Aerial Video Processing Project

Final Report

Sponsor

Dorotha Biernesser

Alex Johnson, Capt

George Mason University

SYST ENG 699

Michael Gibb

Brian Jones

Mingming Lei

Minh Truong

This page intentionally left blank.

Table of Contents

Executive Summary 4
1.0 Introduction 5
- **1.1 Stakeholders** 6
- **1.2 Problem Statement** 7
- **1.3 Project Scope** 7
  - **1.3.1 In Scope** 8
  - **1.3.2 Out of Scope** 9
- **1.4 Assumptions** 10
- **1.5 Literature Search** 11

2.0 Strategy 12
- **2.1 Approach** 12
- **2.2 Objectives** 13
- **2.3 Outline System Design** 13
- **2.4 Developing CONOPS** 14

3.0 Deliverables 15
- **3.1 CONOPS document** 15
- **3.2 Feasibility Study document** 16
- **3.3 Website** 17
- **3.4 Final Report** 17

4.0 Project Management 18
- **4.1 Risk Mitigation Strategy** 18
- **4.2 Challenges** 20
- **4.3 Success Criteria** 21
- **4.4 Conclusions** 22
- **4.5 Next Steps** 24

References 26
Appendices .... Error! Bookmark not defined.
Schedule .... Error! Bookmark not defined.

Executive Summary

This project explores the feasibility of three different algorithm methods and an additional combination of algorithms useable for analyzing video feeds for the purpose of supporting post-disaster response. Following natural disasters, the Air Force's current operational scenario is to load up to 12 hours of Unmanned Aerial Vehicle (UAV) video footage onto a server and have an analyst watch that video for up to 12 hours at a time. This leads to fatigue in the human analyst and results in poor performance in detecting structural damage on the ground following disasters such as hurricane, tornados, floods, and earthquakes. The three algorithms developed for this project offer different approaches to use computer programs to analyze the video and alert the user when anomalies are detected. Anomalies refer to irregularities or differences in the video. Some methods use pre-disaster video to compare before and after results while others use ontology to classify groups of structures. This document explores the many facets of problem solving that went into finishing this project in the single semester time frame. The problem approach, challenges, risk mitigation, and final conclusion are all included to conclude the results derived from the feasibility study and to reflect the realizations that were concluded from the completion of the Concept of Operations (CONOPS) document.

1.0 Introduction

In 1965, the cofounder of Intel, Gordon Moore, predicted that the number of transistors on integrated circuits would double approximately every two years, which in fact turned out to be approximately every 18 months. During this same period of time, the space per unit cost of hard drive memory doubled even faster at every 14 months. The price of one Gigabyte of storage shrank from $300,000 in 1971, to $.10 in 2010. In the 20th century, the pace of Gordon's law has contributed the two necessary ingredients, processor speed and cheap data storage, for what will be the 21st century's primary technological horizon – Big Data.

The pace of development in the sensory industry has reflected the pace of technology in general; and with the burgeoning demand for data that these systems can deliver presents the daunting task of how to handle it. The struggle to process, analyze, and most importantly glean an enhanced insight from massive amounts of "Big Data" is a commonality between many diverse industries in developed countries around the world. Every day, 2.5 quintillion bytes of data is accumulated from everywhere; climate sensors, posts to social media, digital media, purchase transaction records, cell phone GPS coordinates just to name a few. In fact, 90% of the data that exists today has been created in the last two years alone. This leads us to the conclusion that optimal "Big Data" analytics will be one of the primary obstacles that humanity will try to overcome in the 21st century.

The Air Force's current solution is to have UAVs gather video up to 12 hours in length, which will be uploaded, to a Remotely Operated Video Enhancement Receiver (ROVER) system. This video is watched by human analysts at the UAVs home base and checked for areas of interest. If anything of significance is found, the human analyst follows the government protocols to declassify the video segment and then follows procedures to send it up the chain of command who then distributes it to those that need it such as the Red Cross and the Federal Emergency Management Agency (FEMA). The main part of our project's endeavor is to examine a couple ways to bypass the AF's current practice, by designing a versatile video system designed to process up to 12 hours of video feed, automate identification of anomalies at near real-time, and rapidly transfer video down to human analysts on the ground for more precise and thorough analysis of system identified anomalies.

