En-Stor: Energy-Aware Hybrid Mobile Storage System using Predictive Prefetching and Data Mining Engine

Mais Nijim, *Member, IAENG*, Hisham Albataineh

***Abstract*—In this paper, we proposed an energy-aware data mining predictive prefetching technique for hybrid storage systems called En-Stor that uses data mining predictive prefetching to save energy. We used as an example of the hybrid storage systems mobile hard disk drives (MHDDs) and solid-state disks (SSDs). As the SSDs are much more energy-efficient than MHDDs, aggressive prefetching for data from MHDDs will enable them to be in the standby mode as much as possible in order to save power. En-Stor differs from existing energy-aware prefetching techniques in two ways. First, En-Stor is implemented in hybrid storage system where MDDS and SSDs are used. Second, it used data mining predictive prefetching techniques to prefetch the data from MDDs to SSDs to increase the standby time of the MDDs, hence reduce the energy consumption. The data mining predictive prefetching techniques will also increase the performance of the system because most of the requested data will be stored in the SSDs, which offer much faster access time than the MDDs. A simulator was created to evaluate the performance of the En-Stor. Our results show that En-Stor reduces the power consumption of the mobile disk drives by at least 85% when compared with non En-Stor system.**

***Index Terms*—data mining engine, predictive prefetching, solid-state disks, mobile disks**

I. Introduction

MOBILE storage system is defined as an array of hard disk drives with a small form factor that are connected to a host using a storage interface such as SCASI in a mobile computing platform [1]. Mobile storage systems play an important part of several significant applications, such as mobile data centers for disaster response; where these centers have the capability to hastily implement computing and communications infrastructure to assist local officials and relief workers in major disasters [2]. For example, the several mobile data centers developed by Cisco (CSCO) in the wake of disaster recovery. Mobile storage systems has been widely used in various applications, such that in cars, trucks, aircraft, airplanes, or ships that have onboard generators, UPS, multiple high-capacity servers, and satellite Internet links [9].

Manuscript received August 4, 2017; revised April 27, 2018.

M.N. Author is with the Department of Electrical Engineering and Computer Science, Texas A&M University – Kingsville, Kingsville, TX, 78363 USA. (e-mail: mais.nijim@tamuk.edu), H.A is with the Department of Physics and Geoscience , Texas A&M University-Kingsville, TX, 78363, USA.

One of the advantages that mobile disk arrays have is that they face several new challenges including harsh operating environments, and very limited power supply unlike the stationary data centers that are roofed in building [10]. Mobile disk arrays antithesis stationary disk arrays where electrical power is assured, Mobile disk arrays usually don't have direct power supply, and it is mainly depends on batteries or gasoline generators. Consequently, energy-saving issues becomes more crucial for mobile disk arrays because their energy consumption can significantly affect the lifespan of the entire mobile systems [10].

Hybrid storage systems provide cost-effective solutions for large-scale data centers without significantly affecting I/O response times. Usually, the top level of the hybrid storage systems is considered to be the most energy efficient. The Hybrid storage system comprises of a two level system [4][12], the Mobile Disk Drives (MDDs) at the lower level and an array of Solid State Disks (SSDs) is at the higher level, which is more energy efficient [4,5]. The data request will be checked first in the solid-state disks, if it is not found, the lower MDDs will be checked for data retrieval.

Mobile data centers applications offer high mobility and a fast large-volume data processing capabilities. One way to achieve efficient power consumption is to spin down disks when they are sitting in idle state [6]. However, spinning down disks is effective only if the disks are to remain in the standby state for long time periods. Even though, the research on energy efficient mobile storage system is still in early stage, we try to reduce energy dissipation in hybrid mobile disk systems while maintaining high I/O performance using data mining predictive prefetching. To remedy this problem, we proposed an En-Stor, a hybrid storage systems that consist of two levels, the first level consists an array of Mobile disk drives (MDDs), and the second consists of an array of solid-state disks (SSDs). Prefetching the data from MDDS to SDDs will vastly enable the array of MDDs to be in the standby mode as long as possible. Predictive prefetching with the help of the data-mining engine will predict which data needs to be transferred from MHDSs to SSDs before it is requested.

The rest of the paper is organized as follows. Section two describes the proposed system design that consists of data placement, dynamic solid-state disks partitioning, and the proposed data mining engine for predictive prefetching. Section three evaluates the performance of the proposed En-

Stor technique by comparing it with a system that does not employ En-Stor. Finally section four concludes the paper.

