

Towards Distributed Wilderness Search using a Reliable Distributed Storage Device built from a Swarm of Miniature UAVs

Paul Gaynor, Daniel Coore

Abstract—Searching a wilderness environment for missing persons is a task that is well suited to a swarm of miniature UAVs, given the high availability of low cost drones and the large geographic area that is to be covered. The use of a swarm for detection provides an alternative to automated detection solutions that require high processing power to search all images encountered for all possible combinations of items that may be detected. Challenges faced in implementation include limited storage and processing capacity per node, the designation of data storage points when the swarm is far from any base station, and inhospitable wilderness environments that may destroy UAVs. The solution that we employ leverages on the use of distributed in-network storage that allows each UAV to behave as a sensor node which is specialized to the task of detecting certain features. Our data storage strategy self organizes in an efficient way using a randomized selection of backup addresses, and is robust to the loss of data even after events that cause the destruction of a large number of nodes. Analysis and simulations show the survivability of data, how the communication latency varies with the number of backup addresses and that a set of simple processors can be used to collaboratively perform distributed searches with relatively high accuracy.

I. INTRODUCTION

WILDERNESS Search and Rescue (WiSAR) involves finding and assisting humans who are lost or injured in remote areas such as deserts and forests[1]. The high availability of low cost miniature UAVs facilitates increased search effectiveness and reduced risk to personnel by allowing rapid assembly and deployment of a swarm. For the swarm to reduce the burden on the search party, it should manage itself. The swarm should also be responsible for analysis of the data it encounters. It is possible to develop specialized drones that will be have sufficient storage and processing power to provide full service feature detection, however the cost in developing a fleet of high powered UAVs could prove prohibitive. Our solution involves a set of simple drones equipped with a camera, a radio, and the default RAM capacity and processing speed available with a Raspberry Pi[2](ie. 512 MB and 700 MHz respectively). Each drone has code running on its processor that allows it to take part

in random backup address assignment and store and forward communication[23]. Feature detection is accomplished using the Voila-Jones method[3] and a bag-of-features tally[4] is used to determine if an area contains sufficient evidence to warrant further investigation. The self-organizing data storage strategy is robust to the loss of nodes. Our major contribution is to show how a self organizing, scalable, robust, reliable, distributed data storage device built from simple components, can be employed to provide a platform on which distributed search can be executed.

The paper is organized as follows: Section II looks at related work. Section III presents our solution and Section IV presents our results and discussion. Sections V presents our conclusion and future work.

II. BACKGROUND WORK

A. Wilderness Search and Rescue

A WiSAR process often begins when a report is made by a concerned friend or relative of a missing person, to someone in authority at the missing person's last assumed location[1]. A WiSAR team, including an incident commander is then assembled primarily from volunteers[5]. The WiSAR team then locates the assumed incident scene, assesses environmental conditions and determines the equipment necessary for response. A perimeter is then established, and a search plan is designed that may employ one of several search patterns. Documented patterns include the hasty search, the constraining search, the high probability region search and the exhaustive search[5]. A hasty search involves quickly checking locations close to the missing person's last known location such as tents and trails. Speed is important in checking these locations as the probability of locating useful information reduces as time passes[5]. The constraining search pattern is used to help establish a perimeter by checking possible exit routes from the assumed search area. The high probability region search is used to check sections of the search area that the missing person is likely to be, based on evidence from the hasty and constraining searches. Exhaustive search patterns involve members of the WiSAR team combing the search area looking for clues that suggest the missing person had been at a specific location. Significant clues that result from an exhaustive search may trigger a high probability region search[5].

P. Gaynor is an Assistant Lecturer in the Department of Computing, The University of the West Indies, Mona Campus e-mail:paul.gaynor@uwimona.edu.jm.

D. Coore is a Senior Lecturer in the Department of Computing, The University of the West Indies, Mona Campus.

Manuscript received April 17, 2014.