1.1 Stakeholders

The primary stakeholders will be the Secretary of the Air Force's Acquisition (SAF/AQ) department. Mrs. Dorotha Biernesser is the program monitor element for this department and has expressed the interest for the team to complete the project. Mrs. Biernesser plans to share the analysis with the A2C department which oversees short range planning for the Air Force. A system that automates video processing benefits the Air Force by speeding up the time to process video data and freeing up analysts to work on more pressing work. Both of these parties have a vested interest in the details and satisfactory completion of the project.

Additionally, disaster response and disaster relief groups such as the Red Cross, The Department of Homeland Security, and the Federal Emergency Management agency are stakeholders in the project. These individuals will depend on the information provided by the final system when it is deployed to provide accurate and timely processing of video footage to develop disaster response and relief strategies.

The Systems Engineering and Operations Research (SEOR) department at George Mason University (GMU) is a stakeholder in the project because they act like the umbrella corporation overseeing our small team's project. GMU has a fine tradition of delivering quality-engineering products to customers and project sponsors alike. Their stake in the project will be to ensure reliable and intelligent systems engineering practices have been executed on the project to provide a feasible and well thought out solution for the sponsor to uphold the reputation of the department and university.

1.2 Problem Statement

The United States Air Force (USAF) needs a faster way to analyze video feeds to support post-disaster relief efforts after natural disasters such as hurricanes, tornadoes, floods, and earthquakes. Quicker and more automated methods are needed to free up human resources and to leverage the technological performance of automated tools to identify anomalies in video footage to aid first responders in giving assistance to those who need it most.

1.3 Project Scope

This project will focus on defining a concept of operations (CONOPS) for a software system to process and analyze twelve hours of video for the purpose of disaster relief. The main deliverable for the project will be a CONOPS document and a minimum of 2 alternatives to recommend for executing the task. The project is only concerned with defining how the software system should work. No recommendations for hardware solutions, network nodes, communication link, or video sensors will be provided under the project scope. The project team plans to deliver a robust CONOPS document to document how the system will function at a high-level. This will allow the team to come up with two feasible alternatives which will go into more details about the flow of data and potential algorithms. The alternatives will be evaluated using a predetermined and sponsor approved criteria model and the best recommendation will be provided.

1.3.1 In Scope

Software alternatives and algorithm development are the main points of focus for this project. Figure 1 below illustrates the system's role in the overall operation lifecycle the Air Force uses. The only area in scope for this project is the data processing element.