II. THE SYSTEM DESIGN

This section describes the En-Stor architecture and the set of its significant features which includes: parallel data transfer using data mining predictive prefetching, intelligent solid-state disks partitioning, and data placement. Figure 1 illustrates the hybrid architecture for En-Stor, which contains an array of solid-state disks (SSDs), an array of mobile hard disks drives (MDDs), and an energy-aware solid state/mobile disk controller. Remember that the number of SSDs and MDDs are independent of each other. The energy aware controller coordinates multiple modules, including power management, data partitioning, disk request processing, and data mining predictive perfecting schemes. The solid-state disks perform as a non-volatile cache to boost up the I/O performance and to improve the energy efficiency by absorbing the disk traffic fluctuations. The solid-state disks respond to read and write requests. A read miss on the solid-state disks causes a hit at the mobile disks. To boost up the performance, the disk request will be prefetched from the mobile disk to the solid-state disks before it is requested. The mostly frequent used data will be placed in the SSDs first. If the solid-state disks are full, the write request will be directed back to the mobile disks using the least recent used policy.

A data mining predictive prefetching scheme is designed to bring data into solid-state disks before its use. Apart from the prefetching scheme, we developed a write strategy to energy-efficiently handle writes using solid-state disk disks. The write I/O load imposed well balanced by equally distributed write request to all the active solid-state disks. We developed a disk manager that is responsible for the following activities. First, the disk manager aims to minimize the number of active disks while maintaining reasonably quick response time for disk requests. Second, the manager must deal with the read and write requests redirected from the solid-state disk in an energy-efficient way. Third, the manager has to energy efficiently move data among the solid-state disk, and the mobile data disks [4].

The energy aware controller is considered as the milestone of the proposed architecture. It consists of the data placement technique; solid-state disks partitioning, data-mining engine predictive prefetching for power management.

A. Data Placement

We assume that the data initially is distributed between the MDDs and the SSDs. More data will be allocated to the MDDs level since the MDDs have larger storage capacity than the SSDs. The less urgent data will be placed in the MDDs while the top priority data will be placed in the SSDs. In our previous study [4,5,6], we indicated that the lower level of a hierarchy in storage systems has a lower priority while the higher level of the hierarchy has higher priority. This implies that our proposed data mining predictive prefetching algorithm move that data that is likely to be accessed in the future to the higher priority level that is in this case the SSDs.

B. Dynamic Solid State Partitioning

The performance and the energy consumption of the En-Stor can be significantly improved by employing a number of solid-state disks. SSDs can increase the performance of the En-Stor by shortening the average response time of the disk request. Solid state partitioning is of importance because solid-state management can judiciously cash data blocks for future access. In addition to boosting performance, by prefetching data blocks from disk drives to the corresponding solid state partition, for example, the disk drive that has the highest load will get the largest partition of the SSD, a solid state management scheme will shorten the average response time of disk requests, making it possible to reduce energy consumption by having the disks stays in sleep mode as much as possible. In this paper, we use the conventional solid-state management algorithm (Least Recently Used algorithm or LRU for short). However, our En-Stor algorithm does not limit itself to the use of the LRU policy; rather, our approach can be employed in combination with other replacement policies.

The SSDs dynamically partition the SSDs blocks in such a way to increase the performance and also reduce the energy consumption. The partitioning is motivated by the observation that the storage systems workload is not equally distributed between the MMDs, i.e. a mobile disk that has a higher workload than other MMDs would get more SSD space than the others MMDs. Each SSD partition of the disk is managed separately using LRU replacement algorithm. If the workload of the mobile disk is high, give it more solid-state disk space, which will keep it in a sleep state as much as possible.

C. Data Mining for Predictive Prefetching

Data mining algorithms can be applied to predictive prefetching techniques, allowing storage systems to identify the data access patterns and predict trends with past request information [7][8][11]. The current data mining algorithms are not applicable for hybrid storage systems where mobile disk systems and solid-state disks are integrated. To remedy this problem, we proposed a data-mining engine for hybrid storage systems where MMDs and SSDs are used. Figure 2 shows the architecture of the data-mining engine of the hybrid storage systems. It consists of two parts, the FTP server, and the storage area network (SAN). The FTP server which contains a set of solid-state disks will keep the most highly accessed data, while the normally accessed files will be kept in the SAN.

