2015 Robo-Ops Rover Technical Report

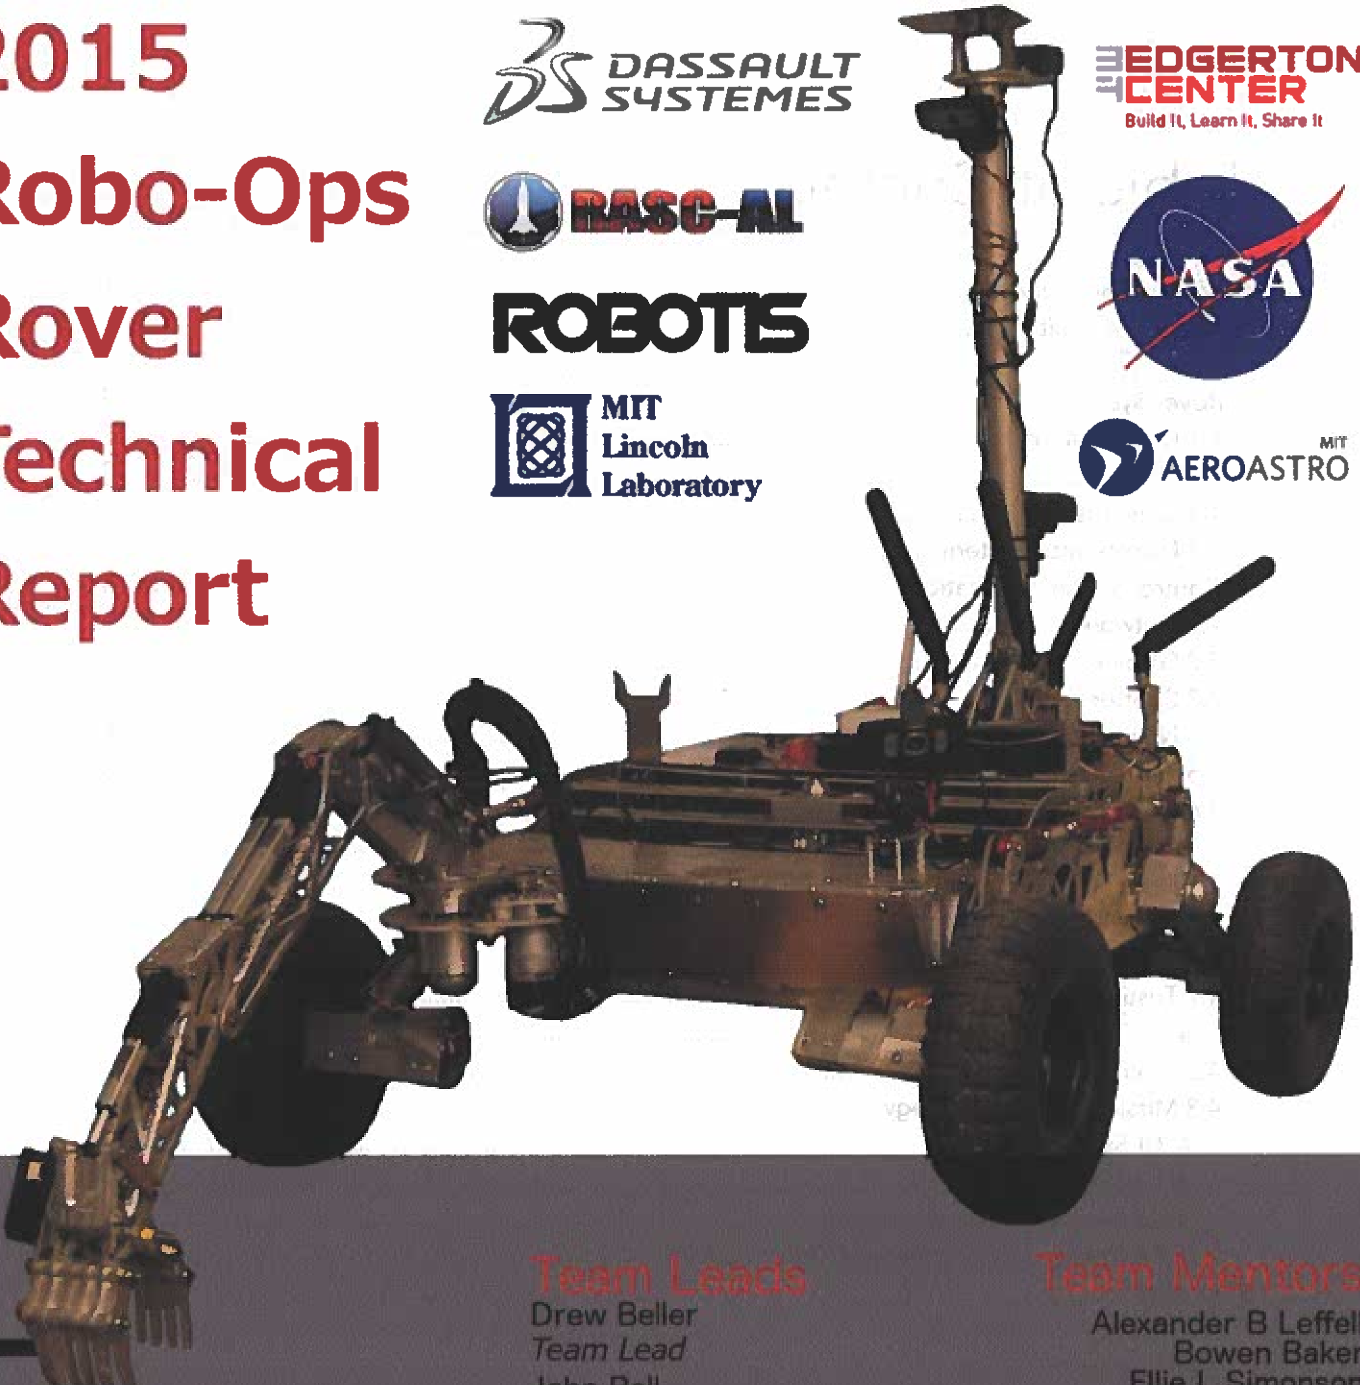

MIT ROBO TEAM

Team Leads

Drew Beller
Team Lead
John Bell
Mechanical Lead
Henry Wu
Software Lead
Diana Dmitrichenko
Electrical Lead

Faculty Advisor

Dr. Olivier de Weck
Professor in the Department of Aeronautics & Astronautics and Engineering Systems

Team Mentors

Alexander B Leffell
Bowen Baker
Ellie L Simonson
Justin J Cheung
Raichelle J Aniceto
Reo Baird
Rita Ainane
Roger D Lo

Team Members

Alice Jin
Aya K Suzuki
Juan A Romero
Kelsey K Chan
Michael A Feffer
Nayoung Lee

Table of Contents

Executive Summary

The following report describes the design and implementation of the MIT Robotics Team's entry in the 2015 RASC-AL/NASA Exploration Robo-Ops competition. The competition is centered around building a planetary rover prototype capable of traversing diverse terrain and collecting rock samples at the Johnson Space Center while being tele-operated from a team's home campus. The main restrictions and requirements for the competition include:

- a maximum size and weight of 1.0m*1.0m*0.5m and 45.0kg respectively
- the ability to navigate obstacles of up to 10cm in height and grades up to 33% (including loose sand)
- the ability to identify and collect colored rock samples ranging in size from 2 to 8 cm and weight from 20 to 150 grams
- the ability to stream back a high quality audio and video stream from the rover

Rover Capabilities

As a second year entrant to the Robo-Ops competition, the MIT Robotics Team seeks to improve upon the design from last year. The main focuses in the new design have been on increasing reliability and performance of last year's systems. Highlights of last year's design include "a flexible multi-pivot averaging suspension system, a robust five degree of freedom arm, a powerful onboard image processing system, and an intuitive piloting interface that takes advantage of the latest real-time video streaming and mobile connection technologies."[1] Key improvements made this year include a more reliable, lighter arm, reduced onboard motor control latency, improved rock identification integrated into Mission Control's GUI, and higher quality video streaming. In addition, the camera system has been upgraded to send back high resolution photos for image processing by Mission Control.

Team Structure

The MIT Robotics Team has grown significantly in the past year, attracting both undergraduate and graduate students from various disciplines. The team has been divided into 3 engineering divisions: electrical, mechanical, and software; the divisions have worked closely together to carry out key improvements to the rover. Graduate students from the MIT Sloan School of Management have also joined the team as members of the business division; they have been instrumental in securing partnerships and sponsorships as well as organizing the team's successful crowdfunding campaign.

---

[1] Source: MIT Robotics Team 2014 Robo-Ops Final Report

1. Rover Systems

1.1 Drive System

The drive system uses a passive averaging suspension system, which allows for the rover to maintain four points of contact over large obstacles and uneven terrain. The passive averaging system also has five pivot joints to distribute load evenly across the chassis. The averaging is controlled by an aluminum center linkage that rotates about an aluminum spindle centered in the chassis. The center linkage is then connected to the four drive legs by helm-joint rods and made adjustable through a connection by male threaded rods. Each drive leg is also connected to an aluminum spindle on the chassis via custom dual-taper-roller bearing assemblies. The bearing assemblies are designed to withstand axial and thrust loads on the drive leg. The suspension system provides great stability and traction; however, it does not reduce much vibration when turning, which greatly affects the rover's video compression algorithm. To combat vibrations when turning or skidding, the rover uses Kyosho Zeal Vibration Absorption Mounting Gel on its camera mounts.

The wheels used on the rover are 10" diameter foam-filled turf tires, which provide excellent traction on a variety of terrains. The tread pattern on the wheels is bidirectional and symmetrical, which is optimal for skid steer control, as the wheels must be able to slide laterally when turning. The team decided to use skid steer control type wheels for ease of turning at high speeds on loose terrain, such as sand or gravel. When operating the arm, skid steering allows the driver to be very precise in aligning the rover, as it allows for zero-radius turning.

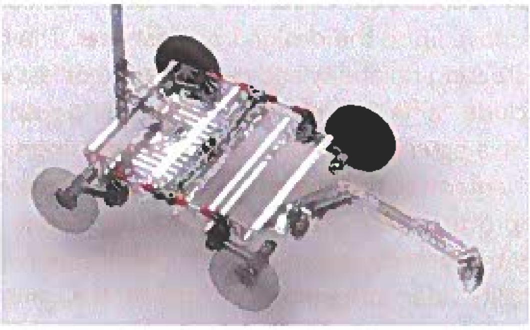
FIGURE 1: DRIVE SYSTEM

1.2 Chassis

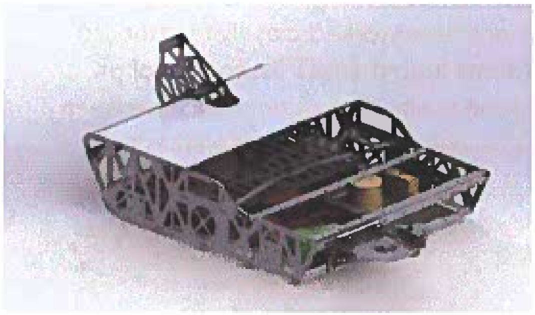

The chassis design is very similar to the design of last year's chassis, but overall weight has been reduced by ~30%. The chassis was made out of tig-welded water-jetted eighth inch aluminum sheet. The final dimensions of the chassis are 89cm x 53cm x 17cm with a final weight of 41 kg.

FIGURE 2: CHASSIS

1.3 Cameras

Four Logitech C930e webcams are used as main driving cameras for the pilot. The C930e webcam is capable of capturing 1080p video at 30fps. On the rover, the cameras run at a resolution of 640 x 360 pixels; this resolution provides excellent quality for driving, ensures a wide view of the terrain, and is fast to compress.

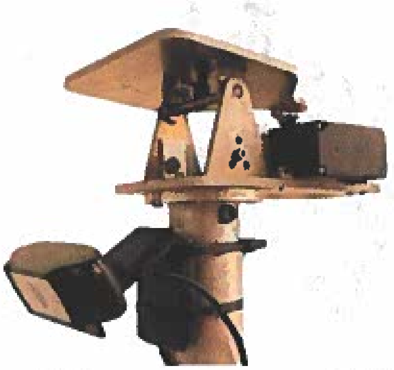
FIGURE 3: GH3 CAMERA MOUNT

The rover also incorporates a folding camera mast, which can be open and closed remotely using two Firgelli 100mm L16 linear actuators. The mast carries a Panasonic Lumix DMC-GH3 camera, which offers much higher resolution than the three Teledyne Dalsa BOA PRO cameras used last year. The maximum height of the GH3 camera is 120cm. The GH3 also provides onboard autofocus and auto-white-balance adjustments as well as zoom control. The camera mount uses a four bar linkage system and a Hitec HS-755MG servo, which provides Mission Control with control of the camera's tilt, ranging from -30° to 15°. The GH3 sends back a low-resolution live preview (640 x 480 pixels) and high-resolution photos (4608 x 2592 pixels).

1.4 Manipulator System

Based off the success of last year's robust arm design, the team decided to continue using the same base design, modeled off of a commercial excavator. It has five degrees of freedom: a base rotation joint, a boom, a stick, a claw, and a thumb. The rotation joint is controlled by a Pololu 131:1 metal gearmotor, which operates at 80 rpm and allows for 360 degrees of rotation. The boom is controlled by a 150 pound force linear RoboTeq actuator and can raise and lower the overall height of the arm by 20 degrees in each direction, thus resulting in a 40 degree range in total. The stick joint is controlled by two 100 mm F16-P Firgelli linear actuators with a total load capacity of 45 pounds. The stick joint is capable of rotating 120 degrees below horizontal. The claw joint is controlled by a 50 mm F16-P Firgelli actuator with a load capacity of 22.5 pounds, and is capable of rotating the claw 80 degrees below the horizontal.