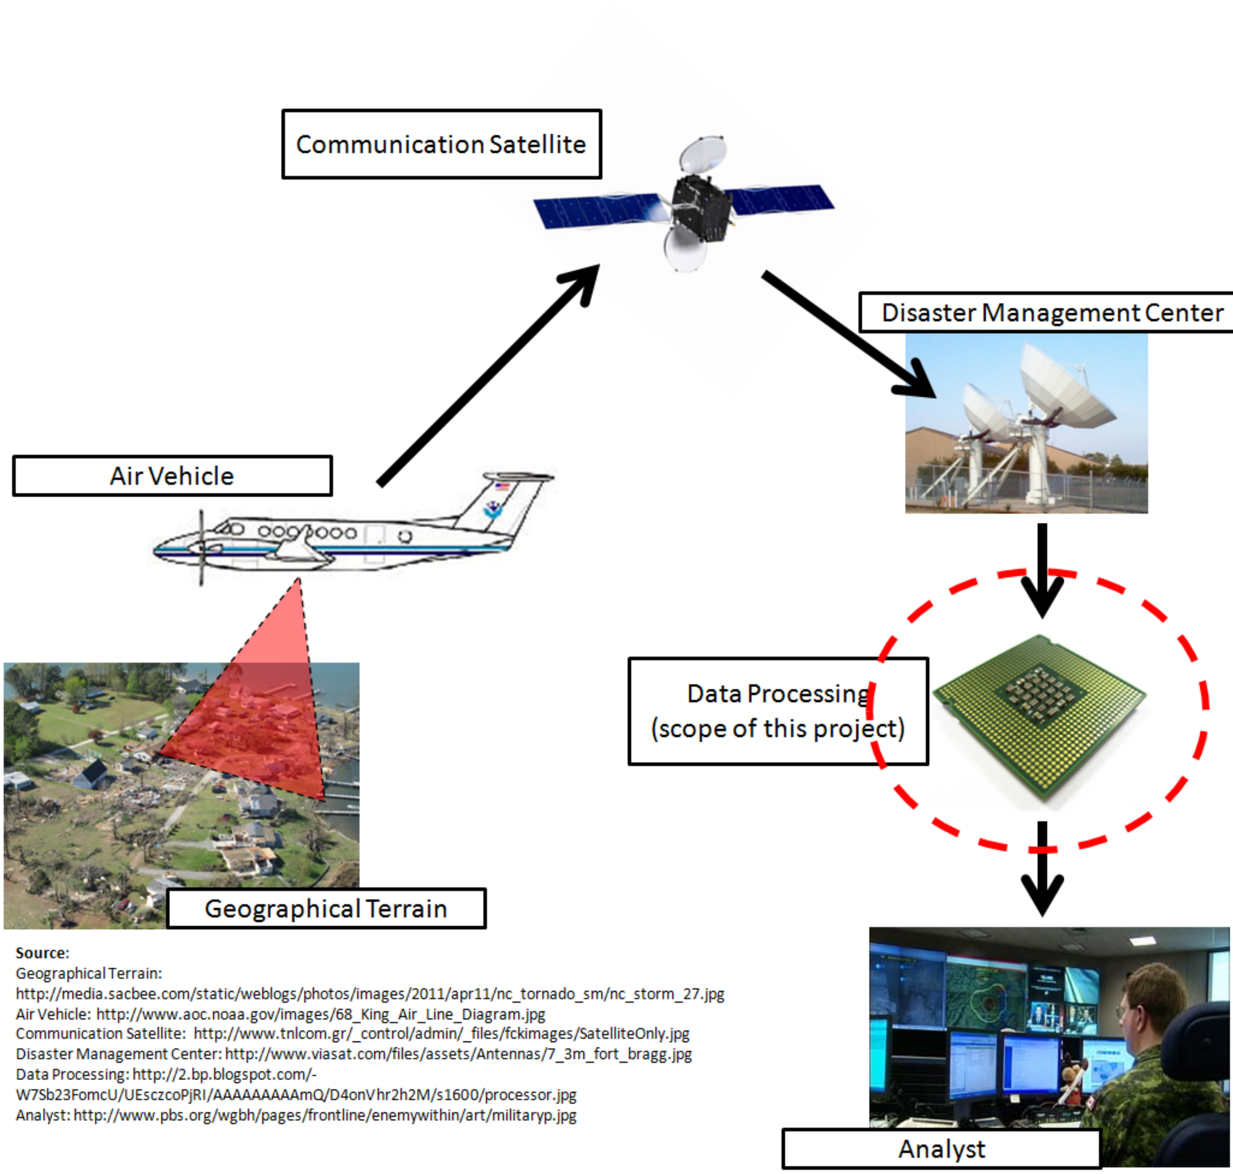

Figure 1. Operational Concept (OpCon)

1.3.2 Out of Scope

The data processing element of the operational envelope is only concerned with software solutions and algorithm development. Hardware components are not to be considered or analyzed because the Air Force has its own classified network of hardware already in use. If they choose to implement any of the algorithm development solutions the project defines, they will assign a hardware engineer to optimize the hardware configuration for their mission. The other aspects of the mission that are not in scope are the data collection mission specifications, relaying of video feeds, or loading of said video feeds onto servers for analysis. These pieces of the project are all proprietary to Air Force operations and are out of scope for this project.

1.4 Assumptions

This project will operate under several assumptions. These assumptions are designed to provide boundaries and guidance for the concept of operations and the analysis of alternatives. The first assumption is that all videos incoming will be following NGA MISB standards. This means the system will not need to account for any unknown video formats or interfaces. We will assume that we are processing twelve hours of data at a given time. This assumption means that the video is being loaded after it has been collected and that in loading the video, users can provide additional inputs to the system to provide more information on the video itself. It also means sometimes the videos will be after-the-fact.

It is assumed that any before-disaster videos taken of an area are of the same format and taken at the same altitude as the post-disaster videos. This means that pre and post imagery should align without additional processing.

