

Software Defined Multi-Spectral Imaging for Arctic Sensor Networks

Sam Siewert^{*1,2}, Vivek Angoth², Ramnarayan Krishnamurthy², Karthikeyan Mani², Kenrick Mock^{3,4}, Surjith B. Singh², Saurav Srivistava², Chris Wagner², Ryan Claus¹, Matthew Demi Vis¹.

¹Embry Riddle Aeronautical University Prescott, 3700 Willow Creek Road, Prescott, Arizona, USA 86301; ²University of Colorado Boulder, Campus Box 425, Boulder, Colorado, USA 80309-0425; ³Arctic Domain Awareness Center, 3211 Providence Drive, Anchorage, Alaska, USA 99508-4614; ⁴University of Alaska Anchorage, 3211 Providence Drive, Anchorage, Alaska, USA 99508-4614

ABSTRACT

Availability of off-the-shelf infrared sensors combined with high definition visible cameras has made possible the construction of a Software Defined Multi-Spectral Imager (SDMSI) combining long-wave, near-infrared and visible imaging. The SDMSI requires a real-time embedded processor to fuse images and to create real-time depth maps for opportunistic uplink in sensor networks. Researchers at Embry Riddle Aeronautical University working with University of Alaska Anchorage at the Arctic Domain Awareness Center and the University of Colorado Boulder have built several versions of a low-cost drop-in-place SDMSI to test alternatives for power efficient image fusion. The SDMSI is intended for use in field applications including marine security, search and rescue operations and environmental surveys in the Arctic region. Based on Arctic marine sensor network mission goals, the team has designed the SDMSI to include features to rank images based on saliency and to provide on camera fusion and depth mapping. A major challenge has been the design of the camera computing system to operate within a 10 to 20 Watt power budget. This paper presents a power analysis of three options: 1) multi-core, 2) field programmable gate array with multi-core, and 3) graphics processing units with multi-core. For each test, power consumed for common fusion workloads has been measured at a range of frame rates and resolutions. Detailed analyses from our power efficiency comparison for workloads specific to stereo depth mapping and sensor fusion are summarized. Preliminary mission feasibility results from testing with off-the-shelf long-wave infrared and visible cameras in Alaska and Arizona are also summarized to demonstrate the value of the SDMSI for applications such as ice tracking, ocean color, soil moisture, animal and marine vessel detection and tracking. The goal is to select the most power efficient solution for the SDMSI for use on UAVs (Unoccupied Aerial Vehicles) and other drop-in-place installations in the Arctic. The prototype selected will be field tested in Alaska in the summer of 2016.

Keywords: software defined, multi-spectral imager, field programmable gate array, general purpose graphics co-processor, sensor fusion, stereo vision, long-wave infrared, near-infrared, Arctic, saliency, power efficient, real-time

1. INTRODUCTION

The Software Defined Multi-Spectral Imager (SDMSI) combines long-wave, near-infrared and visible imaging into a single smart camera with a wide range of spatial, spectral and temporal resolution integrating two or more off-the-shelf detectors with continuous image processing. The SDMSI integrates a real-time embedded processor to fuse images, detect targets and to create real-time depth maps for use on aerial platforms or as fixed node in a sensor network. Several SDMSI configurations have been built on a test-bench to explore alternatives for power efficient continuous image fusion with the idea that most of the primitive transform operations can be used in common with stereo matching when two visible and one infrared camera is used. For example common pixel or feature-level image registration, pyramidal transforms for resolution conversion, segmentation, and gradient edge detection. The SDMSI concept integrates co-processing for continuous transform such that the SDMSI main processor, which runs embedded Linux, can be devoted to the sensor network interface, user interfaces, and higher-level scene processing algorithms such as saliency and target recognition.