The first step of the data mining engine is the pattern discoveries that access the user history table and generate the candidate pattern table based on specified business rules such as that generate the urgent pattern table and the non urgent pattern table. If the data is in the SSDs and was found in the non-urgent pattern table, it would be transferred from the SSDs to the MDDs. If the data is located in the MMDs and was found in the urgent table pattern, it would be prefetched from the MMDs to the SSDs. To discover the patterns, we use the temporal market basket analysis (MBA) to generate the candidate patterns.

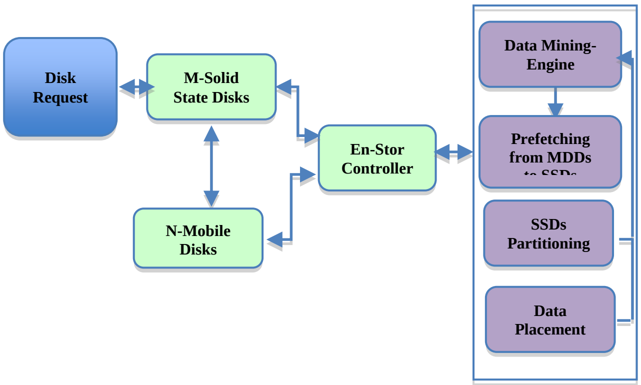

Figure 1. En-Stor Hardware Architecture

The second step is the pattern search algorithm, which goes through all the user requests and produces the requests to be prefetched from MMDs to SSDs. The responsibility of the pattern search algorithm is to take the temporary online request window as the input, look up the activated pattern table, generate the prefetch file list and pass the prefetch file list into the prefetching function. We built a meta-data management module as shown in figure 3 to interact with the data-mining engine for the predictive prefetching. This module could be applied to any kind of hybrid storage systems. The meta-data manager will provide required information that helps in managing data-intensive applications, thereby achieving optimal I/O performance. We seamlessly integrated the core part of the meta-data management module with the En-Stor architecture.

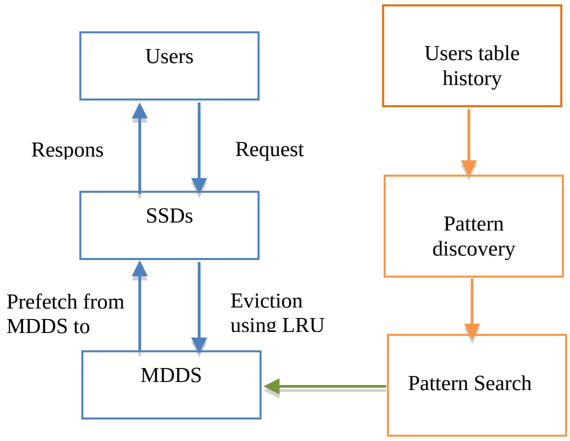

Figure 2: Data Mining Architecture for E2M-Stor

Figure 3 shows a meta-data management framework, which consists of clients, a meta-data management layer, and the En-Stor architecture that contains two layers; the MDDs and the SDDs. The meta-data management module will be used to keep track of what data sets are accessed by clients, making data access information available for the data-mining engine. With the meta-data management module in place, our proposed data-mining predictive prefetching will be able to quickly locate data sets from solid state disks and mobile disks. Furthermore, the meta-data manager will provide the predictive prefetching mechanism important information regarding performance characteristics like I/O bandwidth and capacity. Our meta-data management module will adaptively read and modify meta-data on the fly after performing data-mining-based predictive prefetching algorithms. There are two approaches to maintaining and monitoring data locations. First, we increase the number of empty fields of the metadata before it is stored on the MDDS along with data content. After data prefetching is performed, corresponding metadata on storage devices must be modified dynamically. Second, we developed a meta data repository to record the meta-data of data sets. The repository is implemented to store and keep track of system level and domain independent meta-data using a uniform interface.

III. PERFORMANCE EVALUATION

To evaluate the performance of En-Stor, we implemented a simulation toolkit, which consists of a set of core components and supplementary components. The core components include pattern discovery algorithms, pattern search algorithms, a meta-data management module, and parallel data transfer modules between the MDDs and the SDDs. The supplementary components consist of all the functions that are necessary to implement the En-tor algorithm. We will compare the performance of mobile disk systems with En-Stor against another system without employing En-Stor. In this study, a system that does not employ En-Stor is a standard mobile disk system.

In our simulation, we made use of the following two performance metrics to demonstrate the effectiveness of the En-Stor algorithm:

(1) Energy Consumption: is the total energy consumed by the mobile disk systems.