We will also assume that humans will be viewing the results of the video processing software and will need to validate the identified video anomalies. This also means that users will not likely rewatch the videos to catch any anomalies that the software has missed. We will also assume the technology exists for the software components and algorithms to function as intended and that the USAF possesses or will possess the necessary hardware to provide the needed computational power.

1.5 Literature Search

To start the feasibility study into possible algorithms, the group needed to become more acquainted with algorithms that applied to image and video comparison. None of the team members came to the project as subject matter experts (SME) in the field of software design or algorithm applications for image comparison. As a team, the group members first decided the direction in which the project would take. After receiving initial guidance from the sponsor, the group was able to obtain the scope and boundaries of the project. The sponsor indicated that she primarily cared with HOW the system could theoretically work. The group had team meetings in which brainstormed sessions were held to determine how this problem could be solved and if there were any other real world applications that were relatable. A common operational medium was found by comparing war time operation practices of UAVs surveying territory. Current USAF UAV systems have the technology to fly over terrain in one pass, then fly over again and overlay the images to see what has changed on the ground. This lets military leaders see enemy troop movements, vehicle movements, and even tire tracks through sand and mud. This realization of current technology led the group to research image comparison algorithms to determine how these processes work. Each group member was tasked with researching at least one feasible algorithm and becoming the resident expert on it to contribute to the group. Literature searches into scientific journals were the main focal point of references. Searches included image comparison algorithm as well as disaster relief applications to determine if someone else had already tried the method and determining whether or not it was successful. Multiple algorithms were discovered for comparing images and video frames but the final algorithms chosen for the project were chosen by the group because of the amount of research that existed for them. More research and references for a specified algorithm allowed the team to gain a vast knowledge base about the algorithm, which would help understand the total picture, and help defend against any questions into its feasibility. The literature search continued until the team had the final algorithms chosen and detailed knowledge was known about each one.

2.0 Strategy

2.1 Approach

The team selected a technical approach that allows the team to achieve the objectives for the development and evaluation of the feasibility of a system capable of processing full motion video. The approach is composed of three parts:

1) Research historic studies on various image analysis algorithms that are relevant to this study
2) Identify key metrics to evaluate the performance of each alternative
3) Develop a high order conceptual design of operations that outlines the tests, scope, strategy, and evaluation criteria.

2.2 Objectives

The important objectives for this project effort as it relates to the evaluation of alternative video processing algorithms:

- Derive requirements that address the sponsor's problem specific to processing up to 12 hours or more of full motion video feeds
- Document overall system design and engineer comprehensive feasibility study
- Compile recommendations as part of a report to the sponsor
- Deliver the final project report and present project details and findings to the sponsor and GMU faculty

Although these objectives can address the sponsor's difficulty, a comprehensive investigation at the subsystem level or lower is required to fully appraise these cumulative means. Given the nature of this project, evidence regarding low-level mechanisms is routinely classified or proprietary; assumptions that best satisfy areas where information has not been made available are described in the sections that follow.

2.3 Outline System Design

The operational approach to detecting areas-of-interests, such as structural damages, can be categorized into four stages:

1. **Input stage**: when the system receives aerial video footage and the system gathers image/video data from other sources corresponding the geographical areas being assessed.
2. **Processing stage**: when the system analyzes the data collected.
3. **Pairing meta-data stage**: when the system pairs the meta-data corresponding to what the processing stage outputs as the areas-of-interest.
4. **Alert stage**: when the system alerts an operator of the areas-of-interest.