FIGURE 4: THE ARM

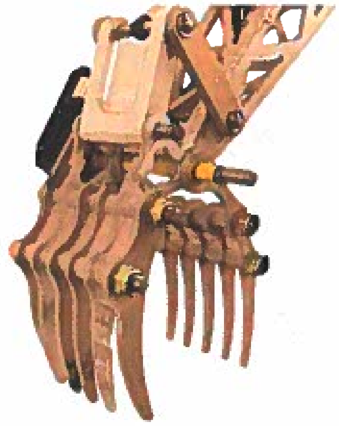

FIGURE 5: THE REDESIGNED CLAW

The main upgrade to the arm this year has been the redesigned claw. The thumb joint can rotate 150 degrees and is controlled by a Dynamixel RX-64 Smart Servo, which has a stall torque of about 50 inch-pounds. The Dynamixel also allows for variable torque control, which reduces the risk of the servo burning out when grasping samples. The claw has been redesigned to be lighter in weight in order to reduce servo stress and have thinner flanges. The thinner flanges make the claw much more effective at securing smaller samples more easily and digging under samples to retrieve them. Each flange is also curved and tapered to guide samples into a secure pocket during retrieval.

When fully extended, the arm has a reach of 70 centimeters and can most effectively pick up samples within a 50 centimeter range. All structural components on the manipulator system were manufactured by the team in-house out of aluminum sheet and stock. The boom stick and claw are made from 3/16 inch and 1/8 inch water jetted aluminum sheet.

2. Control & Communications

2.1 Network

The rover is able to establish high-bandwidth, low-latency communication with Mission Control through the use of Peplink's SpeedFusion technology. Peplink's technology performs packet-level load balancing, hot failover, and bandwidth aggregation on multiple WAN sources. This allows the rover to treat its two broadband wireless connections as a single reliable connection.

The rover will use two 4G broadband connections (through providers AT&T and Verizon) as its WAN sources. The broadband connections will be integrated with the SpeedFusion Technology through the Peplink MAX HD2 Mini modem. This will relay back to a Peplink Balance 580 router at Mission Control, which will re-bond the disassembled packets from the rover.

During the competition, an internal WiFi network will also be set up through a Peplink Surf On-The-Go (OTG) router in order to connect to Panasonic's proprietary control software on the Panasonic Lumix GH3 camera.

A full networking schematic is available in Appendix II: Rover Networking Schematic.

2.2 GH3 GUI

The GH3 GUI controls the GH3 camera at the top of the rover's mast. The GUI gives Mission Control access to the camera's high quality live feed as well as control over the camera's tilt and all of the camera's settings (e.g. focus, white balance, exposure, zoom).

The GH3 operator uses a combination of the live feed and input from the spotter to take pictures of area of interest. These pictures are then sent to multiple spotters for image processing and rock identification.

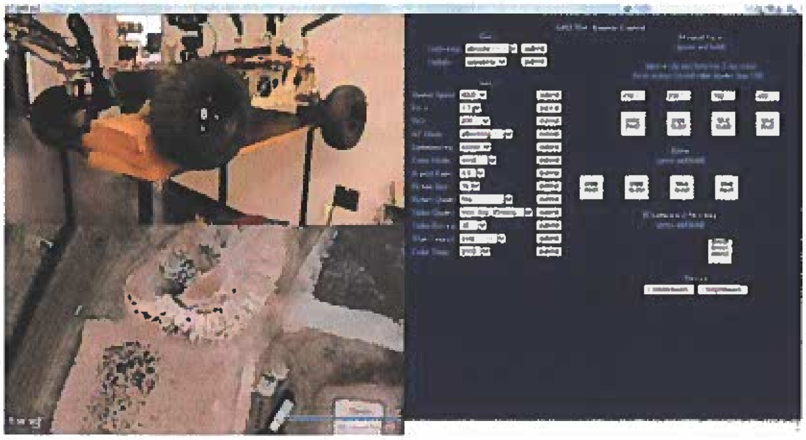

FIGURE 6: GH3 GUI SCREENSHOT

2.3 Spotter GUI

FIGURE 7: SPOTTER GUI SCREENSHOT

The Spotter GUI searches through the high quality photos sent back from the rover. The photos are first processed by Mission Control's computer vision system (Section 2.6), which identifies the location of rocks within each image. The GUI uses this information to mark the location of each rock in the image with a cross. However, false positives and negatives must be avoided; so, computer vision is only used to assist human spotters.

The spotters examine each image and find where the rocks, if any, are in the image. To reduce error, spotters can view a zoomed in version of any area they select. Finally, spotters mark each rock found with a color and an estimated position in the competition field. This information is then sent to the Navigation GUI, which displays it on a large map of the field.

2.4 Navigation GUI

The Navigation GUI determines the path the driver should take. The navigator will decide this based on suspected rock locations from the spotter. The Navigation GUI will track the rover using GPS and implement a traveling salesman algorithm to determine the intended path for the driver to follow. The traveling salesman algorithm will take Mission Control's overall strategy into consideration when planning a path by accounting for variables like time remaining, color of each rock, and the risk of the terrain at each rock location. As the rover approaches an area of interest marked by the spotter, the Navigation GUI brings up the picture that the rock was spotted in, to give the navigator a clear view of the area.

2.5 Pilot GUI

The Pilot GUI is the main interface for the rover driver. The driver works with the navigator to follow the planned path and find the rocks in the areas identified by the spotters. The Pilot GUI displays four camera feeds across two HD monitors for the driver at Mission Control, as well as diagnostic information such as battery level, ping time, roll, pitch, and tilt.

The Pilot GUI offers a variety of control modes to optimize the driver's performance based on the current task. The various driving modes include a tank drive mode for precision, an arcade drive mode for speed, and an arm control mode mapped to the standard SAE (Society of Automotive Engineers) Control pattern for efficiency. This can be seen in Appendix V: SAE Control Chart.[2]

The rover is controlled using two Logitech 3-axis joysticks. The joysticks have multiple buttons used for camera control and driving mode selection. Proportional and exponential gains on all controls can be changed by the driver on the fly during competition to favor either precision or speed when controlling the rover. The Pilot GUI also has button activated autonomous maneuvers for depositing a rock, resetting cameras, stowing the arm, and resetting the arm position, all of which increase driving efficiency.

2.6 Computer Vision System

Mission Control's computer vision system consists of a custom blob detection algorithm. Images are split into RGB components, and a score for the color intensity of each pixel is computed. These values are processed using an adaptable threshold to select the set of points with the greatest color intensity. A specialized flood fill then merges sets of points into blobs, and determines the location and size of each one. Finally, excessively small or large blobs are filtered and rejected as noise.

---

[2] Source: SAE Pattern

2.7 Media Processing

The rover's media processing systems have been greatly overhauled. From an architectural standpoint, last year's encoding software was inefficient and difficult to work with. All video frames were combined into a single image, which was then compressed using ffmpeg and sent over the Internet.

This year, the rover uses a GStreamer backend, which manages each audio and video stream. Each stream automatically pulls data from the rover's cameras and microphones, encodes this data, and sends it over UDP. As an added benefit, streams are automatically multithreaded. This allows the onboard CPU to compress video and audio at higher qualities.