B. Automated detection

Camera equipped UAVs have achieved highly publicized success in surveillance operations. The aerial imagery that they supply can be used to support WiSARs tasks by reducing the need for humans to traverse all possible search areas. Search time and exposure to environmental risks are therefore reduced. It is suggested in [6] that several people, each with a specific role are required for UAV assisted WiSAR. The suggested roles include operating the UAV and analysing the images obtained. Employing humans to satisfy these roles for a swarm of UAVs will however require a significant amount of human resources. Automated flight path management for swarms of UAVs is discussed in [7]. The approach used in [8] allows each UAV in a swarm to be queried as a part of a remote distributed database using the COUGAR[9] infrastructure through the aid of a query proxy layer. The implementation used in [8] allows images captured by remote UAVs to be available, however results show that connectivity varies with network configuration and that communication costs are a significant burden on the infrastructure. We will not assume that any of our miniature UAVs operating in a remote location will possess sufficient power to directly communicate with a base station.

As a continuous monitoring network, the best strategy for communication with a base station is to record data in batches and transmit relevant data in a hop by hop, store and forward nature[10]. Limited energy reserves prohibit the transmission of every detected image through the network. It is however possible to perform automated image analysis on a node to detect features of interest using a cascade of wavelet transforms[3].

Object detection is commonly accomplished using the Viola-Jones object detection framework[3]. The Viola-Jones framework applies filters to Integral Images that are calculated with one pass over the image. Each filter behaves as a weak classifying function that has an accuracy slightly better than random. A combination of several of these functions results in a detector that demonstrates high accuracy. Figure 1 shows how filters can be used to test if a human outline exists in an image. An example of the application of features involves testing if a section of an image is likely to contain a human outline, by first testing if the part of the image tested looks like the cascade in part (a), that is, if the center part of the tested area is darker than the area to the left and right. If the initial test is true, then the area of the image that passed the first test will have other tests applied to it. The test applied in part(b) checks if the lower part of the image will have legs that are separated.

An entire cascade of features produces very accurate results with low computational requirements. Viola and Jones were able to achieve an accuracy rate of 95% for human face detection using 200 simple features[3]. Cascades are normally implemented as XML files. Image processing libraries like OpenCV[11] commonly includes Haar cascades that can detect specified features. Instructions on the creation of

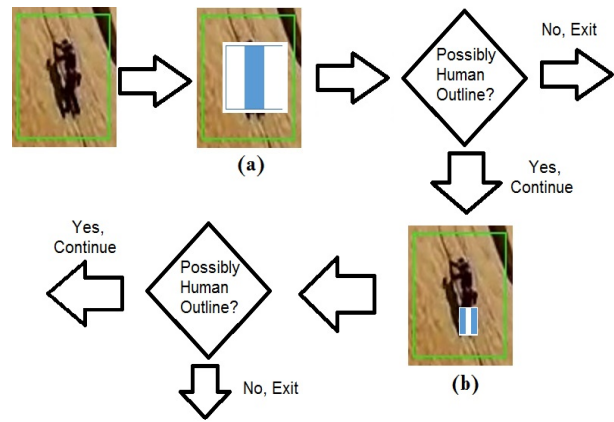


Fig. 1. Haar features that may help to detect a human outline.

customized cascades are also provided. Generated cascades will have a range of file sizes. Cascades exist with file sizes of up to 5 MB. If each cascade is specialized to extract a single feature, and there are numerous features to be detected, it becomes impractical to expect a platform with the limitations of a Raspberry Pi to be able to perform detection for more than a few features and still have space to store image data. Evidence of missing persons in a wilderness may have no pre-defined spatial organization between features and therefore a Bag-of-Features method[4] is practical. Bag of Features methods represent images as un-ordered sets composed of local features. An analogy can be drawn to the Bag-Of-Words representation for document comparison[4] where a document is represented as a normalized histogram of word counts. While training the Bag-Of-Features representation, millions of local features are sampled from the training data, and clusters of features are considered to be visual words[12]. It is therefore possible to create an image representation comprised of an unordered composite of features that may give an indication of the presence of a missing person. An example is a set of images of a hiker's clothing or tools.