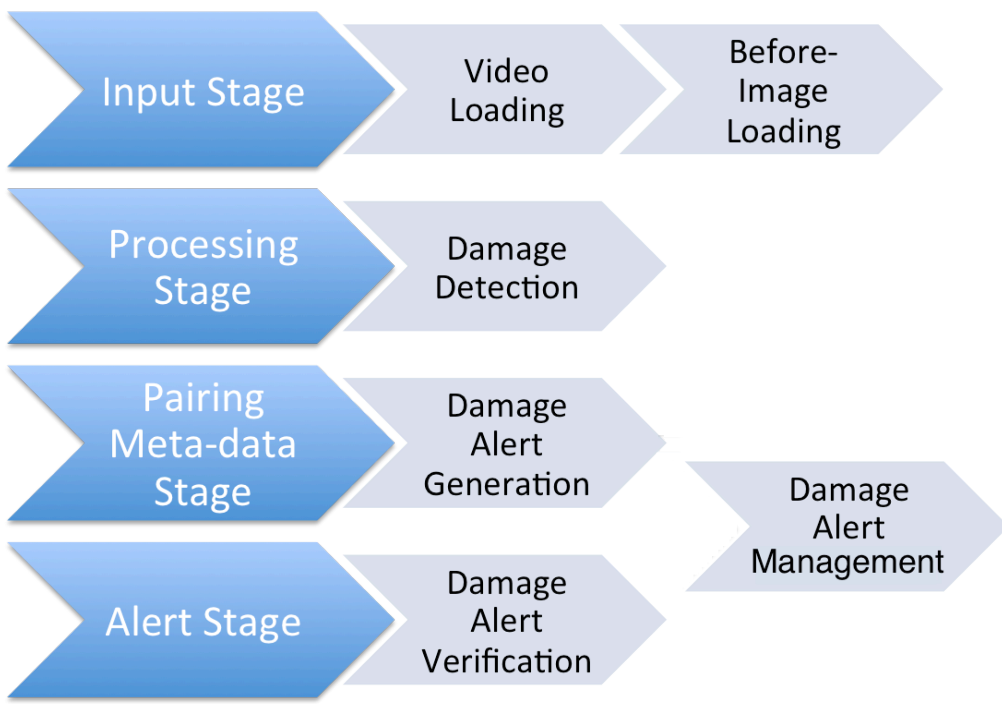

Figure 2. Correlation of Stages to Functional Capabilities

2.4 Developing CONOPS

This piece of our examination displayed our conceptual understanding of the system design. After decomposing the system and identifying components, we conducted a thorough analysis of historical findings that provided several alternatives and comparative performance outcomes of various designs of video processing algorithms.

3.0 Deliverables

3.1 CONOPS document

The Concept of Operations document is the culmination of the work performed by the team. It discusses in depth the background of the project, the need for the capabilities of this project, the operations the potential system would perform, and the functional capabilities required that the system perform.

The background covers the high-level concepts associated with this project including what the problem is, what the scope of the project is, and an overview of the stakeholders. The need section covers what the business needs are what that gap is currently. The operations section begins the core of the work and research put into this project. It covers the missions that are occurring that this system would be part of. It discusses some of the assumptions and constraints that the system will face in operations. It gives detail into what the operational concept would look like and what environment the system would operate in. It also delves into the image processing algorithms and how they would work in the operations as different configurations to consider. The final section covers functional capabilities of the system. The key functions the system will provide to the operations is somewhere to load videos, gather prior-imagery if applicable to compare to the loaded video, detect anomalies in the video, generate alerts for all the anomalies detected and verify the accuracy of the detections, and manage the videos and the anomalies associated.

3.2 Feasibility Study document

The feasibility study document is integrated into the Concept of Operations document. It takes the concepts embedded into the ConOps and delves deeper into evaluating the feasibility of the system described in the ConOps. The feasibility study covers the evaluation criteria used to evaluate the system, the operation scenarios the system would function under, recommendations for system designs for each scenario, and finally recommendations for future work followed with a conclusion.

There are three key evaluation criteria that the group decided to use: cost, speed, and accuracy of the system. We use the existing system's metrics as a baseline for what a new system must meet. The two operational scenarios we came up with are what we term mitigatable disasters and unmitigatable disasters. The key difference between the two scenarios is that in one case, we know the disaster will occur. In the non-mitigatable case, there is no warning prior to the disaster. For mitigatable disasters, algorithms that compared prior imagery to the loaded footage could be used. In non-mitigatable disasters we recommended that algorithms that did not require prior imagery be used. The group thought the greatest limitation of the system was the processing time if done linearly; however as hardware was scoped out of the problem, parallel computing was not considered. Therefore, we recommended future research focus into looking at improving processing time of the image processing algorithms.