The rover uses one GStreamer stream for each webcam. Each video stream is compressed using the x264 video encoder with the H.264 standard. The x264 compression currently takes less than 10ms per frame. To account for accumulated error in the H.264 output caused by vibration, the x264 encoder refreshes with a much greater frequency than last year, thus clearing away old video data much faster.

To conserve valuable bandwidth, Mission Control's Pilot GUI sends messages to the Streaming Server on the rover indicating which streams are being used and at what resolution. The server then dynamically scales the bitrate of each stream, reducing and increasing each one as necessary. In addition, a "Low Bitrate Mode" can be manually enabled, which further reduces the bitrate of each video stream. This mode allows for more effective operation in areas with poor cellular reception.

GStreamer is also used to stream audio data. Sound is captured from the microphones built into the top mast webcam and processed using a high-pass filter, which reduces vibrational noise and decreases bandwidth usage. The audio is compressed using the Opus audio codec.

2.8 Latency Information

Control latency has been reduced from last year by fully threading the software drivers used to control hardware. All blocking system calls or otherwise expensive operations can be offloaded to a driver thread. As a result, the main execution thread can process events faster and more consistently than before. Qualitatively, the rover arm and drivetrain are significantly more responsive to user control.

Audio will be encoded using the Opus audio codec due to its low latency and high quality operation. The codec is set to 20 milliseconds of internal delay; the resulting end-to-end latency of the audio stream is similar to the latency of the video stream.

Video processing has been improved by using GStreamer to automatically thread both the video encoding and decoding processes. Frames from each video stream can be compressed simultaneously. In the Pilot GUI, streams are decoded simultaneously and independently of the main process.

Over WiFi, the nominal end-to-end latency of the video stream is between 170 and 200 milliseconds. Further testing shows that using a 4G LTE cellular connection results in an additional latency of about 100 milliseconds. Thus, the team expects that the video streams will run with less than 300 milliseconds of delay during the competition.

3. Technical Specifications

TABLE 1: TECHNICAL SPECIFICATIONS

DIMENSIONS	1m x 0.7m x 0.5m
WEIGHT	41 kg
RATED PAYLOAD	45 kg
MAX SPEED	2.3 m/s
MAX OBSTACLE SIZE	20 cm
OPERATING TIME	1.5 hours
DRIVE POWER	~50 W (idle) ~200 W (typical operation) ~1600 W (peak)
BATTERIES	5x 4-cell, 14.8 V, 5000 mAh, Lithium Polymer packs
ONBOARD COMPUTER	Intel i5 4570 based x86-64 Micro-ATX PC
COMMUNICATIONS INTERFACE	Verizon 4G LTE AT&T 4G LTE
CAMERAS	4x Logitech C930e Pilot Cameras 1x Panasonic Lumix DMC-GH3
SOFTWARE	Ubuntu 12.04 Robot Operation System (ROS): "Groovy"

4. Strategy

4.1 Testing Strategy

The team has been testing systematically since April, with testing sessions occurring after build sessions on a weekly basis. The rover has been driven extensively around the MIT campus, providing the driver with practice in speed control and turning. The rover's performance was tested across a variety of terrains, including the stairs of the entrance to MIT's Lobby 7, sand pits on the MIT campus, rocky construction areas, and the uneven landscapes in Boston Common.

These tests provided Mission Control and onsite crew members with invaluable practice in driving, spotting, navigating, and onsite debugging.

4.1.1 Testing Results

The rover performed well in all of the terrains tested so far by the team. The vision system is consistent across all four terrains, and the drive train is diverse enough as well. The sample retrieval system is generally more efficient in softer terrain, such as loose dirt or sand, in which the claw can dig under the desired sample much like an excavator. The slip steer drive train performs extremely well on gravel and dirt, though can tend to over slip on sand. Another problem the rover faces when traversing sand is the risk of burying a wheel. This problem occurs when accelerating too fast before the wheels gain traction. Overall, the system is effective in all terrains, but the rover's speed must be kept lower in the sand pits.

4.2 Overall Strategy

Due to the success of last year's optimization strategy, this strategy will be used as the basis for this year's competition. From last year's technical report,[3]

$$K_p = \frac{T_{remaining}}{t_{pickup} \cdot P(score_i \geq p) + t_{find}}$$

where K_p represents the expected number of rocks that can be picked up in each area, given p, the minimum score of the rocks that should be picked up. The resulting expected total score is

$$E[totalscore] = \sum_i score_i \cdot P(rock_i)$$

The main modification required is taking into account the capabilities of the Lumix DMC-GH3 camera used as the spotter camera. In testing, even small rocks could be spotted at distances of around 30 meters; in comparison, the pilot cameras are typically limited to less than 3 meters. As the characteristic length scale of the rock field is 100 meters, surveying territory will yield significant information on rock distribution.

The team augmented the model by considering an array of cells A[x][y] representing the probability of a rock at (x,y). So, A[x][y] = P(rock at (x,y)).

PARAMETERS

p_0	prior probability of a rock at a cell
p_{pass}	probability of missing a rock given the image covers the rock location
p_{hit}	probability of observing a rock given the image covers the rock location

---

[3] Source: MIT Robotics Team 2014 Robo-Ops Final Report

INPUT

$(x_i,y_i),\sigma_i$	points where rocks were observed, and estimated error in position

Using the GH3 image data, spotters provide a set of points where rocks were detected; call the map cells within σ of (x,y) image "hits." Call the rest of the cells image "passes," areas where no rocks were visible. The rock probability distribution can then be updated using Bayes' Rule.[4]

Image Hit: $A[x][y] \leftarrow A[x][y] * p_{pass} / (p_{pass} + p_{hit})$
Image Pass: $A[x][y] \leftarrow A[x][y] * p_{hit} / (p_{pass} + p_{hit})$

The array A[x][y] represents the probability of a rock within each cell. A modified travelling salesman algorithm can then be run, taking into account the expected number of rocks detectable along a route, the colors of the rocks, time required for travel, and bonuses such as visiting each terrain at least once.

4.3 Mission Control Strategy

Mission Control will be operated out of MIT room 33-132. The room offers considerable space, two large projectors, and 16 available monitors for GUI use. The room also has a back wall of windows where spectators can watch from unobtrusively.

4.3.1 Staffing

Mission Control requires the operation of at least five people: a driver, at least one spotter, a navigator, a GH3 camera controller, and a rover monitor. The driver will operate the Pilot GUI and work directly with the navigator to plan the best route based on the information gathered by the spotters, with the driver having the final say. The spotters will work directly with the camera controller in order to focus Mission Control's search efforts on areas of high interest or to confirm locations of rocks. The rover monitor will be staffed by the team's software lead and is responsible for monitoring the status of Mission Control's connection to the rover as well as monitoring the rover's internal communication system to ensure that everything is functioning properly. In the case of poor connection, the rover monitor will also be able to adjust settings in the rover's streaming servers to reduce latency.