Evolution of SDMSI from Previous Work

The SDMSI architecture has evolved from previous work by the same research team to develop a multi-spectral smart camera using an FPGA (Field Programmable Gate Array), a custom PCB (Printed Circuit Board) and lower-level programming using HDL (Hardware Design Language) [7]. The previous work resulted in a prototype that was bench

tested, but it became clear that new SoC (System on Chip) technology and higher-level FPGA and GP-GPU (General Purpose Graphics Processing Unit) co-processors enabled the ability to focus more on algorithms and system design research and less on the hardware development to interface cameras to image transform processing. Based upon this earlier work, the team decided to place focus on power efficiency and development of image transform primitives that can be off-loaded to a co-processor using off-the-shelf detectors and computing, such that the current instrument is really defined only by software. The software definition focuses on image transforms for the FPGA or GP-GPU co-processor and common embedded Linux software for scene processing and SDMSI configuration and use. As such, the concept which was referred to as a Computational Photometer has now been re-designed and is referred to as the SDMSI.

Use of SDMSI for Mapping, Search and Monitoring

The SDMSI is envisioned to be a versatile smart camera that can be used on UAVs, in marine environments on vessels and in fixed locations for applications such as port perimeter security. Ideally the design will enable scientific missions as well as security. The SDMSI is intended to use off-the-shelf LWIR (Long Wave Infrared, 10-14 micron), NIR (Near Infrared 0.7-1.4 micron) and visible optics and detectors (panchromatic 0.39-0.70 micron) and as such can be adapted well to a range of applications. The core of the SDMSI is the continuous image processing capability at low power to provide stereo mapping as well as image fusion for heterogeneous (distinct and dissimilar cameras in real-time at frame rates of 30Hz or higher). The design for the Arctic region limits the use of grid and battery power, so a primary focus of the work presented here is to adapt existing well-known 3D mapping and fusion transform primitives to low-power computation such as FPGA and GP-GPU. The primary goal is in fact to determine the computing solution with the lowest continuous power requirement as a function of frame rates, resolution and number of cameras.

SDMSI Challenges

A major challenge has been the design of the camera computing system to operate within a 10 to 20 Watt power budget for continuous 30Hz fusion and 3D mapping. To put this goal in perspective, the 2015 NASA technology roadmap TA4: Robotics and Autonomous Systems, has set a goal for planetary rovers at 1 Hz at 5 Watts or less, stating that as “an example is the need for sensors that produce range images on the order of once per second for rover navigation, with 512 x 512 pixels or over 1 steradian field of view with range resolution on the order of 10 centimeters (cm) at 10 meters (m), with a power consumption on the order of 5 watts (W) or less” [20]. The SDMSI has a 30Hz rate at 640x480 resolution or higher with continuous transform to register and fuse heterogeneous visible and LWIR images while simultaneously providing 3D mapping in visible, using a total of 3 cameras (1.2 to 1.5 Watts per camera) using between 2.5 and 7.5 Watts for computation, for a total of 6.1 to 12 Watts worst-case. The off-grid power supply and UAV batteries have been designed for a maximum output of 20 Watts. While the two goals are not directly comparable, the SDMSI is a power efficient design that will be tested in the harsh Arctic environment, similar to requirements for deep space instruments (other than radiation hardening).

SDMSI Power Efficiency for Fusion and 3D Mapping

This paper presents a power analysis of three options: 1) multi-core, 2) field programmable gate array co-processors with multi-core, and 3) graphics co-processing units with multi-core. The multi-core software only results are omitted since performance analysis is easily found using OpenCV and existing software libraries and tools that do not make use of co-processors [43]. In fact, the project research goal is to efficiently off-load common image transforms so that CPU (Central Processing Unit) cores can be used for machine learning, sensor network participation and human interfaces rather than pixel-level transforms, with the goal to optimize overall SDMSI power efficiency. For each test, power consumed for common fusion workloads has been measured at a range of frame rates and resolutions. The detailed analysis from our power efficiency comparison for workloads specific to stereo depth mapping and sensor fusion are being used to select the most ideal configuration for testing with the ADAC (Arctic Domain Awareness Center at University of Alaska) in summer 2016. Preliminary mission feasibility results from testing with off-the-shelf LWIR and visible cameras in Alaska and Arizona will also be presented to demonstrate the value of the SDMSI for applications such as ice tracking, soil moisture, ocean color mapping, animal and marine vessel detection and tracking. Future work will depend upon power efficient continuous transforms describe in the paper (many common to both fusion and 3D mapping applications).

