data rates, the packet length and the error correction capability of the code needs to be changed (strength adaptive BCH code). In this paper, we present a linear feedback shift register (LFSR) based architecture with a critical path bounded by (logN) (where N is the packet length), independent of the strength of the code and without any penalty in latency at a reasonable additional cost. The differentiating feature of the proposed architecture is the agility at which the strength can be changed at a very competitive cost.
pp. 1962-1966

Rebuilding for Array Codes in Distributed Storage Systems

Zhiying Wang (California Institute of Technology, USA); Alex Dimakis (University of Southern California, USA); Jehoshua Bruck (California Institute of Technology, USA)

In distributed storage systems that use coding, the issue of minimizing the communication required to rebuild a storage node after a failure arises. We consider the problem of repairing an erased node in a distributed storage system that uses an EVENODD code. EVENODD codes are maximum distance separable (MDS) array codes that are used to protect against erasures, and only require XOR operations for encoding and decoding. We show that when there are two redundancy nodes, to rebuild one erased systematic node, only 3/4 of the information needs to be transmitted. Interestingly, in many cases, the required disk I/O is also minimized.
pp. 1967-1971

Modeling, Detection, and LDPC Codes for Bit-Patterned Media Recording

Kui Cai (Data Storage Institute, Singapore); Zhiliang Qin (Data Storage Institute, Singapore); Songhua Zhang (Data Storage Institute, Singapore); Yibin Ng (Data Storage Institute, Singapore); Rathnakumar Radhakrishnan (Data Storage Institute, Singapore)

In this paper, we present a thorough and comprehensive study for bit-patterned media recording (BPMR), from a signal processing and coding perspective. We first propose a recording-physics-based generic channel model for BPMR, which includes all the major characteristics and impairments of the system. It also provides a fair basis for the performance comparison of different coding and detection schemes. We further propose various channel algorithms and techniques for BPMR, including a two-dimensional (2D) equalization scheme with one-dimensional (1D) generalized partial response (GPR) target to mitigate inter-track interference (ITI) and media noise, a maximum a posteriori (MAP) detector for BPMR with write errors, various low-density parity-check (LDPC) codes, as well as the iterative detection and decoding schemes. The corresponding performance gains are illustrated at 4Tb/in2.
pp. 1972-1976

4:00 PM - 5:30 PM

ACTEMT 06: Device Technologies

Room: Orchid AB
Chair: Brian Michael Kurkoski (University of Electro-Communications, Japan)

Techniques for Embracing Intra-Cell Unbalanced Bit Error Characteristics in MLC NAND Flash Memory

Guiqiang Dong (Rensselaer Polytechnic Institute, USA); Ningde Xie (Intel Inc., USA); Tong Zhang (Rensselaer Polytechnic Institute, USA)

Multi-level per cell (MLC) technique has been widely used to improve the storage density of NAND flash memory. However, bits stored in each MLC memory cell are subject to different bit error rates. In current practice, bits stored in each cell belong to different pages and all the pages are protected using the same ECC tuned for the worst-case scenario, which results in over-protection for many pages and hence storage capacity waste. In this work, we first develop a flash memory channel model to capture the dominant noise sources such as cell-to-cell interference and random telegraph noise. Using this model, we demonstrate the significant intra-cell unbalanced bit error characteristics for 3bits/cell NAND flash memory. We further develop two techniques that can better address this issue to minimize the overall redundancy overhead and hence improve effective capacity. First, we propose an aggregated page programming scheme by modifying the recently emerging full-sequence MLC NAND flash memory programming strategy, which can ensure all the pages experience the same overall bit error rates. Secondly, in the implementation of non-binary ECC such as RS code, we propose to combine a bit-error-rate-aware symbol grouping scheme in order to further reduce the required coding redundancy.
pp. 1977-1982

The Role of Non-Volatile Memory from an Application Perspective

Brett Kettering (DoD, USA); James Nunez (Los Alamos National Lab, USA)