---

[4] Source: "Introduction to Electrical Engineering and Computer Science I"

4.3.2 Practicing

In the testing day before competition, the team will focus on testing the rover's vision and network connections in all four terrain areas. The tests will be conducted much like the practice runs the team has conducted at home: on-site members will hide rocks, and Mission Control will then search and retrieve them.

4.3.3 Decision Making Strategy

Decisions will be based around the team's overall strategy (Section 4.2). Overall decisions will be made by a combination of the team's driver and navigator, since they have the most information on hand, with the driver having the final say. Spotters will also make the decisions of which areas the camera controller should focus on, based on what they are able to see from previous received photos.

4.3.4 Plan for Contingencies / Redundancies

Through the team's many test outings so far, both the on-site team and Mission Control have had experience with live debugging. Mission Control is also currently planning to have a number of unassigned team members who will be ready to carry out necessary work or take on extra jobs. Testing in practice Mission Control sessions will determine the final number of spotters that are used to maximize efficiency and decrease redundancy.

5. Budget

5.1 Distribution of Funds

TABLE 2: DISTRIBUTION OF FUNDS

ITEM	COST ($)
MECHANICAL MATERIALS	~500
WATER JETTING	~630
PANASONIC LUMIX GH3	1500
MOTHERBOARD & HARD DRIVE	180
UPGRADED MOTORS & SERVOS	480
UPGRADED ELECTRONICS	~430
ROBO-OPS REGISTRATION	920
OUTREACH & MEDIA	~1700
TRAVEL EXPENSES	2500
ROVER SHIPPING	~2000
TOTAL	~10840

5.2 Sponsors / Grants

TABLE 3: SPONSORS / GRANTS

COMPANY	CONTRIBUTION TYPE
MIT AERO/ASTRO	Monetary
ROBOTIS	In-kind
MIT LINCOLN LABS	Monetary / Working Space
MIT EDGERTON CENTER	Monetary / Working Space
INDIVIDUAL DONATIONS*	Monetary

* through the team's crowdfunding campaign; see Appendix VI: Crowdfunding Backers

6. Public / Stakeholder Engagement

This year the team has undergone many outreach events and has made a serious effort to connect to MIT undergraduate students, MIT graduate students, MIT alumni, and K-12 students. The most consistent method of outreach has been the recently launched website (http://roboteam.mit.edu/) and the team's Facebook page "MIT Robotics Team." The Facebook page has over 4,500 "likes" and has had a peak reach this year of about 7,800 users.

Along with social media outreach, the team has also been very active in many STEM related events this year. These events include:

Maker Faire (October 4th, 2014)

The team had a booth at the MIT Maker Faire, where team members demoed the previous year's rover to hundreds of people, ranging from five-year-old aspiring engineers to graduate students. Visitors really enjoyed driving the rover's arm to pick up rocks and stack blocks.

Sloan C-Function (November 13th, 2014)

The team demoed the previous year's rover at a MIT Sloan School of Management event. The event focused on connecting engineers with business students and trying to foster relationships between the two. Team members drove the rover around the dining area and participated in a mini show about a love affair between an MBA student and the rover! The team generated a lot of interest and was able to recruit people for the team's business development.

Science on Saturday (November 15th, 2014)

The MIT Underwater Robotics team hosted an event for about one hundred elementary and middle school students. Members representing the team staffed a booth, where they encouraged visitors to drive the rover's arm and explained to young aspiring engineers what it really means to be part of a robotics team.

Thanksgiving Mentoring (November 25th-December 5th, 2014)

Many team members returned to their homes over Thanksgiving weekend, where they worked with local FIRST FTC and FRC robotics teams in cities ranging from Hewlett, New York to Diamond Bar, California. Members were encouraged to mentor their former teams and share what they've learned from being part of the MIT Robotics Team.

MIT Venture Capital & Innovation Conference (December 12th, 2014)

The team had a booth, where members demonstrated the technology used on the rover. Many visiting companies and investors showed interest in the networking technology used. Visitors also had the chance to practice picking up hockey pucks on stage. During the event, more MIT Sloan business students and even MIT alumni expressed interest in joining the team in January.

Vecna Human and Robot 5K (April 12, 2015)

The team entered the rover into a Robot Spring Challenge organized by Vecna Cares as part of their Human 5K and Robot Race. The event focused on promoting innovation through technology, while also supporting the Vecna Cares Charitable Trust. The rover garnered significant exposure and added to the diversity of robots at the race. Attendees at the race, including families, marathon runners, volunteers, MIT alumni, and members of other local robotics teams, all had the chance to observe a robot designed for a non-racing purpose. After the race, the rover was featured in an MIT alumni video production about the event.

CPW 2015 Activities Midway & Edgerton Center Open House (April 18, 2015)

The team had a booth at the MIT Campus Preview Weekend (CPW) Activities Midway, where team members showcased the rover to over 1000 incoming students with an interest in STEM. Potential students who visited the booth had the chance to practice stacking hockey pucks with the rover's arm. After the Midway, visiting students were encouraged to visit the team's workspace at the MIT Edgerton Center, where team members described a typical work day.

The team also ran a successful crowdfunding campaign, with significant influence coming from announcements made on our social media page. During the 30-day campaign, the team raised $29,959 from approximately 170 donors. "Thank you" rewards included personalized machined thank you plaques, custom-made robotic parts, opportunities to remotely drive the rover, as well as options to visit the team and directly interact with team members during a work session.

References

"Introduction to Electrical Engineering and Computer Science I." 2011. MIT OpenCourseWare. <http://ocw.mit.edu/courses/electrical-engineering-and-computer-science/6-01sc-introduction-to-electrical-engineering-and-computer-science-i-spring-2011/unit-4-probability-and-planning/discrete-probability/MIT6_01SCS11_chap07.pdf>.

MIT Robotics Team. 2014. <http://nia-cms.nianet.org/getattachment/RoboOps/Previous-Winning-Teams/2014-Winners/Robo-Ops-2014---Final-Report---MIT.pdf.aspx>.

SAE Pattern. <http://www.heavyequipmentforums.com/attachment.php?attachmentid=6267&stc=1&d=1181880306>.