*siewerts@erau.edu; phone:1-928-777-6929; fax:1-928-777-6945; <http://mercury.pr.erau.edu/~siewerts/>

2. MISSION OBJECTIVES

To date, the ADAC has provided the primary mission objectives for the SDMSI. ADAC has outlined key mission objectives to include: 1) coastal storm surge and ice tracking, 2) response to environmental threats in the region, 3) search and rescue operations, such those common to the USCG (United States Coast Guard), 4) port perimeter security. The SDMSI prototype for ADAC is intended for use in field applications including marine security, search and rescue operations and environmental surveys in the Arctic region. Based on the Arctic marine sensor network mission goals, the team has designed the SDMSI to include continuous mapping and image fusion features, with the goal to off-load the primary camera CPU so it is free provide intelligent features such as ranking images based on saliency, targets of interest and potential threats to security. Based on use in the Arctic, the research team has focused on low power continuous transforms to enable operation using 20 Watt or less power sources such as fuel cells, alternative (wind, solar in summer), and UAV battery or fuel cell systems, with a goal for the basic three camera configuration to use less than 10 Watts in continuous operation.

Primary Arctic Domain Awareness Mission

Deployment in marine environment for homeland security (ADAC) for use in ports, towers, vessels, buoys with long-term unattended operation. Use on vessels and UAVs for interactive search and rescue operations. In the summer of 2015, basic Arctic and maritime feasibility analysis was completed establishing that marine vessels and ice could be tracked well and for vessels many miles off shore, especially with the LWIR camera which can detect and distinguish vessels based on engine and exhaust stack configurations fused with S-AIS data (Satellite Automatic Identification System) [36]. For the ADAC maritime security and tracking missions, visible with LWIR fusion can combine thermal imaging cues for situational awareness with visible images and S-AIS so that agencies such as the U.S. Coast Guard can verify reported positions in S-AIS (or failure to report) with partially visible vessels in fog or darkness with expected thermal signatures of the engine systems similar to security examples noted in related research [10]. The U.S. Coast Guard has used forward looking infrared [38], specifically finding LWIR to be of most use during the day to find targets in the water with less clutter (reflection of visible light from waves) [37]. Studies related to expand use of long-wave infrared and infrared by the U.S. Coast Guard have included surface oil spill tracking [39] and through ADAC new uses are being investigated such as use in coastal erosion mapping from UAV platforms and use in Arctic ports and straits for vessel monitoring and ice tracking. Feasibility studies planned and field trials to determine value of use on USCG vessels, buoys, and light-stations for SAR (Search and Rescue) as well as port and marine security operations are planned for summer 2016 and 2017. The authors believe that the SDMSI will simplify integration into a wide range of maritime security and safety applications in open water and in ports, but more field tests with the SDMSI is needed to fully validate this hypothesis. Low-cost, convenience (self-powered) and sensor fusion image processing with S-AIS in both visible and LWIR along with potential for UAV and buoy deployments based on low-power smart sensors has invigorated renewed interest in infrared [37].

Related Arctic Missions

For SAR and mapping missions the SDMSI will operate with continuous data collection for shorter periods of time with battery power, most likely on UAVs or light aircraft. Operating off of battery power significantly limits mapping time. Today, the flight duration of quad-copters and other battery powered UAVs is typically tens of minutes, such that the payload operational limits (with a dedicated payload battery) are not the overall mission duration limit. For longer duration, higher altitude UAVs and emerging fuel cell hydrogen powered electric UAV systems, the duration of payload operation must also significantly improve, to hours of operation to match new UAV technologies.

For example, the authors intend to map coastal erosion regions in Alaska to produce digital elevation maps in summer 2016 or 2017. The goal of field testing is to prove out value not only for the primary security and safety uses, but also for parallel scientific surveys, such as marine environmental studies in the Arctic, so the SDMSI has been designed to allow for software configuration of the Linux-based camera system to work in detection and track and identification modes as an extension to basic fusion and 3D mapping. Finally, for security, the camera will most likely be mounted on coastal towers or on vessels and provide periodic surveillance to detect targets and track and identify targets only after visual saliency is determined for targets of interest. For fixed mounts, the camera is being designed to operate on a limited power budget for long duration unattended operation up to one year, potentially running off of fuel cell power rather than batteries. For all of the intended missions, the key is low power consumption in continuous transform modes

as well as low idle and standby power use. Parallel development and integration testing with off-grid power solutions including batteries, but also renewable and fuel cell sources for Arctic conditions (operation below 40 degrees Celsius) are being considered as overall SDMSI mission objectives and design challenges, as summarized in Table 1.