HPC (High Performance Computing) applications spend a great deal of time waiting. The focus of our research and work has been to find ways to reduce the amount of time these applications spend waiting for the I/O subsystem to write or read their data in persistent storage. Current, emerging, and future NVM (non-volatile memory) technologies give us hope that we will be able to architect HPC systems that initially use these technologies in a memory and storage hierarchy, and eventually use these technologies as the memory and storage for the system. Our ultimate vision is a system where NVM will increase in density and performance, and decrease in cost such that it is both the active memory and persistent storage for HPC application data. There will be no more wasting time transforming data between its in-memory structure and its file-based structure; no more wasting time waiting for relatively slow persistent storage devices to record or provide this transformed data. Applications will use the NVM as they do DRAM-resident data today and rely on the same NVM to provide persistence with ownership and protections as does an HDD-based (hard disk drive) file system today.

pp. 1983-1987

Predicting Disk I/O Time of HPC Applications on Flash Drives

Mitesh R Meswani (San Diego Supercomputer Center, USA); Pietro Cicotti (UC San Diego, USA); Jiahua He (University of California, San Diego, USA); Allan Snavely (University of California, San Diego, USA)

As the gap between the speed of computing elements and the disk subsystem widens it becomes increasingly important to understand and model disk I/O. While the speed of computational resources continues to grow, potentially scaling to multiple peta flops and millions of cores, the growth in the performance of I/O systems lags well behind. In this context, data-intensive applications that run on current and future systems depend on the ability of the I/O system to move data to the distributed memories. As a result, the I/O system becomes a bottleneck for application performance. Additionally, due to the higher risk of component failure that results from larger scales, the frequency of application check pointing is expected to grow and put an additional burden on the disk I/O system [1]. Emergence of new technologies such as flash-based Solid State Drives (SSDs) presents an opportunity to narrow the gap between speed of computing and I/O systems. With this in mind, SDSC's PMAC lab is investigating the use of flash drives in a new prototype system called DASH [8, 9]. In this paper we apply and extend a modeling methodology developed for spinning disk and use it to model disk I/O time on DASH. We studied two data-intensive applications, MADbench2 [6] and RTM [5]. Our results show that the prediction errors for total I/O time are 14.79% for MADbench2 and and our efforts for RTM yield error of 9% for one category of read calls; a total of 3 categories of read/write are made by RTM. We are still investigating RTM, and in this paper we present our results thus far for both applications.

pp. 1988-1992

Beyond the Datasheet: Using Test Beds to Probe Non-Volatile Memories' Dark Secrets

Laura M Grupp (University of California, San Diego, USA); Adrian M. Caulfield (UCSD, USA); Joel Coburn (UCSD, USA); John Davis (Microsoft Research, USA); Steven Swanson (University of California, San Diego, USA)

Non-volatile memories (such as NAND flash and phase change memories) have the potential to revolutionize computer systems. However, these technologies have complex behavior in terms of performance, reliability, and energy consumption that make fully exploiting their potential a complicated task. As device engineers push bit densities higher, this complexity will only increase. Managing and exploiting the complex and at times surprising behavior of these memories requires a deep understanding of the devices grounded in experimental results. Our research groups have developed several hardware test beds for flash and other memories that allow us to both characterize these memories and experimentally evaluate their performance on full-scale computer systems. We describe several of these test bed systems, outline some of the research findings they have enabled, and discuss some of the methodological challenges they raise.

pp. 1993-1998

ACTEMT 07: Robust Memory Design

Room: Orchid C

Chair: Anxiao Andrew Jiang (Texas A&M University, USA)

Towards Longer Lifetime of Emerging Memory Technologies Using Number Theory

Lara Dolecek (UCLA, USA)

Emerging non-volatile memory devices show tremendous promise for a wide variety of applications, ranging from consumer electronics to server technologies. The advent of such multi-scale opportunities also carries a unique set of challenges. Increasingly popular Flash memory devices possess an intrinsic asymmetry during the write operation: programming memory cells to values lower than currently stored values is considerably slower and more costly than programming to higher values. It is critical to keep this cost low, as it directly affects memory lifetime and performance. Concurrently, demands for higher densities under reduced technology sizing make the data reliability a formidable objective. To address the compound issue of sustained and low-cost data reliability and high performance, in this work we propose a methodology to provide guaranteed immunity to a prescribed number of asymmetric errors, while having asymptotically negligible redundancy. Our construction uses ideas from additive and combinatorial number theory, and builds upon recently introduced coding schemes. We first show how this construction can be used in the single level cell (SLC) set-up, and subsequently extend the construction to the setting with several levels per cell, including multi-level cell (MLC) and triple-level (TLC) setting. We also discuss practical aspects of such schemes, including methods for systematic encoding, correction of limited-magnitude errors, and the additional protection under a certain number of bidirectional errors. This number-theoretic based approach is a promising direction for extending the lifetime of memories at sustained reliability.