Appendix I: Onboard Control Schematic

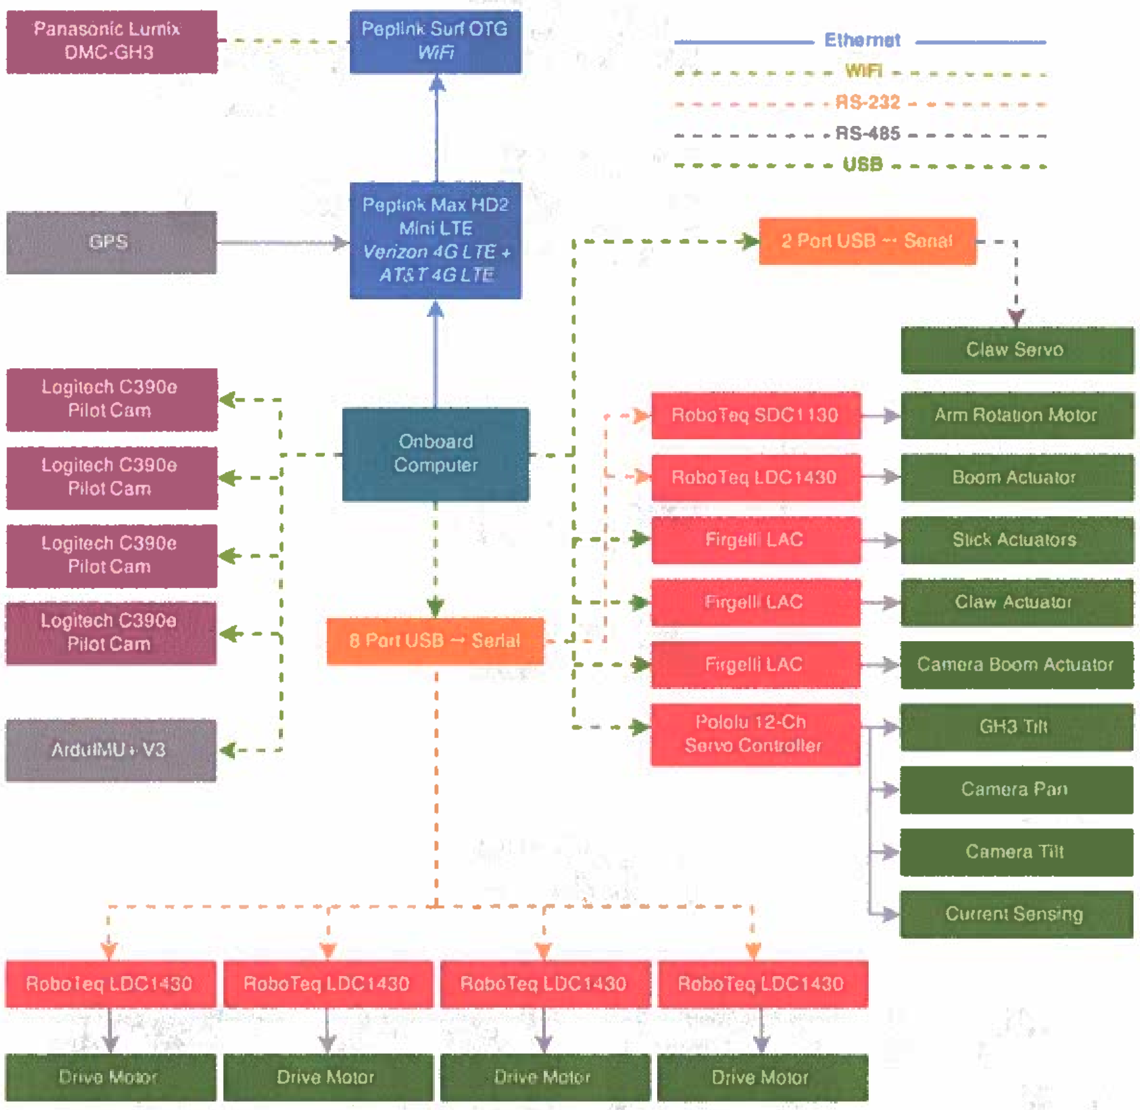

Appendix II: Rover Networking Schematic

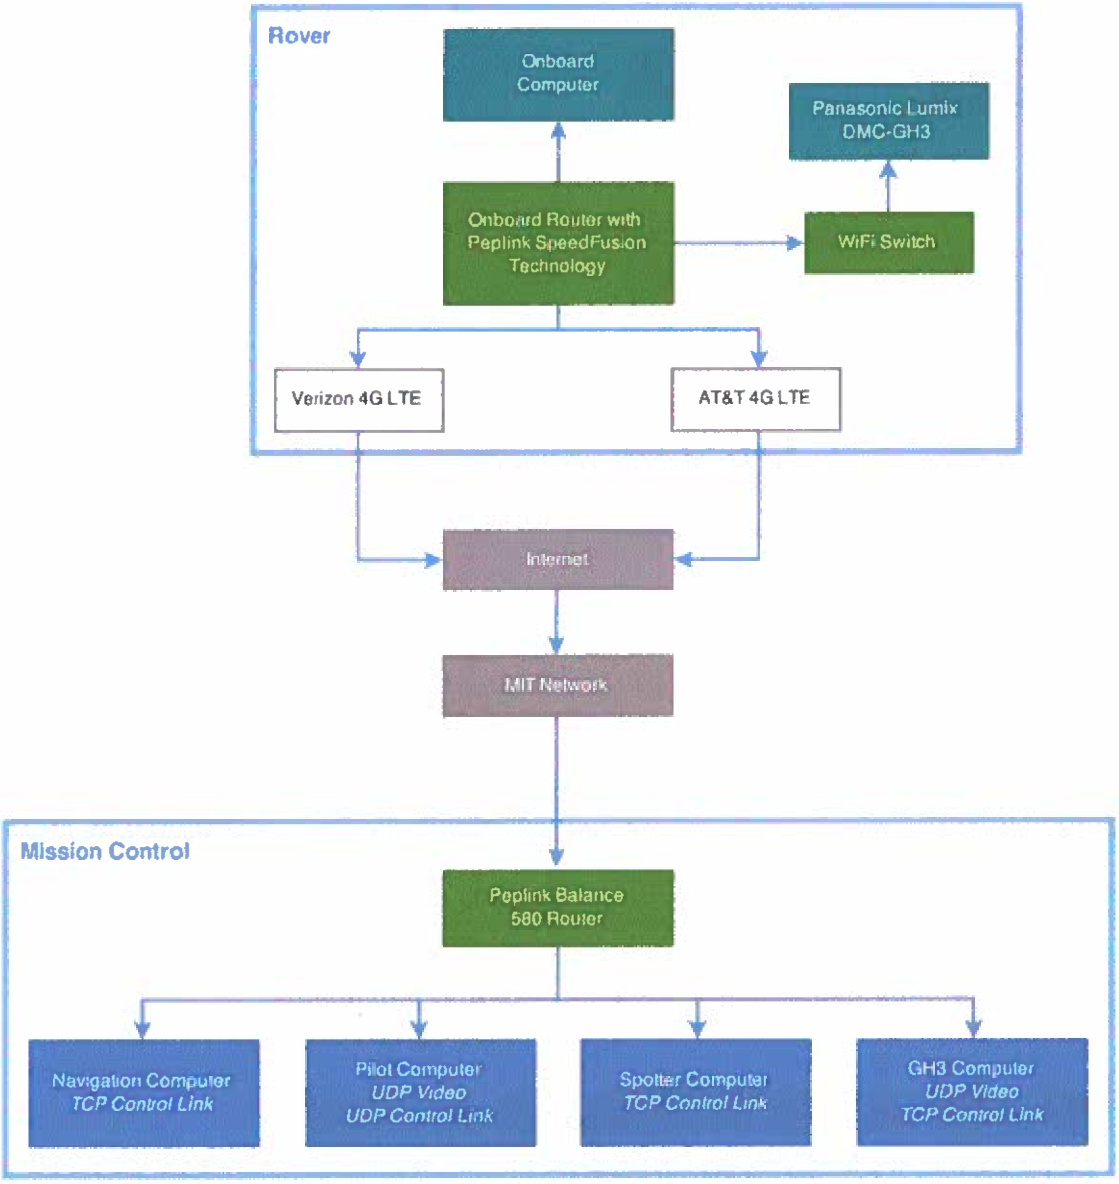

Appendix III: Software Schematic

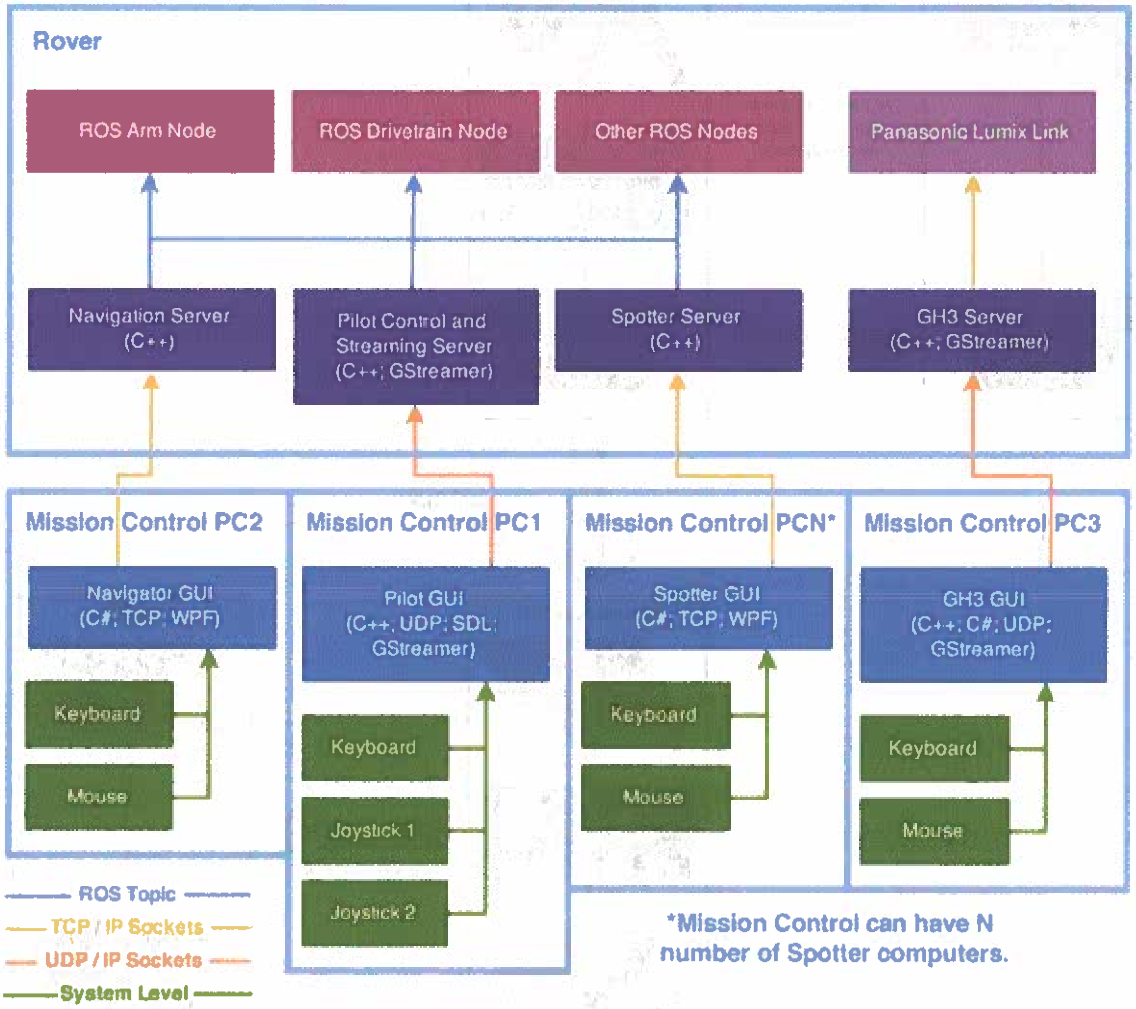

Appendix IV: GUI Control Flowchart

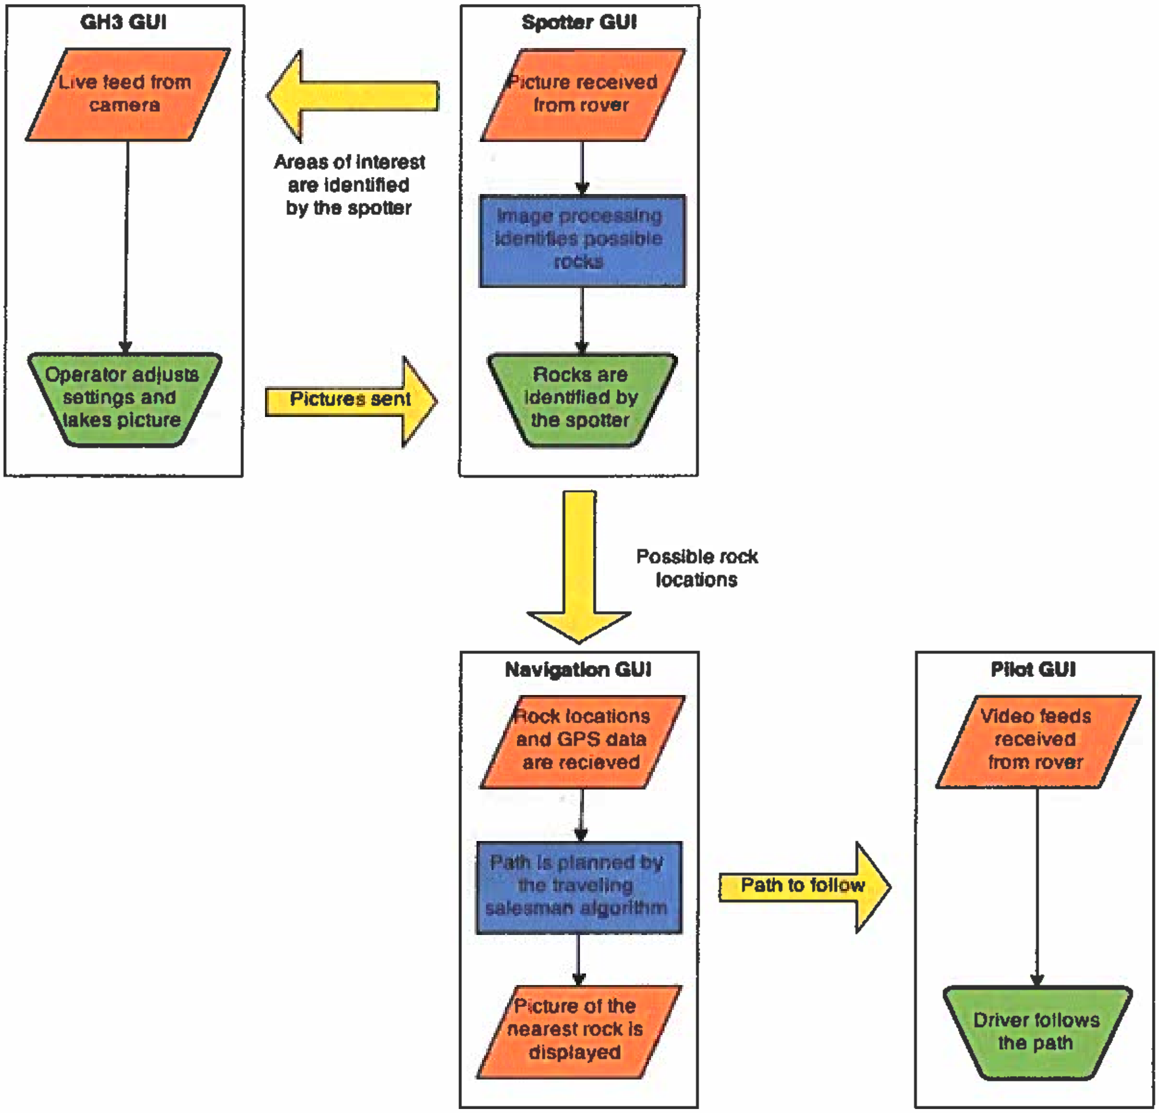

Appendix V: SAE Control Chart

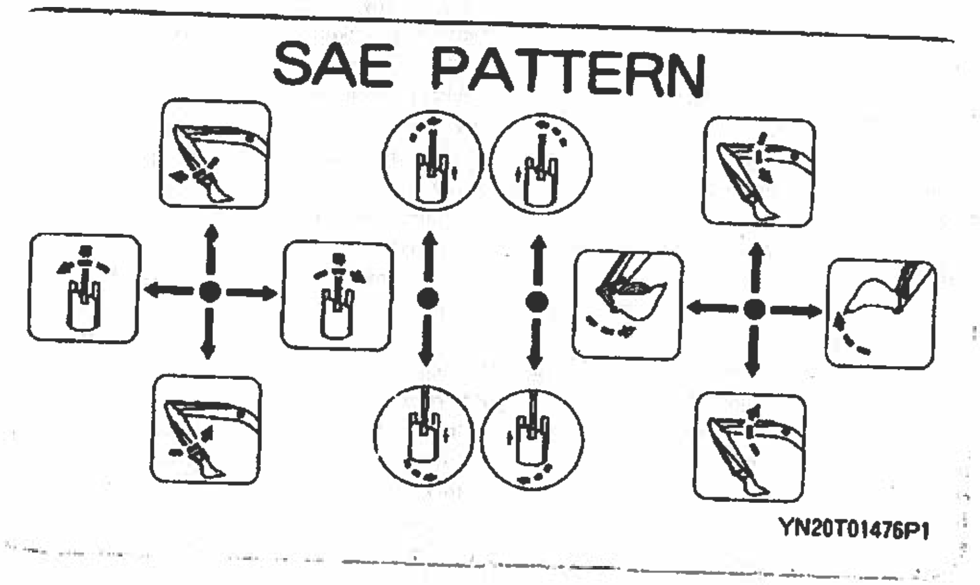

Appendix VI: Crowdfunding Donor List

Alan and Marilyn Cisar
Albert Hong
Alberto Giovannini
Alejandro Sánchez
Ali Hagen
Alvin Dulcan
Amarajeewa Jayananda Withanage
Amy Sutherland
Andres Sanchez
Andrew Hodgdon
Andy Rahman
Ann McLean-Muse
Anna Lim
Anonymous x13
Arthur Olivan
Asher Hecht-Bernstein
Baldomero Llorens-Ortiz
Bette Hochberger
Beverly Gibson
Beverly Gonda
Carl Wikstrom
Carol Williams
Chad Carpenter
Chee Wee Ang
Chris Holmes
Chris Mohr
Chuan Zhang
Craig, Ruth, Bari and Charlie Gruber
Dan Graham
Dave Humphreys