Table 1. Mission Objectives and Problem Summary

Challenge area	Near term Target	Long term Target	Current status
Self-powered	< 20 Watts with batteries for UAV applications or maritime/port security with rechargeable batteries	< 10 Watts with fuel cell + renewable source and one year of operation at fixed location or buoy	< 20 Watts continuous operation on bench tests
No battery operation in Arctic conditions	< 20 Watts for testing in cold chamber with hydrogen PEM fuel cell	< 10 Watts with renewable and fuel cell power to extend lifetime off-grid	< 10 Watts tested as 600 gram payload on AscTec UAV with RF link
Minimize image storage and uplink through multi-spectral fusion and saliency	30Hz image capture for 15 minute UAV flights with capture and fusion on ground laptop	30Hz fusion and selected image downlink from embedded low-power FPGA or GP-GPU processor over periods of up to one year	Visible and LWIR detectors operating at less than 1.5 Watts each (4.5 total), and co-processor operation below 10 Watts
Expansion of sensor fusion to include satellite sources such as S-AIS and satellite remote sensing multi-spectral imaging	Use of S-AIS to verify maritime vessel detection, tracking and recognition. Use of Worldview 2 & 3, NASA MODIS and USGS Landsat images to compare to low altitude remote sensing from aircraft and UAVs	Full integration with S-AIS, GPS and active LIDAR sensors for maritime security and safety. Scientific investigation of value of the SDMSI to correlate with satellite panchromatic and multi-spectral remote sensing	Preliminary field testing completed in 2015 with LTE/GSM S-AIS data (mobile phone), Linux laptop, DRS 8-14 bit LWIR camera and SLR images of marine vessels in San Pedro California and Valdez Alaska.
Active tracking of human, animal and biodiversity in marine environments	Detect and track port security threats (trespassers), dangerous animals (bear, moose) and apply to detection and tracking of marine mammals and phytoplankton.	Test on-site at the port of Anchorage and other Arctic ports (Nome, Barrow). Fly sensor over Arctic regions with significant marine mammal and biodiversity interest.	Testing at San Pedro Coast Guard facility and identification of expected marine traffic as noted by S-AIS.
Active tracking of ocean ice and tidal glacier ice	Feasibility analysis of ice tracking from aircraft and US Coast Guard vessels as available to ADAC.	Ice hazard tracking north of the Aleutian islands from buoys, US Coast Guard vessels, aircraft and UAVs.	Feasibility tested at Portage glacier and in Valdez in summer 2015 with 8-bit LWIR.
Mapping Arctic regions with coastal erosion and sea ice dangers	UAV 3D mapping missions to combine soil moisture (LWIR) with visible stereo mapping and use at fixed locations near Anchorage.	UAV and aircraft 3D mapping to combine LWIR with visible stereo mapping in Arctic regions correlated to satellite remote sensing.	Preliminary data collected on AscTec UAV in Arizona.

2. SOFTWARE-DEFINED MULTI-SPECTRAL IMAGER REQUIREMENTS

Key Requirements

Based on the mission objectives for the SDMSI related to security, safety and environmental science missions for ADAC, the derived requirements for the SDMSI have been drafted and include:

1. No-battery power system design using solar, ultra-capacitors and fuel cells in Arctic environment for long term unattended operation at fixed coastal sites and on buoys – up to one year.
2. Battery powered for short duration UAV and small aircraft use for SAR and environment mapping.
3. Design for unattended operation, drop-in-place, with intelligent mapping, detection, tracking and identification modes of operation. Opportunistic uplink of image data to vessels and UAVs in sensor networks with efficient, selective image storage with all data sources fused (S-AIS, GPS, visible, LWIR, NIR and optional LIDAR).
4. Low-cost such that use on buoys and UAVs is not prohibitive (an order of magnitude lower cost).
5. Integrated detectors with 8 bit or better dynamic range for visible and LWIR sensing.
6. Storage of all images for up to one year for opportunistic uplink in a sensor network.
7. Less than 2 kilograms for use on UAV systems including self-powering batteries.
8. Continuous resolution up and down conversion for LWIR, NIR and visible sources between 640x480 and higher 4:3 aspect ratios (based on commonly available LWIR cameras).
9. Off-load of CPUs on SDMSI for use by scene processing algorithms such as saliency.
10. Continuous transforms to register both homogenous visible cameras images (for 3D) and heterogeneous for multi-spectral fusion including pixel level registration, feature level, and pixel and feature level fusion.