(2) Energy conservation ration.

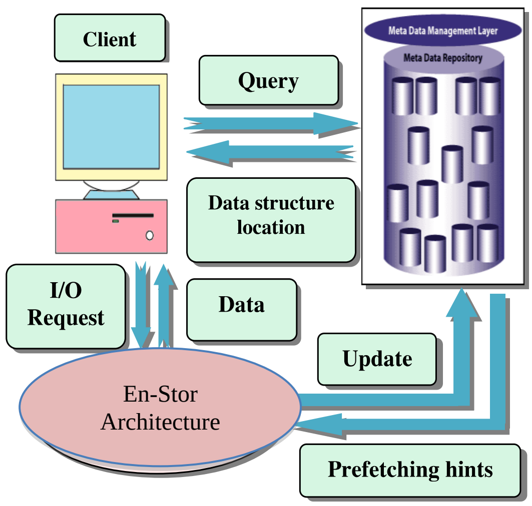

Figure 3: Meta Data Management for En-Stor

A. *Impact of Arrival Rate*

In this experiment, we evaluate the impact of mobile disk arrival rate on the performance of the system, energy consumption, and energy conservation. The arrival rate ranges from 0.1 to 0.5 No./sec. Figure 4(a) shows that En-

Stor algorithm significantly reduces the energy consumption of the mobile disk system by an average of 80%. The improvement of energy consumption is due to that fact that solid state disks is used which will let the mobile disk systems to stay in sleep mode as much as possible, also the data mining engine is used to predict the data that needs to be transferred from the mobile disk system to the solid state disks. Figure 4(b) show that the energy conservation ratio is reduced as the arrival rate increases. The energy conservation ration for the En-Stor is higher than the energy conservation in the mobile disk system that does not employ En-Stor.

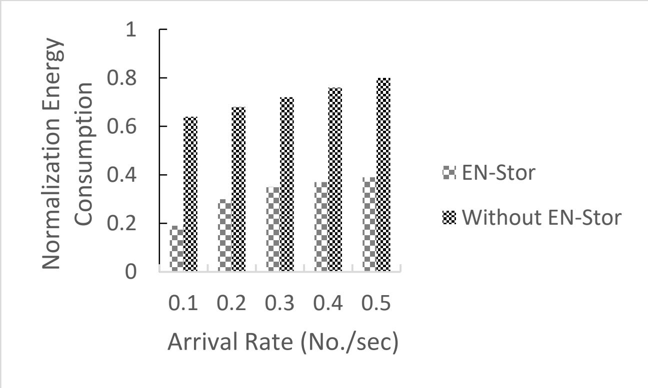

Figure 4(a). Impact of Arrival rate on normalized energy consumption

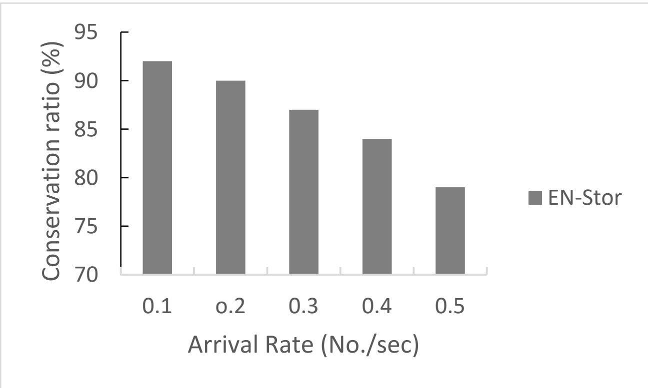

Figure 4(b). Impact of Arrival rate on energy conservation ratio

B. Impact of Mobile Disk Bandwidth

In the second experiment we measure the performance impacts of disk bandwidth on the energy efficiency of En-Stor. We vary the disk bandwidth from 10 to 50 MB/s, with an increment of 10 MB/s.

Figure 5(a) obviously reveals that the normalized energy consumption decreases as the mobile disk bandwidth increases. This is due to the fact that energy consumption is a product of power and mobile I/O processing time, which is remarkably reduced with the increasing disk bandwidth.

Figure 5(b) show that the energy conservation ratio is reduced as the disk bandwidth increases. The energy conservation ratio for the En-Stor is higher than the energy conservation in the mobile disk system that does not employ En-Stor.