3.3 Website

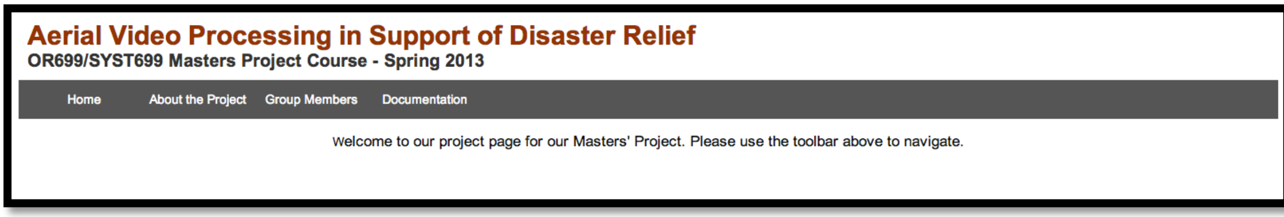

Figure 3. Screen shot of Website

Our website can be found at: http://mason.gmu.edu/~ylei/syst699/index.html. The website is the outward-facing representation of the work that we've performed this semester. The design was kept simple to make it easier to use and speed up loading times. The website consist of four main sections: a welcome page, a project description page, a page about the team members, and a page containing all the work turned in for the project.

The welcome page is basic and simply provides a place to enter the site on. The project description page gives a high-level overview of what the problem is and how our project addresses this problem. The page on the team member provides the names of who worked on this project and also what our backgrounds are. The final page just provides links to all the work done for this project to include the proposal, the concept of operations and feasibility study documents, the final presentation and also the final report.

3.4 Final Report

The final report wraps up the programmatics of this project. It discusses the process and approach the group took to put together a product. It covers the major components of the system that need to be accounted for. It also discusses the deliverables that were put together for this project and the work that went into them. The final report additionally covers our schedule and the risk management plan we put together.

All members of the team worked on putting together the final report. It contains information on how the project was managed and laid out and what work done was.

4.0 Project Management

4.1 Risk Mitigation Strategy

Project risks were primarily focused around the literature search. The literature search for the project would make or break the final product because of the necessary information needed to be obtained from literature references. Since the project was primarily a conceptual design with an accompanying feasibility study, the literature search into how to make the system would was crucial to overall project success. In order to mitigate this, the group focused a large amount of time and effort on researching algorithms. Only the algorithms with the most comprehensive research made the final cut and were included as feasible solutions to the problem. One major risk that was unable to be mitigated was the lack of access to proprietary and industry best practices for these algorithm applications. As explained in the literature section above, it is well known that the military currently uses image comparison technology to compare images of battlefields obtained from flying platforms. Therefore there is very likely classified or proprietary documentation detailing the use and design of these systems. Since the group knew it would not be able to access this information, it was noted as an unmitigatable risk that we may not recommend the optimal solution since we did not have access to all research from the field. Other risks for the project included the schedule risk from only having one semester to complete the project. This risk was mitigated by properly bounding the problem up front to scope it down into a manageable problem that could be solved within the one semester time frame.

As a group, the team faced challenges and risks in completed the project related to group dynamics and scheduling. These problems are inherent to any team collaboration and were well mitigated through discussions, coordination, and shared responsibilities. The risk of not completing the project on time was mitigated by having group meetings twice a week. The group met on Mondays during regular class times and once one weekends. The biweekly meetings helped to ensure group members were meeting their deadlines and were able to ask questions before deliverables were due. The group recognized that team members were not solely focused on school. All team members were working professionals with full time jobs and all had other graduate classes to worry about. To mitigate the risk of one person not completing an assignment on time, important tasks were often assigned to more than one individual. This ensured that there was a shared responsibility for tasks and that assignment would likely never be completely overlooked. Meeting minutes and follow up emails were also sent out after every meeting to ensure that every team member was working towards the same goal. Meeting minutes allowed the team to hold each other accountable for what was discussed and what was expected to be completed by the next team meeting. File sharing websites and online collaboration tools such as Google Drive were also implemented to allow team members to simultaneously collaborate on documents. On line repositories helped the team work online at one time while also allowing for configuration management and version control of project documents. To help the team mitigate the risk of deviating from the project path, the sponsor was consulted with prior to every major deliverable to make sure she was content with the team's progress.