C. Self Organizing Reliable Data Storage

Sensor nodes operating in dangerous or inaccessible environments face a high probability that they may be destroyed, and that communication links to them will fail. One commonly implemented solution to maintain data availability is the use of redundancy. Added redundancy creates other challenges however, including the need to copy data between the redundant points in a way that preserves battery life[13]. The persistent node concept [14] organizes regions of processing nodes into groups that implement redundant storage, and moves data away from high risk areas. A high risk area is defined as a region where a node has recently failed. The reactive strategy increases the probability of survival of data in a network for a small number of failures but there will still exist a risk of data loss if data not copied before the area is deemed high risk. An alternative strategy, such as the one employed with the use of growth codes [15], proactively aims to achieve redundancy in nodes by copying

coded information to neighbors. Before copying, each node encodes previously received data with current information, resulting in successively higher levels of coding. A growth code strategy preserves the data on a node in the event of small scale node failures. The high information density attained using the growth code strategy is however vulnerable to a disaster affecting a geographic region, where a source data node, along with the set of neighbors, can be destroyed without warning, resulting in data loss[15]. One tried and tested method of acquiring reliability through redundancy is shown with RAID technology[16]. The supporting infrastructure to deploy RAID includes a high level of connectivity, which may not be available for wireless deployments. There have however been implementations that abstract RAID-like services over a sensor network such as S-RAID[17], implemented as a layer that abstracts redundancy. Experiences from RAID implementations highlight the need for tradeoffs, such as those between the amount of redundancy deployed and the level of throughput to be expected. A convincing argument on whether reliance on centralized storage should be preferred to multiple distributed copies of data is provided in [18]. The argument presented is that communication vs storage tradeoffs should be scrutinized. It is posited that if the storage energy cost is less than the communication energy cost, then in-network storage should be exploited to save energy. The author then points out that storage energy costs using low power flash devices have become much lower than communication energy costs in conventional systems.

Limited power reserves on sensor nodes prescribes that sensor networks deployed for extended periods of time will need to take measures to reduce power usage. Research in [19] that tested power usage when nodes were sleeping, receiving or transmitting, show that the power usage of a node varies depending on the state of the node. A node that is sleeping uses a small amount of power. A node that is receiving data uses more power, and a node that is transmitting uses even more power. To prolong lifetime, nodes should (as far as is practical for application purposes), spend as much time as possible sleeping, less time receiving, and even less time in transmitting mode. Of course, there will be trade-offs between latency and power efficiency as discussed in [20]. Allowing nodes to wait for extended periods of time before transmitting received data could give rise to the risk that data may become corrupted in a way that it is not recoverable. There has been sufficient recent work on index coding [21] however to allay those concerns. Index coding was used to describe a strategy for video distribution network over assymmetric satellite links. The basic index coding problem involves a sender, with all the information that is to be transmitted to receivers. Receivers can collect and store information apart from the information that they are waiting on. Eventually each receiver builds up a "Has Set" (set of packets that it has), as well as a "Wants Set" (set of packets that it wants). The source will be informed about the Has sets and the Wants sets of all the destination nodes. Index coding seeks to determine the most efficient transmission that will allow each node to satisfy its Wants

set. Index coding has also been used to implement distributed storage, by taking advantage of local repair using previously overheard side information.

III. WILDERNESS SEARCH USING DISTRIBUTED DATA IN THE SWARM OF MINI-UAVS

A. The storage protocol

We have coined the term **Address Coding** for our storage protocol, as it is based on the exchanging of encoded information between backup addresses. The coding applied to messages transmitted between nodes is a function of the backup addresses of the communicating nodes. The Address Coding protocol allows a node to be used as an interface point to the wireless sensor network. An interface node exposes three operations to an external user, namely, `init(k)`, `store(a, v)` and `get(a, responseTimeout)`. The interface node also exposes the variable `response`. The function `init(k)` initializes the system, and sets the number of backup addresses in the network to the value k . Subsequently, the user may store a value v at address a with the command `store(a, v)`, where $0 \leq a \leq k - 1$. To later retrieve what is stored at address a , a user issues the command `get(a, responseTimeout)` at the interface node, and then after a user specified interval `responseTimeout`, tests the value of `response`.