David Gruber
David Lam
David Pan
David Vogan
Dawn & Stuart Gulland
Dee Landergren
Dena Stein
Dennis Waldman
Diane Madrid
Edoh Amiran
Elaine Gruber
Eleanor and Harvey Sand
Elizabeth Pi
Elyse Cooper
Emily Bianchini
Emma Boykin
Enakshi Singh
Erin Artin
Flaviu Lepure
Frederick Shane
Gary Wingo
Gautam Kapur
Gay Hardy
Glenn Nelson
Hans Peter Brondmo
Harvey Golomb
Hideho Narita
Isitri Modak
James Lee
James Sholer
James Stameson
Javier De la Rosa Maura
Jean Hannon
Jennifer and Wah Cheung
Jennifer Lou
Jenny Wunderly
Jim Slattery
Jin Kim
Jinqiu Chen
Jinyong Lee
John Baird
John Bono
Jondean Haley
Jorge Amador
Jose DeLeon
Josh Copp
Julian Chan
Katherine Paseman
Kaushik Sinha
Kelsey Chan
Kevin Curran
Larissa Rodriguez
Lazarus Malabey
Leon Wheeless
Lin Floresca
Lisa Le
Lizhe Sun
Manu Capoor
Mary Muse
MaryAnn Dietrich
Matt Kaminer
Meg Ramsey
Mehmet Efe Akengin
Michael Beller
Michael Simonson
Mike McHugh
Mike Russell, MD
