Design and Prototyping of a High-speed Camera System

J. Tyler Laseter III
In expected proposal of 1st Senior Project, BSEE
April 14, 2013

Abstract

A project for Fall 2013 in fulfillment of one of the Senior Project requirements is requested through this proposal involving the design and prototyping of a system for the acquisition of high-frame rate image series. This proposal has the objective of creating a modular high-speed camera system to be used in other research involving the examination of short-duration events. The camera system proposed is unique in that it will focus on reducing costs associated with such camera systems by utilizing commodity computing components (primarily memory and storage) where applicable, and using a modular construction approach to allow for a realistic prototyping timeline and enabling possibilities for further improvement, while also relying on field-programmable gate arrays (FPGAs) in the design to allow for a modular approach to the camera firmware. This proposal includes a description of the project, a review of literature involved in research on the topic, a plan of the design and a prototyping timeline, a discussion of Tyler's qualifications, the budget for this project, material referenced, and an appendix defining terms and acronyms.

Project Description

I propose to design and construct a high-speed camera utilizing several novel properties. A high-speed camera is a device used to record events which occur in short periods of time, with special characteristics to allow for those events to be then replayed as if they occurred much more slowly. This is accomplished by recording at a much higher framerate than is typically utilized for video, then playing that footage back again at one of those typical framerates. Most camcorders and television shows run at 29.97 frames per second (in the United States) and movies run at 24 frames per second (fps), whereas high speed cameras typically run between 500fps to over 10,000fps. When played back at 24fps, a 500fps camera will portray an event as if it happened 20.8 times slower than it actually did (one second is dilated to 20.8 seconds). This is an extremely useful quality, leading to near-universal utility in a large variety of research areas: crash testing, autonomous vehicles, machine vision, and preventative maintenance in industry; armor and explosives analysis in defense; speech analysis, motion capture, eye tracking, and insect and fish observations in biomedical applications. High speed cameras have started to become prevalent in television, making appearances in television shows such as Mythbusters and others (there are some on the Discovery and Science channels devoted exclusively to high speed footage). Despite the increased visibility to the average American, high speed cameras remain a highly rarefied and expensive device, typically ranging from between $30,000 to over $150,000, depending on performance characteristics. While there are a variety of characteristics that make them difficult to produce inexpensively, there are a variety of alterations to standard designs that could be accomplished to reduce the cost of high-speed cameras significantly. It is for this reason that it is proposed to design and build a high-speed camera system. An appendix is attached at the end of this proposal elaborating any abbreviations and terms nonstandard to those unfamiliar to the field of embedded design and imaging within electrical engineering.

Unique Challenges

In order to reduce the costs associated with high-speed cameras, an understanding of the difficulties presented by their unique properties must be reached. Unlike video cameras, high speed cameras – even on the low end of the market – typically have much higher sensor resolutions, typically one to four megapixels (720p is .9 megapixels, 1080p is 2.1 megapixels -- 4 megapixels is 4K, rarely seen out of the highest-end feature-film cameras). This presents an increase in the amount of data needed to be processed per frame, but the requirements are even more steeply increased by the number of frames per second recorded by the camera, as the amount of data required to record per second is determined by:

$$\frac{gigabytes}{second} = \frac{frames}{second} * \frac{pixels}{frame} * \frac{bits}{pixel} * \frac{1\ gigabyte}{8589934592\ bits}$$

The majority of high-speed camera sensors are greyscale, as it allows for data which does not need to be de-interlaced and colorized from the Bayer filter typically present on a color sensor, presenting a higher resolution without interpolation. Additionally, each pixel is typically ten bits rather than the eight of a color sensor, as ten presents a compromise between ballooning data rate and analog-to-digital converter (ADC) complexity while allowing for higher dynamic range (as there is still less than half as much data to process than 8-bit red-green-blue (RGB) in 24-bit color). Thus, for a typical two-megapixel high-speed sensor operating at 500 frames/second, we arrive at:

$$500\frac{frames}{second} * 2000000\frac{pixels}{frame} * 10\frac{bits}{pixel} * \frac{1\ gigabyte}{8589934592\ bits} = 1.16\ \frac{gigabytes}{second}$$

To compare, a typical high-end 3-sensor (with trichroic prism for greater light sensitivity without a Bayer filter) 1080p video camera, such as a Sony PWM-EX1, is around five megabytes/second. The challenges presented by these high data rates are the primary parameter which sets high-speed camera systems apart and makes them a unique design challenge. The plot below illustrates the rapid increase of data rate encountered in these systems.