3. RELATED RESEARCH

Related Research

The authors have consulted a wide range of related research based on the interdisciplinary nature of the SDMSI use for the ADAC mission objectives and prior work on sensor fusion, including the following:

1. Multi-spectral instrumentation for remote sensing [17], [18], [19] that currently provides panchromatic and infrared images from satellite systems with spatial, temporal and spectral resolution comparable to what could be collected by the SDMSI at much lower altitudes, most likely with higher resolution.
2. Sensor fusion algorithms are well known at the pixel level [10], [11], [12], [14] and general methods of pixel-level fusion based on distinct homogeneous and heterogeneous cameras in terms of spatial, spectral and temporal resolutions.
3. Environmental and physical-based fusion [1], [2], [5], [6], [13], [15] that takes into account models of target physics and spectral absorption, emission, and reflectance.
4. Security sensor fusion [4] and more specifically face recognition [3] and perimeter security [9], [16].
5. Software defined sensor networks with semi-autonomous operation [8] used in extreme environments.
6. Computation of image saliency, passive depth maps, and fusion of images [21], [22], [23], [24], [25], [26] have well-known algorithms the author's intend to use.
7. System architecture studies of low power image processing [20] using FPGAs or GP-GPUs.

While the related research is helpful, the authors did not find any single source solving the set of problems that are of interest to ADAC and meeting the requirements outlined for the SDMSI. Based on this related research, the biggest challenges for the SDMSI are heterogeneous continuous fusion and saliency determination for targets of interest, self-powering for long duration, size and mass for buoy and UAV integration, and general image handling and processing in low-power remote sensor networks used in harsh Arctic environments. The related research validates the concept that the SDMSI should largely integrate off-the-shelf detectors, optics, and co-processor solutions with emphasis on meeting spectral, spatial and temporal fusion requirements through a combination of embedded SDMSI software.

4. SYSTEM DESIGN STUDY

Continuous Infrared and Visible Image Fusion

Continuous fusion can be based upon physical features found in common between LWIR (and other infrared bands) that are in common with visible images, but this is an application specific approach that requires significant knowledge of the physical properties of the targets being observed. Such algorithms were considered for the SDMSI, but based upon the goal to support a wide range of ADAC missions, the ability to characterize targets by physical expectations of emission, absorption and reflection in various spectral bands does not generalize well, so these approaches are best used to tune specific multi-spectral fusion applications such as facial recognition [3] or remote sensing of well-known environments [1] as discussed in the previous section on related research.

The SDMSI intends to support heterogeneous LWIR and visible cameras, but with well-known extrinsic mounting differences and well-characterized differences in fields-of-view such that relatively accurate image registration between the LWIR and visible images can be achieved through geometric transformation for pixel-level registration and direct correspondence through resolution conversion (with Laplacian or Gaussian Pyramidal transforms) to determine corresponding pixels [10][11][12][14], however, it is envisioned that the SDMSI will also support feature correspondence which can serve as verification for robust registration. How well this will work must be further investigated for each proposed use of the SDMSI and solving general correspondence and registration for multi-spectral images (from a wide spectral range) remains an area that requires more research – no single general method is known for feature correspondence due to the vastly different appearance of feature in the thermal band compared to visible.

For this initial power efficiency study, the focus is therefore on pixel-level methods of correspondence and registration that make use of Pyramidal transforms for resolution matching. This study does not yet include the well-known 3D affine geometric transformations and lens radial distortion transforms that would also be needed [31] for detectors that do not share a common bore sight.