pp. 1999-2003

Scrubbing with Partial Side Information for Radiation-Tolerant Memory

Euiseok Hwang (Carnegie Mellon University, USA); Seungjune Jeon (Carnegie Mellon University, USA); Rohit Negi (Carnegie Mellon University, USA); B. V. K. Vijaya Kumar (Carnegie Mellon University, USA); Michael K. Cheng (Jet Propulsion Laboratory, USA)

Memory systems used in space applications suffer from radiation-induced errors, either temporary upsets (soft errors) or permanent defects (hard errors or stuck-at errors). Scrubbing is a method to protect memory contents by periodically decoding the stored data to correct those soft and stuck-at errors then rewriting the corrected data back into memory. However, defective cells will remain and accumulate over time. Conventional coding disregards defective cells, however this may be inefficient for memory protection in space. In this study, alternative coding schemes for scrubbing are investigated, where the channel model depends on the cell states, defective or not, and the encoder uses channel state information (CSI) or side information. At every scrubbing, the error correcting code (ECC) decoder provides partial CSI back to the encoder and the encoder uses the CSI to improve the performance of memory systems with scrubbing. Information theoretic limits of the channel with partial CSI are investigated and several coding schemes are introduced to mitigate the effects of defective cells, particularly those caused by stuck-at defects. In addition, coding schemes with partial CSI are concatenated with binary Bose-Chaudhuri-Hocquenghem (BCH) codes to protect memory contents from both soft and stuck-at errors in space radiation environments. Numerical simulation results show that scrubbing with partial CSI improves reliability over the state-agnostic approaches.

pp. 2004-2008

FFT Processing Through Faulty Memories in OFDM based Systems

Muhammad S Khairy (University of California, Irvnie, USA); Amin Khajeh (University of California, Irvine, USA); Ahmed Eltawil (University of California, Irvine, USA); Fadi J Kurdahi (University of California, Irvine, USA)

In the widely used OFDM (Orthogonal Frequency Division Multiplexing) systems, the FFT and IFFT pair are integral parts used to modulate and demodulate the data constellation on the sub-carriers. Within such systems, embedded buffering memories occupy a large portion of the area and hence directly control the overall metrics of the system including power consumption and cost. This paper presents a unified statistical model that accurately reflects the impact of random embedded memory failures due to power management policies on the overall performance of an OFDM-based communication system. The proposed model expands the design space, by allowing the designer to replace the faulty hardware with perfect hardware while propagating the resulting distribution (due to voltage overscaling) through the system. The proposed framework enables system designers to efficiently and accurately determine the effectiveness of novel power management techniques and algorithms that are designed to manage both hardware failures and communication channel noise, without the added cost of lengthy system simulations that are inherently limited and suffer from lack of scalability.

pp. 2009-2014

Peer-to-peer Technologies Applied to Data Warehouses

Simone Cirani (University of Parma, Italy); Lorenzo Melegari (University of Parma, Italy); Luca Veltri (University of Parma, Italy)

Data mining and user data collection applications, like Facebook and Yahoo, make dealing with huge amounts of data more and more frequent. A solution to cope with this problem is to spread data over multiple network-connected physical devices. Having more devices, though, means increasing system complexity and introducing additional possible points of failure. Moreover, despite the capacity of hard drives as massive storage systems has increased extremely during years, the speed at which data can be accessed has not. In order to address this problem, over the years, distributed file systems, such as NFS and HDFS, have been designed and deployed. Such systems provide access to files stored on multiple hosts connected through a computer network in a transparent way to users. The peer-to-peer network paradigm has been introduced to overcome some limitations of the client-server architecture by adding features, such as scalability, fault-tolerance, and self-organization. In this work, we present a solution that integrates peer-to-peer network support to HDFS in order to realize a flexible, low-cost and, dynamic distributed file system.

pp. 2015-2019