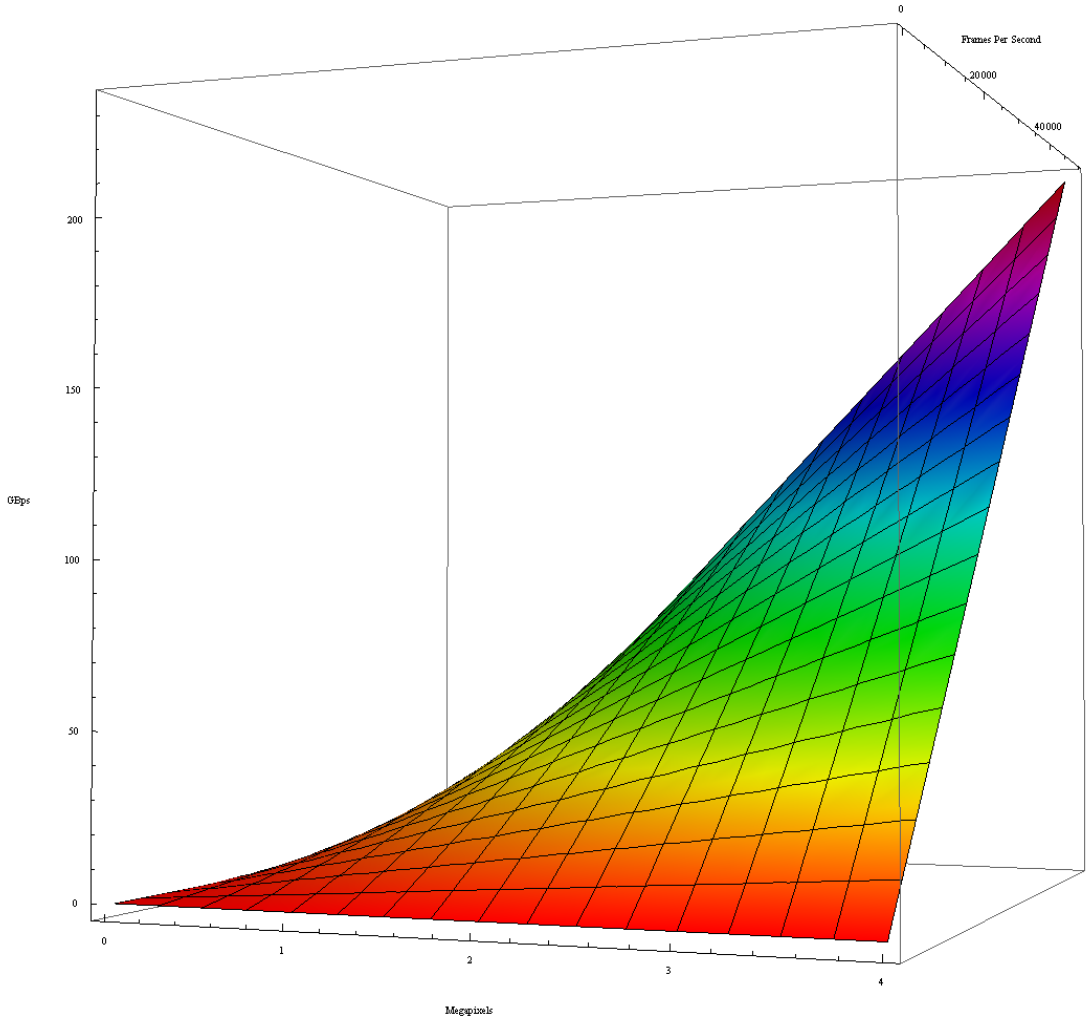

Additionally, there are a variety of other factors which must be attended to that are unique to high-speed cameras. Among them are light sensitivity, triggering, and data interfaces. When your image sensor is shooting 500fps, it can only receive less than 1/500th of a second of light (accounting for the