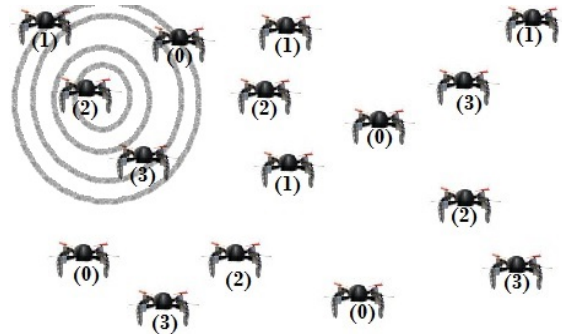


Fig. 2. Sample distribution attained with four backup addresses and sixteen nodes.

1) *Initializing the network:* The initialization process is invoked by injecting an `init` message with parameter k to a node p that acts as an injection point to the network. In response to the injection, p selects a backup address and then broadcasts the `init` message with parameter k to all its neighbours. A node q in the neighbourhood of p that subsequently receives the `init` message will trigger its `receive` method. The `receive` method triggers the `init` message on q , causing q to select a backup address, and then re-broadcast the `init` message to all neighbours of q . Immediately after issuing a broadcast a node will sleep for a short random delay. The `init` message will therefore propagate in an expanding ring from p to the edge of the network.

Let n be the number of nodes in the system. Each node independently selects one of k address classes. A determination of the probability that each class is represented after the

address selection process is given by the Coupon Collector’s problem[22]. Analysis shows that for k classes, the expected number of coupons, n , that will be randomly inspected before all classes have been encountered is $k \ln k$, with variance $< 2k^2$ [22]. Based on Chebychev’s principle[24], we expect full coverage within three standard deviations (99.7% probability) as long as

$$n \geq k \ln k + 3\sqrt{2} \cdot k \quad (1)$$

2) *Saving data to the network:* To save data into the network, the user injects to a node p the `store` function with arguments a , the destination address and v , the value to be stored at the destination address. p will determine whether the data will be stored in its local memory by checking to see if its assigned backup address is a , and then broadcasts a signal to immediate neighbours that invokes a `store` message to be invoked on those neighbours. A neighbour q of p that receives the broadcast message checks if v should be stored at q , and then broadcasts a signal to immediate neighbours that invokes a `store` messages on those neighbours. The `save` message will propagate in an expanding ring from p to the edge of the network.

3) *Extracting data from the network:* A user outside the network will request data by issuing the message `get(a, responseTimeout)` to an interface node p . The argument a denotes the backup address from which the user requires data. `responseTimeout` defines the time interval the user is willing to wait on a response from the system. After `responseTimeout` has elapsed, the user will test the node p to query the variable `response`. The current implementation is optimized for speed of response whereby the user that queries will accept the first response returned to the interface node to be accurate. In this implementation the distributed storage device is operating as a distributed cache, and will be well suited for an application that expects multiple backup points acting as mirror images. With minor modifications, the protocol can be optimized to implement data reliability measures in queries, by ensuring multiple copies of the data are received and corroborated by the entity interacting with the interface node.

B. The Search Process

Our system consists of mini-UAVs that are GPS enabled, therefore know each UAV knows its own coordinates. Each mini-UAV is considered a node in the network, and there are n nodes in the network. The number of features to be searched for are grouped into k classes, with the constraints that $n > k \ln k + 3\sqrt{2}k$ [23], and that the amount of RAM on a node is greater than the memory required to store the configuration files and a vector of matching images of interest. A flash storage module will be incorporated to provide additional storage to save encountered images that are interesting. Each node is configured with a folder called `publicdata`, which contains information that is to be transmitted to all other devices in the network. The Haar



Fig. 3. *Detection of multiple features.*

cascade templates for each of the k address classes along with the noise threshold for the class, are then injected into the network, resulting in each node being configured to detect a certain type of feature. An example of the distribution attained is shown in Figure 2. The number below each node represents the backup address selected, and the outermost of the concentric rings denote the radius of communication of the node that has selected backup address 2 in the upper left corner of the diagram.