Continuous Stereo Mapping

Continuous stereo mapping requires image registration by geometric transformation to account for camera mount extrinsic differences in fields of view and to account for any intrinsic differences in the camera optics or detectors (our study uses homogeneous visible cameras mounted on a common baseline). More costly solutions might design optics with a common bore sight and pick-off mirrors to avoid the computational cost of geometric transformation. Registration with the addition of pixel-level and even feature level correspondence can be used to increase robustness of left-camera and right-camera registration as shown in Figure 1 [30].

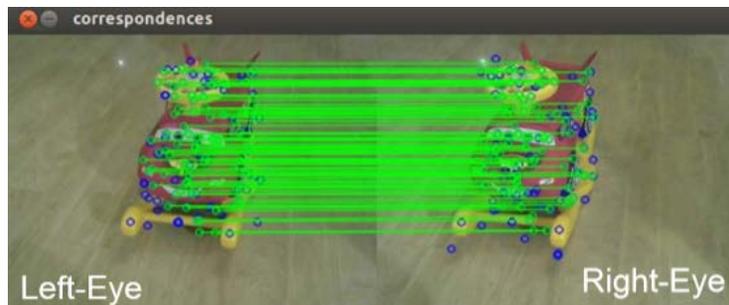


Figure 1. Test image used to find feature level correspondence with robustness for common baseline homogeneous cameras.

Based on the power efficiency and preliminary field testing completed by the authors to date, focus is on traditional pyramidal transforms, feature identification robust registration as well as pixel level, and on image enhancement. This approach will allow for transforms such as Laplacian and Gaussian pyramidal, gradient edge detection and shape-based feature extraction to be used in common for both 3D mapping and heterogeneous sensor fusion. The related research supports this hypothesis, but the authors did not find the integration of 3D and fusion combined with architecture for low-power continuous operation in previous work. The basic analysis presented here is the basis of the SDMSI architecture and the benchmarks developed could be of general use for power efficient computer vision systems.

Strategy to Define Common Transform Primitives

While the identification of common transforms required for both heterogeneous multi-spectral image fusion and common-spectrum visible stereo mapping needs more work to define the best transforms to support and benchmark for the SDMSI, we have started the process by selecting well-known potentially useful transforms to both applications.

The Sobel transform shown in Figure 2 is a simple gradient edge detector [27], which represents a feature transform (for which improved methods such as Canny edge transform could be used), but is typical of the pixel level computation needed for edge feature extraction from images ranging from LWIR (10 to 14 micron) to visible (0.39 to 0.7 micron).

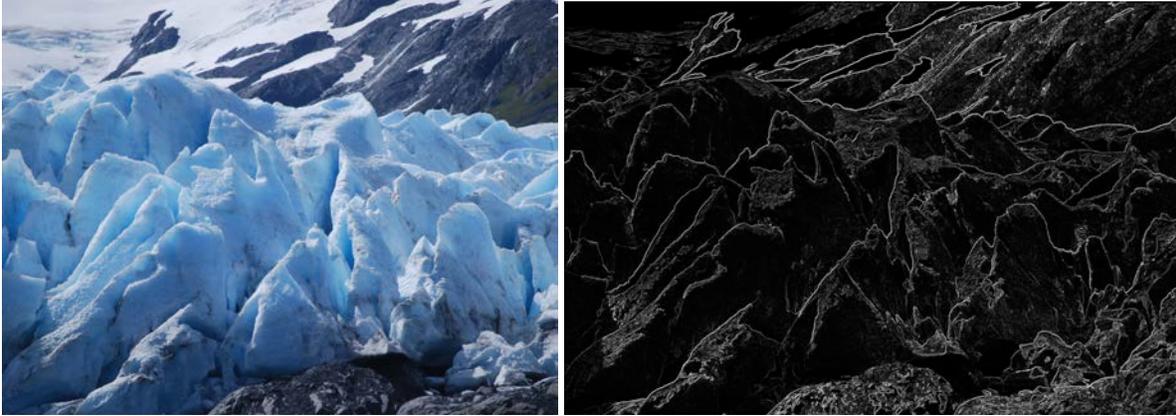


Figure 2. Test image used in Sobel edge detection and Pyramidal transforms. The source image is 1280x960 and was up-converted to 2560x1920 and down-converted to 640x480 and 320x240. The Sobel transform of this image is on the right.