blanking time between frame captures), requiring intense illumination and very fast (high transmissivity, low f/stop) lenses. In some systems researched, the camera was able to synchronize with a light-emitting diode (LED) strobe system, allowing for extremely high intensity light without taxing the thermal characteristics of the power LED's by only emitting light while frames were being captured. Triggering is important as well – if a gigabyte per second is the storage cost of every second recorded, the beginning and ending of the recording is typically controlled through external transducers and signal conditioning equipment (light, sound, contact, acceleration, etc.). There are few data interfaces capable of transferring the high volume of data live, so typically either application-specific interfaces such as Camera Link must be used or very efficient real-time compression must be utilized at a significant processing cost. This issue can be ameliorated by only outputting a low frame rate monitor signal while aligning the camera, then blanking that output and only recording to on-board storage before transferring the data recorded. For this application, gigabit Ethernet would work extremely well, as it can be transported over fiber optic cables extreme distances, and be processed and switched through standard commercial-off-the-shelf (COTS) equipment.

Gigabit Ethernet is an excellent example of how utilization of specific design choices coupled with COTS hardware can reduce the cost of the camera system significantly. As I stated, the primary goal of this project is to not only create a high-speed camera system, but also create one at significantly lower cost than those typically available commercially. To do so, utilization of solid state hard drives and commercial dynamic random access memory (DRAM), both capable of high-speed operation, presents a much cheaper option than what is typically utilized in commercial cameras. Installing laptop-style Dual Data Rate 3 (DDR3) laptop dual inline memory modules (DIMMs) is much cheaper than soldering the memory chips down to one of the custom circuit boards that will compose the camera, and solid-state drives allow for nonvolatile storage that also is extremely fast. Part of the high cost of high-speed cameras stems in part from the market inertia associated with memory costs – as illustrated below, they have dropped drastically over time:

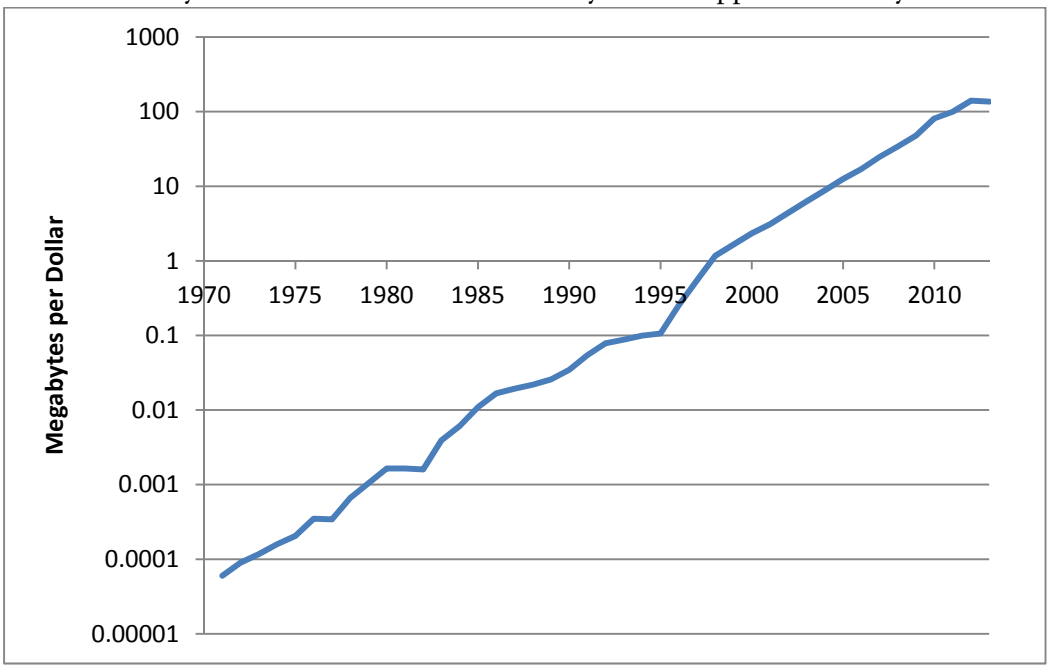

Furthermore, earlier cameras required the use of expensive and inefficient static RAM (SRAM), as DRAM lacked the speed necessary for these applications. All of this has changed, except the market expectation of what a high-speed camera *should* cost has not. Additionally, cameras are typically outfitted with every potential option, rather than embracing a modular model where only capabilities required for any given application contribute expense to the final project.