4.2 Challenges

The challenges consisted of two types: challenges centered around the project and challenges centered around the group itself. The first major challenge related to the project was scoping it down to be a manageable yet challenging problem capable of being solved in one semester. This challenge was overcome by talking with the sponsor and coming to an agreement as a team as to the direction the project would take. The final decision was to complete a high-level conceptual design and feasibility study. Other project related challenges revolved around the algorithms themselves and the research necessary to complete the conceptual design. Most algorithm research required multiple documents to pull all the pieces together. There was not a single algorithm where all the specifications and limitations of the algorithm were pulled from one source. This made it challenging to complete the conceptual design because the team needed certain pieces of information to ensure that algorithm would in fact work the way we said it would. This challenge required the team member primarily responsible for that algorithm to organize their sources and share with the team so all the team members thoroughly understood the recommendation and could defend the algorithm's feasibility during the final presentation.

Group related challenges lasted the entire semester and were well mitigated through communication. The challenges were expected because of usual group dynamic paradoxes and the realization that every team member had other things going on including work and other classes. The first major group challenge was team meeting times. It was determined early on that each team member had a different class schedule and meetings besides Monday class times during the week were improbable. At least one team member was missing Tuesday through Thursday for other graduate classes. The second team meeting was a relatively easy challenge to overcome by scheduling meetings early on Saturday or Sunday mornings. This allowed team members to have the team meeting and then carry on with their other weekend plans. As the weather changed this team meeting moved to Friday nights because of other activities. Team members had family obligations and schedule events so the challenge of meeting on the weekends increased. Friday night meetings also posed a problem for some with other obligations. This challenge meant that whoever missed the meeting would not have the ability to brief the team on their activities for the week and would miss the chance to pick their tasks for the following week. This challenge was overcome by group members proactively emailing scheduling conflicts and taking responsibility for assignments they were given if they missed the meeting. This challenge was expected since it was understood that all team members had jobs, other classes, and families to attend to.

4.3 Success Criteria

This project will be successful if the following can be accomplished:

- Sponsor accepts the in-depth CONOPS based on the criteria that it meets all the baseline requirements and provides the information needed.
    - Identify at least two viable alternatives how the system processes aerial video data.
    - Identify at least two methods how the system alerts an operator.
- Sponsor accepts the Feasibility Study of the at least two viable alternatives.
- Dr. Barry and/or Dr. Loerch accept the Final Report.
- Dr. Barry and/or Dr.Loerch accept the Final Presentation.
- Create a website that is accessible and contains the details of this project.

4.4 Conclusions

This project began as a research study into the various alternatives for automatically processing aerial video images as a mean for detecting damages. The algorithms and operational process outlined in the CONOPS fulfill the project requirements. That is, we have identified approaches as to how the system ingest video data, process the data, and generate alerts for an analyst to review. Conceptually, the approaches make sense, but in reality, we must consider how the system will perform when implemented with real-world limitations - accuracy, speed and cost.

There is a need for the system to have some form of a reference image to compare and determine whether the video frame being analyzed contains any anomalies. The accuracy of the system is reliant on the reference images. To have a repertoire of pre-identified images for the system to learn from may be feasible, but to capture and store every types of damage scenario may not be. The other solution is to obtain pre-disaster videos of the geographical area, but doing so would be costly since that would require a massive amount of resources to fly out and capture videos of all areas. Further, some disasters are impossible to predict and they could occur anywhere. The algorithms as identified are sufficient in detecting damaged areas with acceptable levels of accuracy as high as 86%, but the threshold of the acceptable accuracy level relies on the severity of the damage and what is at stake.

The algorithms are, however, not fast enough to surpass human processing. Currently, the ideal standard is for the system to at least process as fast and as accurate as a human analyst. The current system is not feasible for damage assessment that is reliant on the urgency of time. Such situations include assessment for first responder or for emergency rescue. The system is, however, feasible for research studies such as tracking the changes of coastal lines after a tsunami or determining damages for insurance purposes, for example.

The cost of the system is driven by both software and hardware capabilities. Ideally, the cost to build and maintain the system should be less than that of the manpower required to watch the video in a given amount of time. As such, it is not feasible with the current computing technology to build a system and expect it to be within a reasonable price to meet the ideal requirements.

There are many outlets for improvement to the aerial video processing system. More research into sophisticated approach is within reach that can solve the biggest challenge for this project - processing a large amount of data within a reasonable amount of time with an acceptable accuracy. With respect to accuracy, speed and cost, it is inevitable that at least one element is compromised in exchange for superior performances in the other two.