To implement the application described, cascades for appropriate templates are prepared[11]. The cascades are then grouped into k groups, with the constraint that n , the number of nodes, is related to k such that $n > k \ln k + 3\sqrt{2}k$. `init` is used to determine the number of address classes, and to trigger the selection of backup addresses by nodes. A series of k `store(a, v)` operations are then used to copy each set of templates into the network.

On encountering an image, a node will test for an object of interest by applying the cascades loaded in its RAM. Any node that detects an object of interest will broadcast a local detection signal to the nodes around it indicating its success. The cascades applied test for the human frame, as well as signs of civilization such as cars and houses. In selecting the cascades that are going to be applied, there are two main constraints. The first is the RAM available in the machine, relative to the size of the cascade files. Our implementation needs to operate with an available RAM space of 512 MB, and our cascades are approximately 5 MB. The average size of frames extracted from video streams to be processed is 200 KB. If the system is able to extract 60 frames per second from the video stream it encounters and is configured to buffer approximately 40 seconds of video, then the memory capacity of the device will become exhausted. These space constraints dictate that we set the maximum number of cascades is three on our platform. Another constraint is the processing time per image that a node requires during its analysis. A very powerful processing device with sufficient memory may be able to analyze an image for multiple features in a reasonable time, giving a

result similar to that shown in figure 3.

Sensor nodes to be remotely deployed may be equipped with limited processing facilities to limit the power usage and cooling requirements. Our approach reduces the required processing by prescribing that each node detects a small number of features, giving a detection result similar to that shown in figure 4, and coordinates with other nodes that are testing for a different set of features to determine the number of interesting features in a location.

In the application described, `init` is used to determine the number of address classes, and to trigger the selection of backup addresses by nodes. `store(a, v)` is used to place the initially load data into the network by invoking a series of `store` commands, one for each backup address used, where $a_0 \dots a_{k-1}$ are the backup addresses selected by the nodes, and $v_0 \dots v_{k-1}$ are the templates that are to be placed on the nodes. A user with authority may also transmit through the network for each set of features, a detection threshold to indicate how precise detections are expected to be, as well as a correlation threshold, that tells the number of features in a local neighbourhood that will be considered interesting.

Nodes that experience a local detection signal save an image of the detection watermarked with the GPS location, MAC address and timestamp on the node, and broadcast the detection event along with the backup address to immediate neighbours. Eventually nodes in a neighbourhood that contains multiple features populate a vector of detections, effectively implementing a bag of features search. If the number of detections in the vector on a node surpasses the correlation threshold, the node broadcasts a `store` message that contains the detected data. The `store` message will be propagated through the network. A node that receives a global update signal will retransmit the broadcast in its own time, using index coding to manage the dissemination of information. Receiving nodes that have detected images will re-transmit after a random delay that is greater than time delay configured to prevent broadcast storms. Any receiving node receiving data for its address class will add the image to its `publicdata` folder and retransmit any newly added content to the folder. Over time, all images that have the same address class should have the same data in their `publicdata` folder. An incident commander will then be able to query the network using `get` messages. Over time, the detection and correlation thresholds can be adjusted to finetune the granularity of detections.

C. Data Reliability

Our storage scheme is robust to data loss, and therefore provides data reliability as discussed in [23]. Let $r = \frac{n}{k}$ and let d be the number of nodes are destroyed in a system of n nodes, and k backup addresses. Assuming that address coding is applied, the probability of data loss is given by

$$P(n, r, d) = \sum_{i=1}^k \frac{(-1)^{i+1} \binom{k}{i} \binom{n-ir}{d-ir}}{\binom{n}{d}} \quad (2)$$



Fig. 4. Single detection of features.