Literature Review

Typical Signal Processing Architectures

Such high bandwidths of device-specific data can only be handled well (outside of production volumes) with field-programmable gate arrays (FPGAs). Such devices are used in all referenced research exclusively, as they are the only acceptable method. The primary uses of the FPGA in such a system involve serializing the parallel data output from the sensor [1], applying any desired image processing [2], storing that data in memory, then being able to transfer that data to a computer, either as a file or as streaming data. To send streaming high-speed footage, which has some advantages in processing, significant investments in software development must be made, and the only common computer interface adequate to transfer the quantity of data involved is gigabit Ethernet [3], which also requires extensive software development. For these reasons, an intelligent design methodology involves the camera system as several modules which can be connected together or exchanged [1, 2, 3, 6, 8]. In a research environment, this allows for advancements in some areas to prevent having to remanufacture the entire system. Furthermore, the use of more than one FPGA device allows for lower-performance and much less expensive devices to be used, devoting each to a specific task. The tasks involved that would optimize well to a modular approach would be a power supply, sensor module, processing/memory module, and interface module [8]. For any implementation of the camera, the sensor module and power supply would likely remain the same, as its parameters are primarily governed by performance characteristics of the sensor used. However, optimizations in both processing and interface are myriad, so allowing their exchange allows for a conservative initial design with significant opportunity for expansion, significantly reducing initial cost. For example, an initial design would do little image processing (other than what is required) and only allow for transfer of prerecorded data over an interface such as USB or a memory card attached to the board. Later improvements could involve gigabit Ethernet real-time video transfer, dynamic range correction, video compression, statistics, image recognition, or other processing tasks without sacrificing the initial investment in the sensor processing system by having to redesign the camera entirely.

Performance Characteristics

Performance of such a system is primarily measured in its light sensitivity, noise, and framerate [4], all parameters which can be somewhat optimized by the software written to the FPGA as well as the analog design of the clocks and power supplies. Additionally, dimensional consistency (a lack of significant geometric distortion) is imperative [5], as typically not only is a high-speed record of an event required, but its analysis requires measurements to be made on the captured

images to calculate the dynamics of the scenario. Commercial fixed-focal length lenses are utilized for the purpose with a good balance of transmissivity (f/stop) and distortion to avoid this problem. The chosen sensor will utilize complementary metal oxide silicon (CMOS) rather than charge-coupled device (CCD) construction to allow for on-chip ADCs, greatly simplifying camera design while making the system more temperature-insensitive[7].Achieving all of these performance goals is nontrivial, and any attempt to do so while controlling cost requires a design which is not excessively complex. Thus, modular construction as previously outlined fits closely with meeting these goals, electronically and optically (interchangeable lenses allow for excellent flexibility.)

Plan of Work

The work required for this project is typical of any electronics design – prototype cycle, and the modular design I propose allows for the work to be subdivided into sections which can be designed, built, and tested individually to reduce the amount of work required per cycle. I consider a cycle to incorporate seven discrete steps, and each subassembly will incorporate each of these discrete steps as follows: design requirement specification, component selection, schematic capture, board layout, board population (manufacturing), testing, and development (for assemblies which require software). Below is a graphic demonstrating the subassemblies which will be incorporated into the camera system:

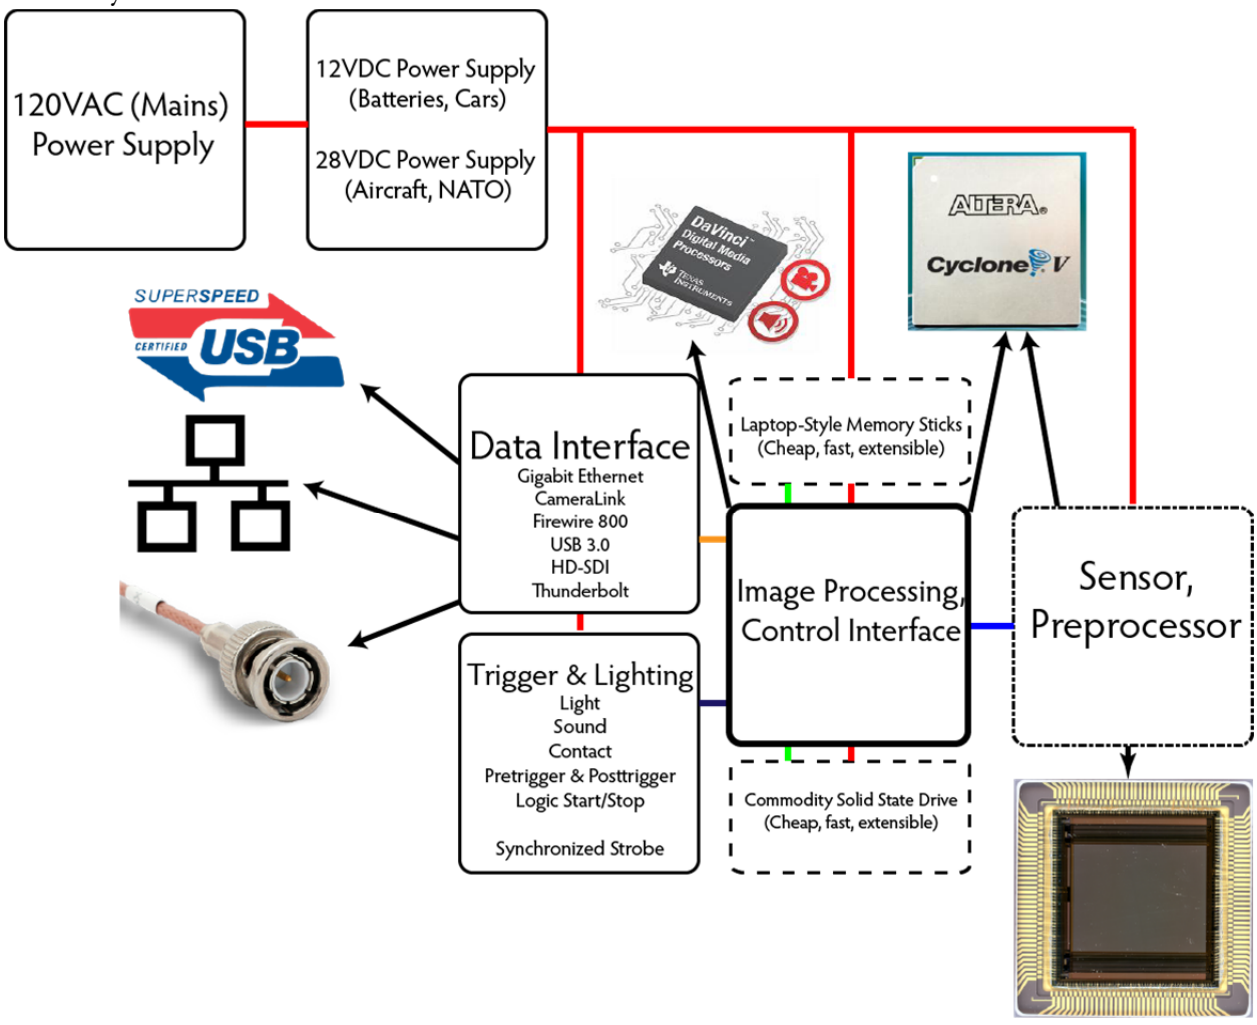

Due to the modular aspect of these subassemblies, different revisions of each subassembly could provide additional or different features from prior revisions. Features must be chosen for the first revision, however, and it is proposed that the first revision modules would be as follows (particularly incorporating a desire to keep software development a reasonable proportion of work required):

- 120VAC power supply
    - COTS 12vdc supply to retain UL approval and to prevent requiring the manufacture of any circuit boards
- DC-DC switching power supply to generate all required voltages from an 11-30vdc input
    - 12v isolated & regulated
    - 5v, 3.3v, 2.5v, 1.8, 0.7v

- Image processing & control
- Image Processing & Control
    - Cyclone V FPGA for buffering data into DRAM and transferring it into storage
    - Microcontroller, potentially TI DaVinci ARM video DSP, to process video and provide a camera control interface over serial, USB, or fast Ethernet (10/100)
    - 8GB DDR3 DRAM in the form of laptop DIMMs
    - SATA solid state drive, likely 160 or 256GB, at least 250MB/s write (such as Crucial M4 series consumer solid state drives)
- Sensor & preprocessor
    - Digital camera sensor, likely Alexima 4MP 350-500fps globally shuttered CMOS greyscale 10b/pixel high speed camera sensor (surplus stock, at no cost)
    - Cyclone V FPGA for clocking the camera sensor and serializing the parallel data output from the sensor (16 10-bit parallel ports) for further processing on the Image Processing & Control board as well as interfacing control commands to the sensor and generating the required clock signals
- Data interface
    - Initially eSATA directly to the SSD, eventually Gigabit Ethernet
- Trigger & Lighting
    - pair of BNC connectors configurable as start, stop, contact closure, or TTL gate