Mirna Romero
Mohamed Hamza
Morgan Beller
Naima Baird
Natalie Pitcher
Nick Schwartz
Olivier Van Dierdonck
Pablo Garcia-Silva
Patricia Rote
Patricia Santoro
Patrick McCuen
Paul Gaddis
Paul Grogan
Paul Philipps
Phil Richardson
Piuter Kogan
Pranam Chatterjee
Quentin Delepine
Rafi & Jacque Schwartz
Reid Sheftall, M.D.
Renee Moss
Reo Baird
Rita McComber
Rita Schwartz
Robert & Andi Moran
Robert Shin
Roger and Miho Baird
Roger Baird
Rose Panarelli
Rose Simonson
Rosey Sand
Ross Faneuf
Ruth Aniceto
Sabina Palisoc
Sabri Sansoy
Sam Lai
Sam Sherry
Sami Ainane
Sara Rogers
Sergio Elizondo
Shawn Baird
Shelley Beller
Shirley Chan
Sid Wax
Sudershan Virdi
Susan Delaney
Susan Simonson
Takako Wolfe
Thatcher Bell
Theo St. Francis
Thomas Fraser
Tiffany Dennen
Tom Janson
Tommy Tam
Trenton Lyke
Vivek Mukhopadhyay
Whitney Lohmeyer
William Chan
Yau Kai Cheung
Yuan-daw Tsai
Zoltan Maliga

Appendix VII: Figures & Tables