A visualization of the probability of data loss events given 300 nodes while varying the number of address classes is given in Figure 6.

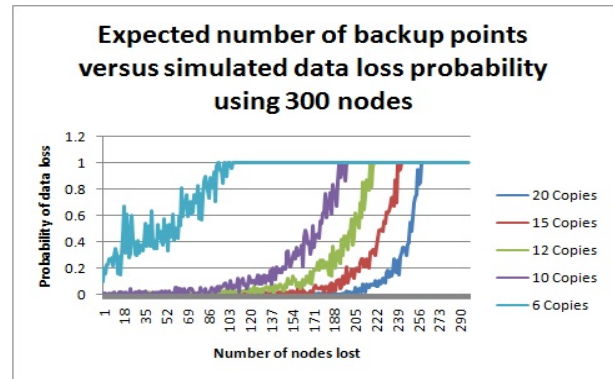


Fig. 5. Probability of data loss in 300 address coded nodes.

We therefore see that our random backup address selection strategy will provide high reliability with relatively low levels of redundancy.

D. Latency Management

In the current configuration, the protocol is designed to operate as a distributed cache, and therefore a query for data at address a issued at a node p propagates to its neighbours, and therefore the closest node to p that belongs to address class a will be the first to respond, propagating the data back to p .

Let A be the area of the network, R be the average range of communication for a one-hop neighbourhood, and assume that the nodes are uniformly distributed over area A . Under those definitions, the average neighbourhood size, δ , is given by $\delta = \frac{n}{A} \cdot \pi R^2$. The area covered by h hops is $\frac{n}{A} \cdot \pi (hR)^2$, and the expected number of nodes in that area is $\frac{n}{A} \cdot \pi (hR)^2 = \frac{n}{A} \cdot \pi R^2 \cdot h^2 = \delta h^2$.

When a search is invoked at node p for a node with address a , the search will propagate through the network as an expanding ring search starting at p . Let m be the number of nodes that are within a distance R from p . The probability that none of those m nodes have selected address

a is $\left(\frac{k-1}{k}\right)^m$, meaning the probability that address a was found is therefore $1 - \left(\frac{k-1}{k}\right)^m$. The probability that address a was found after h hops is therefore the probability that the search for a failed on the first $h-1$ hops, but was successful on the h^{th} hop. In the first $h-1$ hops, the number of nodes encountered is $\delta(h-1)^2$. In the h^{th} hop, the number of nodes encountered is $\delta h^2 - \delta(h-1)^2 = \delta(2h-1)$. be a random variable representing number of hops needed to find a node with address class a . The probability mass function for H is

$$P[H = h] = \left(\left(\frac{k-1}{k} \right)^{\delta(h-1)^2} \right) \left(1 - \left(\frac{k-1}{k} \right)^{\delta(2h-1)^2} \right) \quad (3)$$

IV. RESULTS AND DISCUSSION

Simulations were performed that involved Java based programs processing UAV video footage of wilderness type environments, to check the possibility of successfully locating objects of interest, as well as the likelihood of the generation of false positives or false negatives. A comparison of the time to execute the process on a node for a single cascade versus a series of cascades was also performed. In our simulation we performed feature detection on e. The probability of false positives or negatives were directly related to the images and the cascades presented. The amount of time taken to process an image will is related to the seems to be related to the content of the image rather than the size of the image or depgth of resolution. Images with a high number of undulations such as a forest take almost a second to process a single cascade on a 1.8GHz processor, and the processing of multiple cascades will result in a multiplication of the time required to process the image. In such an environment, the value of parallelizing the processing across multiple processors becomes obvious. Added time spent processing video frames mean that less frames can be processed, and processing less frames increases the risk that frames that contain useful data are not selected for processing. In addition, our simulations applied processing to




Fig. 6. Processing throughput of frames with varying number of cascades.

pre-existing video streams from UAV's but we visualise that a swarm would generate multiple overlapping streams. The net effect that we envision is that the fusion of the resulting video streams could be interpreted as a video sensing system with a very large aperture, giving a deployment similar to a Very Large Array of telescopes.