As parenthetically noted under the entry for the camera sensor, the sensor I have selected was purchased surplus (in volume of 16) and stored unused by a graduate student whom I am friends with and who has arranged with me to donate several sensors to this project in exchange for the information needed to build a duplicate camera system. This is the enabling factor for this project, as typically such sensors range $1000-$5000, a prohibitively high cost for a traditional senior project. The work can be pipelined to some degree, as circuit boards and parts orders incorporate shipping and lead times, and below I outline a plan of work in ten weeks ("order board" includes parts):

- Week 1
    - Evaluate requirements for inter-board communications, develop a connection scheme between boards and document all power rails and interconnections required
    - Design power supply, capture schematic & layout board, order board
- Week 2
    - Select FPGA for image preprocessor board based on sensor requirements
    - Capture schematic & layout board for sensor preprocessor, order board
- Week 3
    - Build & test power supply
    - Select FPGA and microcontroller for main processor board
    - Begin selecting other components and performing schematic capture
- Week 4
    - Layout board for main processor and order board
    - Populate & test FPGA preprocessor board
- Week 5
-
    - Begin writing Verilog code for image processing board
    - Design external interfaces board (combining data & trigger)
- Week 6
    - Finish writing preprocessor Verilog, test sensor data output with logic analyzer
    - Schematic capture and board layout for external interfaces, order board
    - Order SSD & DRAM
- Week 7
    - Populate & test main processor board
    - Begin writing main processor board code
- Week 8
    - Populate & test external interfaces board
    - Continue writing main processor board code
    - Design and build enclosure incorporating a Nikon F lens mount fabricated from aluminum with proper fixturings for all connections and boards
- Week 9
    - Continue code development for camera system

- Week 10
    - Finish development and capture example footage

Qualifications of Designer

Tyler Laseter designed his first circuit board five years ago when an employee of SAND Holdings, LLC in Atlanta, GA for the NetDefib networking appliance (http://www.netdefib.com) . Since then, he has participated in the Tesla Orchesta, an engineering and performance student group. Tyler designed a variety of circuit boards for interfacing with the Tesla coils involved with the project, including MIDI-to-fiber optic converters and the construction of a ten-kilowatt power supply. Tyler has also designed a variety of other systems for personal projects, and designed a MIPS-architecture processor development board with a fellow undergraduate to fulfill a final project requirement in EECS 314 (Computer Architecture) this year (Spring 2013). Tyler is also finishing his photography minor, and his experience with imaging not only includes ordinary photographic process, but also radiography; Tyler currently owns a small business involving industrial radiography of circuit boards, potted assemblies, and electromechanical systems. Through these experiences, Tyler not only can design schematics, but also has skills in circuit board layout CAD and building out circuit boards consisting exclusively of extremely small surface-mount components by hand – an invaluable skill for the project proposed. Tyler's experiences with radiography are also a boon for working with BGA parts, as X-rays are the only method by which the part-to-board connections may be examined. Tyler also took EECS 301 (Digital Design Laboratory), in which he became very familiar with Verilog design for Altera FPGA's using the Quartus software environment.

Anticipated Faculty Involvement

The project proposed is intended to be designed entirely by Tyler and potentially other undergraduates, however faculty assistance is still required. Primarily, this assistance is requested in the form of financial backing for the project at hand. Technical assistance is not anticipated to be required, but is always appreciated. Technical assistance will likely primarily stem from Circuits lab staff member Edwin Burwell, who himself completed a senior project involving an imaging system several years ago and is enthusiastic to assist me with the project proposed here.

Budget

The budget proposed below takes into account donated sensors and all anticipated costs, however a contingency is included as it is not a budget prepared from a final bill-of-materials and will likely evolve with the project (but should not exceed the contingency mentioned.) Cameras researched with similar capabilities exceed $40,000 (list). This budget is for one camera unit, however, two will likely be manufactured as with any prototyping project undertaken (as a contingency itself, such projects are always redundant).