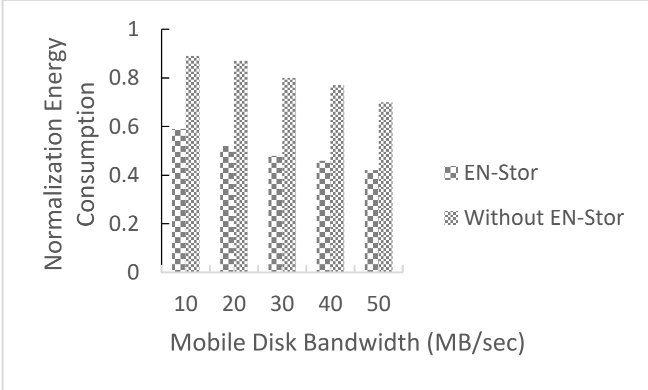

Figure 5(a). Impact of mobile disk bandwidth on conservation ratio when arrival rate is 3no./sec.

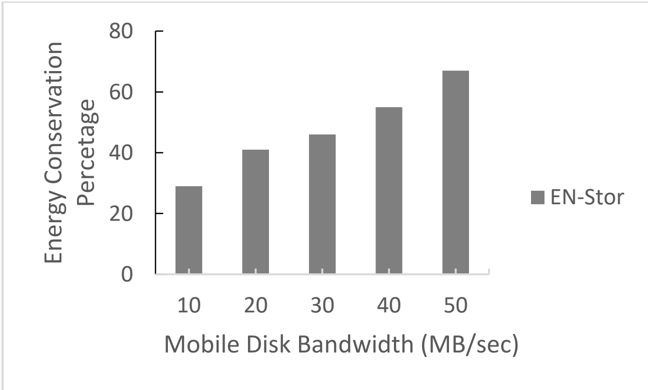

Figure 5(b). Impact of disk bandwidth on energy conservation percentage.

C. Impact of Solid State Disks Size

We measure how energy consumption percentage depends on the size of the solid-state disks size. Figure 6(a) illustrates that the normalized energy consumption percentage relative to the case where the solid-state disks are used as a cache in the En-Stor strategy. We observe that increasing the solid-state disks size, provides the mobile disks more opportunities to stay in inactive mode. As a result, the energy consumption of the mobile disks will be decreased.

Figure 6(b) show that the energy conservation ratio is increased as the solid-state disks size increase.

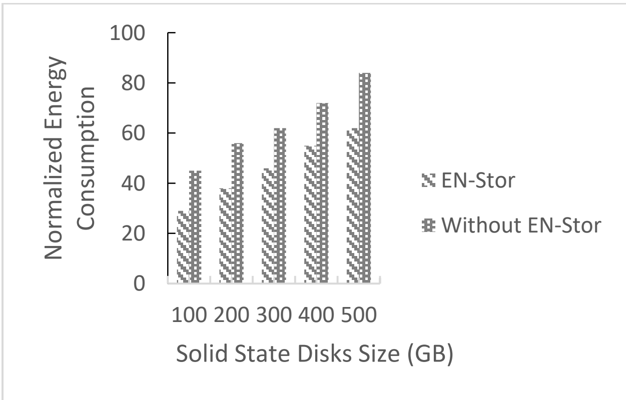

Figure 6(a). Normalized energy consumption against the solid state disks size

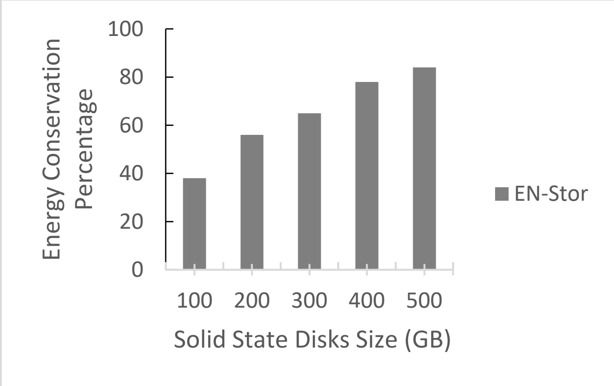

Figure 6(b). Impact of solid state size on energy conservation percentage

IV. CONCLUSION

This paper presented a novel energy aware predictive prefetching algorithm using data mining for mobile storage systems called En-Stor where Mobile disk arrays usually don't have direct power supply, and it is mainly depends on batteries or gasoline generators. Consequently, energy-saving issues becomes more crucial for mobile disk arrays because their energy consumption can significantly affect the lifespan of the entire mobile systems. The main idea behind En-Stor is to prefetch data in the hybrid storage systems from the lower level which consists of a number of Mobile Disk Drives (MHDDs) where it consume lots of energy to the higher level which consists of a number of Solid-State Disks (SSDs) where it is energy efficient. En-Stor used data mining predictive prefetching techniques to prefetch the data from MDDs to SSDs to increase the standby time of the MDDs, hence reduce the energy consumption. The simulation results revealed that when comparing the performance of mobile disk systems with EN-Stor against another system without employing En-Stor, the En-Stor algorithm significantly reduces the energy consumption of the mobile disk system by an average of 80%.