4.5 Next Steps

The USAF needs to invest in a feasibility study of their own to determine if the available technology meets their mission requirements. One possibility is to extend their research into the specific algorithms described in this project. Another, more general approach, is to focus their research on software and hardware capabilities to process big data.

It is recommended that additional research continue with damage detection algorithms. In particular, focus should be on algorithms that can identify non-areas-of-interest. Such capabilities will allow the system to better identify the damaged areas by removing background information such as vegetation or large water mass. This type of logic working in conjunction with damage detection algorithms may serve as a check-and-balance to increase the accuracy of damage detection.

Due to the scope of this project, hardware capabilities were not considered in detail for this damage detection system. Therefore, a recommendation for future research would be in hardware designs to speed up processing power. A potential approach would be to employ parallel computing using redundant hardware for simultaneous video analysis. For example, twelve systems could be running simultaneously and each processing a 1-hour segment of the 12-hour video. This would greatly reduce processing time but may drive up costs.

Crowdsourcing is another approach recommended for future research. A high-level logic may start with the system ingesting the video data and splitting the video into 10-second video segments. After that, the system would distribute the segments to millions of registered users. The users would watch the short video clip and ping back a message of what they see. If multiple messages of a particular video segment include keywords such as "damaged", "collapsed", "broken" and so on, there is a strong indication of a positive damage detection. This entire process may be time consuming, but is highly accurate based on human feedbacks. Needless to say, there will be issues regarding the legality of distributing aerial video to users across the country.

References

1 Uprenda P Poudel, Gongkang Fu, Kian Ye, Structural Damage Detection Using Digital Video Imaging and Wavelet Transformation, Center for Advanced Bridge Engineering, Civil, and Environmental Engineering Department, Wayne State University

2 Daozheng Chen, Mustafa Bilgic, Lise Getoor, David Jacobs, Dynamic Processing Allocation in Video, IEEE Transactions On Pattern Analysis And Machine Intelligence, Vol 33, No 11, Novinger 2011

3 Harold Trease, Daniel Fraser, Rob Farber, Stephen Elbert, Using Transaction Based Parallel Computing to Solve Image Processing and Computational Physics Problems, www.cca08.org

4 Motion Imagery Standards Board, Inserting Time Stamps and Metadata in High Definition Uncompressed Video, MISB STD 0605.3, June 9th 2011

5 Motion Imagery Standards Board, Compression for Infrared Motion Imagery, MISB Standard 0404, December 3rd 2009

6 Motion Imagery Standards Board, MISB Profile for Aerial Surveillance and Photogrammetry Applications (ASPA), MISB Standard 0301.5 Version 1.5, July 26th 2010

7 Ryan Hruska, Greg Lancaster, Jerry Harbour, Shane Cherry, Small UAV-Acquired High Resolution Georeferenced Still Imagery, Idaho National Laboratory, September 2005

8 N. Voisine, S. Dasiopoulou, V. Mezaris, E. Spyroum T. Athanasiadis, I. Kompatsiaris, Y. Avrithis, M. G. Strintzis, Knowledge Assisted Video Analysis Using A Genetic Algorithm, Information Processing Laboratory, Electrical and Computer Engineering Department, Aristotle University of Thessaloniki, Greece.

9 Yuko Roodt, Willem Visser, Willem Clarke, Image Processing on the GPU: Implementing the Canny Edge Detection Algorithm, University of Johannesburg, http://www.prasa.org/proceedings/2007/prasa07-26.pdf

10 Ozisik, D. *Post-earthquake Damage Assessment Using Satellite and Aerial Video Imagery.* Enschede, The Netherlands: International Institute for Geo-information Science and Earth Observation, 2004.

11 Kim, K., Chalidabhongse, T., Harwood, D., & Davis, L. (2005). *Real-Time Imaging*. Retrieved from SciVerse - ScienceDirect: http://www.sciencedirect.com/science/article/pii/S1077201405000057

12 Voisine, N., Dasiopoulou, S., Mezaris, V., Spyrou, E., Athanasiadis, T., Kompatsiaris, I., ... & Strintzis, M. G. (2005, April). Knowledge-assisted video analysis using a genetic algorithm.

In *Proc. of 6th International Workshop on Image Analysis for Multimedia Interactive Services (WIAMIS 2005)*.