Item	Cost	Quantity	Total
Alexima sensor	0*	1	0
Four-layer printed circuit board	66	4	264
Solid-state drive	200	1	200
16GB DRAM	100	1	100
Cyclone V FPGA	30	2	60
TI DaVinci digital video ARM SoC	30	1	30
Connectors	3*	10	30
Passives	1	40	40
Power supply components	3	10	30
Lens	470	1	470
Enclosure	100	1	100
Contingency	300	1	300
(total, compare at $40,000)			$1624

Bibliography

[1] M. Caselle et al., "Ultra-fast streaming camera platform for scientific applications," in 2012 18th IEEE-NPSS Real Time Conference (RT), Berkeley, CA, 2012, pp.1-8. doi: 10.1109/RTC.2012.6418369

[2] B. Yan et al., "Design of CMOS image acquisition system based on FPGA," in Symposium of Industrial Electronics and Applications, Langkawi, 2011, pp.1726-1730. doi: 10.1109/ICIEA.2011.5975870

[3] P. van Hahn and S. Lon, “GigE Advanced Imaging Sensor,” BAE Systems Technology, Nashua, Final Report. DUNS# 00-514-9120 SIC #3679 TIN#52-2268742 under contract HR0011-09-C-0045, Nov. 24, 2009.

[4] Y. Shang et al., "Design and characteristics measurement of a high-speed CMOS imaging system," in IEEE International Workshop on Imaging Systems and Techniques, Shenzhen, 2009, pp.127-131. doi: 10.1109/IST.2009.5071617

[5] Y. Dzenis and R. Feng, “High-Performance Digital Imaging System for Development and Characterization of Novel Materials and Processes,” University of Nebraska-Lincoln, Lincoln, AFOSR report. AFRL-SR-AR-TR-07-0077, Aug. 8, 2006.

[6] U. Muehlmann et al., "A new high speed CMOS camera for real-time tracking applications," in 2004 IEEE International Conference on Robotics and Automation, New Orleans, LA, 2004, pp. 5195-5200. doi: 10.1109/ROBOT.2004.1302542

[7] M. Hillebrand et al., "High speed camera system using a CMOS image sensor," in IEEE Intelligent Vehicles Symposium, Alcalá de Henares, 2000, pp.656-661. doi: 10.1109/IVS.2000.898423

[8] N. Stevanovic et al., "A high speed camera system based on an image sensor in standard CMOS technology," in IEEE International Symposium on Circuits and Systems, Orlando, FL, 1999, pp.148-151. doi: 10.1109/ISCAS.1999.777532

Appendix of Abbreviations and Terms

- ADC: Analog to Digital converter. Converts voltage to digital information.
- ARM: A brand of processor core popular in cell phones and other embedded applications
- Bayer Filter: A filter placed over an image sensor which appears like a red-green-and-blue stained-glass window, filtering the light for each pixel to be proportional to red, green, or blue light. Cheap and the most common technique for imaging in color, it requires sophisticated image processing to colorize and align all the pixels.
- BNC: Bayonet Neil-Concelman, a type of connector (named for its designers)
- Clock: An electronic signal to pace and time internal systems.
- COTS: Commercial Off-the-Shelf
- DDR3: Dual Data Rate Generation 3, the most recent revision of common computer mem.
- DIMM: Dual Inline Memory Module, a format for computer memory
- DRAM: Dynamic RAM, a type of high-density cheap memory common in computers
- eSATA: external Serial ATA, a type of data transfer protocol & interconnect
- FPGA: Field Programmable Gate Array, a type of data processor that is thoroughly reconfigurable, programmed through language which defines not computer code run on the device but the arrangement of gates on the processor itself
- Interlacing: An old term from television technology, where half of each frame was scanned at a time to reduce instantaneous data rate without sacrificing overall picture size.
- Interpolation: Determining a value of a pixel by its neighbors, used with *bayer filters* to have a larger image size than simply the size of the sensor divided by three (RGB).
- Megapixels: A unit of imaging sensor size. Millions of active pixels on the sensor.
- Nikon F: A type of lens mount designed by Nikon in 1959 and used to this day
- SSD: Solid State Drive, a type of hard disk-like device with no moving parts, utilizing memory chips instead
- TI: Texas Instruments
- Trichroic Prism: A type of prism which splits incoming white light into its red, green, and blue components for processing by three greyscale sensors to obtain a color image. Expensive when compared with a *bayer filter*, however each pixel receives more light than it would in a Bayer system.