V. CONCLUSION AND FUTURE WORK

We have shown that it is possible to implement a workable distributed search process using a self organizing data storage protocol over a set of independent UAVs operating as sensor nodes. Assuming that the relationship between the number of nodes and the number of backup addresses is properly set, data robustness and bounded latency in queries are to be expected. The deployment of a physical manifestation of the system is realistic. Additional features may also be incorporated, such as closer inspection of a target by a group of nodes when an object of interest is identified, and the application of interferometry.

REFERENCES

- [1] L. Lin et. al, Supporting Wilderness Search and Rescue with Integrated Intelligence: Autonomy and Information at the Right Time and the Right Place, AAAI, 2010
- [2] Raihan et. al, Raspberry Pi Image Processing Based Economical Automated Toll System , Global Journal of Researches in Engineering Electrical and Electronics Engineering, Volume 13 Issue 13, 2013
- [3] P. Voila, M. Jones, Rapid Object Detection using a Boosted Cascade of Simple Features, IEEE CVPR, 2001
- [4] S. O'Hara, B. Draper, Introduction to the Bag of Features Paradigm for Image Classification and Retrieval, CoRR, 2011
- [5] J. Adams, Camera-Equipped Mini UAVs for Wilderness Search Support: Task Analysis and Lessons from Field Trials, BYUHDMI TECHNICAL REPORT, 2007
- [6] M. Goodrich et. al., Using a Mini-UAV to Support Wilderness Search and Rescue: Practices for Human-Robot Teaming, IEEE SSRR, 2007
- [7] S. Waharte, et. al., Coordinated Search with a Swarm of UAVs, IEEE SECON, 2009.
- [8] W. Lee, Swarm Based Implementation of a Virtual Distributed Database System in a Sensor Network, Air Force Institute of Technology, 2012
- [9] Yao, Gehrke, The Cougar Approach to In-network Query Processing in Sensor , ACM SIGMOD, 2002
- [10] M. Ilyas, I. Magoub, Handbook of Sensor Networks, Compact Wired and Wireless Sensing Systems, CRC Press, 2005
- [11] I. Culjak et. al., A brief introduction to OpenCV, MIPRO, 2012.
- [12] G. Lebanon, et. al., The Locally Weighted Bag of Words Framework for Document Representation, ACM Journal of Machine Learning Research, 2007.
- [13] N. Bhatnagar, et. al., Energy-Reliability Tradeoffs in Sensor Network Storage, HotEmNets08 (ACM), 2008
- [14] J. Beal, Persistent Nodes for Reliable Memory in Geographically Local Networks, MIT, 2003
- [15] A. Kamra et. al., Growth codes: Maximizing Sensor Network Data Persistence, SIGCOMM, 2006
- [16] T. Schwarz and W. Burkhard, Reliability and Performance of RAIDs, IEEE Transactions on Magnetics, 1995
- [17] N Siegmund et. al., Towards Robust Data Storage in Wireless Sensor Networks, IETE, 2009.
- [18] G. Mathur, Ultra Low Power Data Storage for Sensor Networks, IPSN, 2006.
- [19] S. Madden et. al., TinyDB: An Acquisitional Query Processing System for Sensor Networks, ACM SIGMOD, 2005
- [20] Maximizing Data Reliability in Wireless Sensor Networks, Millenial.net whitepaper, Millenial.net, 2005
- [21] Bar-Yossef et. al., Index Coding with Side Information, IEEE FOCS, 2006
- [22] A. Doumas, V. Papanicolaou, The Coupon Collector's Problem Revisited: Asymptotics of the Variance, Advances in Applied Probability, 2012
- [23] P. Gaynor, D. Coore, Towards distributed face recognition on self organized storage built from wireless sensor nodes , WTS 2013
- [24] K. Steliga et. al., On Markov-type inequalities , International Journal of Pure and Applied Mathematics, 2010