REFERENCES

[1] T. Xie, H. Wang, "MICRO: A Multilevel Cashing Based Reconstruction Optimization for Mobile Storage Systems", *IEEE Transactions on Computers*, Vol. 57, Issue 10, pp. 1386-1398, 2008.

[2] B.A. Mah, S. Seshan, K. Keeton, R.H. Katz, and D. Ferrari, "Providing Network Video Service to Mobile Clients," *Proc. Fourth Workshop Workstation Operating Systems*, pp. 48-54, 1993.

[3] Y. Kim, A. Gupta, et. al., "HybridStore: A cost-efficient, high performance storage systems combining SSDs and HDDs", *19th IEEE International Conference on Modeling, Analysis, and Simulations of Computer and Telecommunication Systems*, 2011.

[4] M.Nijim, " HyBuM: Energy Efficient Hybrid Mobile Storage Systems using Solid States and Buffer Disks" Journal of Computer Communication, Vol. 3, Issue 4, 2015.

[5] Mais Nijim, Young Lee' Nuri Yilmazer, Remzi Seker," A DATA MINING ALGORITH FOR MULTI LEVEL PREFETCHING IN STORAGE SYSTEMS", *Proc. of Ubiquitous Computing and Communication Journal*, August 2011, pp. 10-19.

[6] Mais Nijim, Young Lee, Kiran Bellam, HyBuM: Hybrid Energy Efficient Architecture for Mobile Storage Systems, *9th IEEE International Conference on Information Technology: New Generation*, April 2012.

[7] J. Griffioen, R. Appleton, " Reducing file system latency using a predictive approach", *IEEE USENIX*, pp. 197-207, 1994.

[8] A.E. Papathanasiou, M.L.Scott, "Energy efficient prefetching and caching, *IEEE USENIX*, 2004.

[9] Mobile Emergency Datacenter, North Am. Access Technologies,http://www.naat.com/Disaster%20Recovery/mobile_datacenter .htm, 2006.

[10] S. Gurumurthi, J. Zhang, A. Sivasubramaniam, M. Kandemir, H. Fanke, N. Vijaykrishnan, and M. Irwin, "Interplay of Energy and Performance for Disk Arrays Running Transaction Processing Workloads," Proc. IEEE Int'l Symp. Performance Analysis of Systems and Software, pp. 123-132, Mar. 2003.

[11] M. Kuo, H. Chang, D. Chang," A Disk State- Aware Task Scheduler with Energy Efficient Prefetching and Caching", *Proc. IEEE International Conference on Cloud and Green Computing*, 2013.

[12] J. Yoon, E. Nam, Y. Seong, et. al.," Chameleon: A High Performance Flash/Fram Hybrid Solid State Disk Architecture", *IEEE Computer Architecture Letters*, Issue 1, Vol. 7, pp. 17-20, 2008.

[13] S. Wu, G. Chen, K. Chen, *et. al.*, " HM: A Column-Oriented MapReduce System on Hybrid Storage", *IEEE Transactions on Knowledge and Data Engineering*, Vol. 27, Issue 12, pp: 3304-3317, 2015.

Mais Nijim is an Associate Professor in the Department of Electrical Engineering and Computer Science at Texas A&M University-Kingsville. She received her master degree in computer science from New Mexico State University in 2004. Dr. Nijim received her Ph.D. in Computer Science from New Mexico Tech in 2007. Her research interests include parallel and distributed systems, real time computing, storage systems, fault tolerance, and performance evaluation, big data analytics, data mining, and wireless sensor network.

Hisham Albataineh is an Assistant Professor in the Department of Physics & Geosciences at Texas A&M University-Kingsville. He received his master and PhD in high energy physics from New Mexico State University in 2009. His research interest includes but not limited to high and medium energy physics, renewable energy, and big data analysis.