More sophisticated gradient operators might be used, but would not likely change the outcome of power use by the GP-GPU or FPGA unless the operations exceeded co-processor cores available in the case of the GP-GPU or LE (Logic Elements) available on the FPGA. Similarity and use of common transforms used for 3D mapping and sensor fusion forms the basis for the significance of the analysis we present. This hypothesis that benchmarks can assist with architectural computing decisions can only be confirmed by expanding the benchmark suite and completion and testing of applications built from them. Similar work to identify common transforms for direct acceleration in co-processors is in progress by the OpenVX standards effort [32]. Ultimately, it is likely that a full suite of transform primitives for computer and machine vision can be identified so that image transform acceleration becomes much like graphics today, where all fundamental transforms are handled by custom application specific integrated circuit solution co-processors.

The Pyramidal Laplacian (resolution up-conversion) and Gaussian (resolution down-conversion) are fundamental to both multi-spectral fusion and 3D correspondence when cameras of different resolution are used, which is most common between LWIR and visible images, but useful in general for feature analysis [10], [29]. The intent is to expand the basic transforms available for the SDMSI over time and to construct both fusion and 3D stereo matching applications based on these co-processor accelerated primitives. Work is in progress on the Hough linear transform (generalized to elliptical and any template shape) is a future benchmark for feature analysis [28].

FPGA Co-processor Configuration

The FPGA co-processor tested is the Terasic DE1-SoC (Development model 1, System on Chip) which integrates an Altera28 nanometer Cyclone V FPGA with two ARM Cortex A9 central processors and 85,000 programmable logic elements. The benchmark image transforms are run such that they are transferred to the FPGA fabric via a bus interfaced buffer transfer to the logic array which transforms pixels concurrently, making multiple transfers as needed to scale resolution until the transformed image is completed. The benchmark runs under the control of a Linux software application with buffer transfer via an Altera driver, with the logic transform specified by an OpenCL kernel. The use of OpenCL is based upon the advantage of not requiring development of transforms in an HDL (Hardware Design Language) and with the long term goal that all transforms share a common OpenCL kernel format rather than porting transforms between different types of specifications unique to CPUs, GP-GPUs, and FPGAs.

GP-GPU Co-processor Configuration

The GP-GPU co-processor tested is the Jetson TK1 (embedded development kit) which integrates the NVIDIA Tegra K1 System-on-Chip with four ARM Cortex A15 central processors and 192 CUDA (Computer Unified Device Architecture) co-processor cores. The benchmark image transforms are run such that they are transferred to the CUDA cores via a bus interfaced host to device memory transfer to the co-processors which transforms tiles from the image in parallel, making multiple transfers as needed to scale resolution until the transformed image is completed. The benchmark runs under the control of a Linux software application with buffer transfer via the NVIDIA CUDA driver, with the logic transform specified by a CUDA kernel. Ideally we intended to use a common OpenCL kernel, but the Jetson does not yet support OpenCL, so we used a similar CUDA kernel. OpenCL would have the advantage of not requiring reimplementations of transforms when used on either co-processor platform and would arguably result in a more direct comparison.

5. SYSTEM DESIGN ANALYSIS

The analysis of transform power requirements by transform type (Sobel and Pyramidal), resolution, and frame rate are presented here as measured. The focus in the analysis is on the co-processors and not on the main CPU cores. The implementation of the Altera Cyclone V SoC and NVIDIA Tegra K1 benchmarks tested here can be found on GitHub at https://github.com/siewertserau/fusion_coproc_benchmarks.

For this study, we have chosen to ignore the idle power and small power increases associated with scheduling transform kernel execution on the co-processors. Rather than attempting to configure the two SoCs to have identical CPU configuration power characteristics, we simply subtract off the idle and the continuous standby transform power use. A simple DMM (Digital Multi-meter) was used that store peak and average power used over time, with a continuous workload of transformation of a single in-memory image. Power efficiency studies for the NVIDIA Tegra K1 include more detail on lowering idle power while using the GP-GPU stream co-processors [33]. The most relevant blocks used in our study are shown in Figure 3 [41].

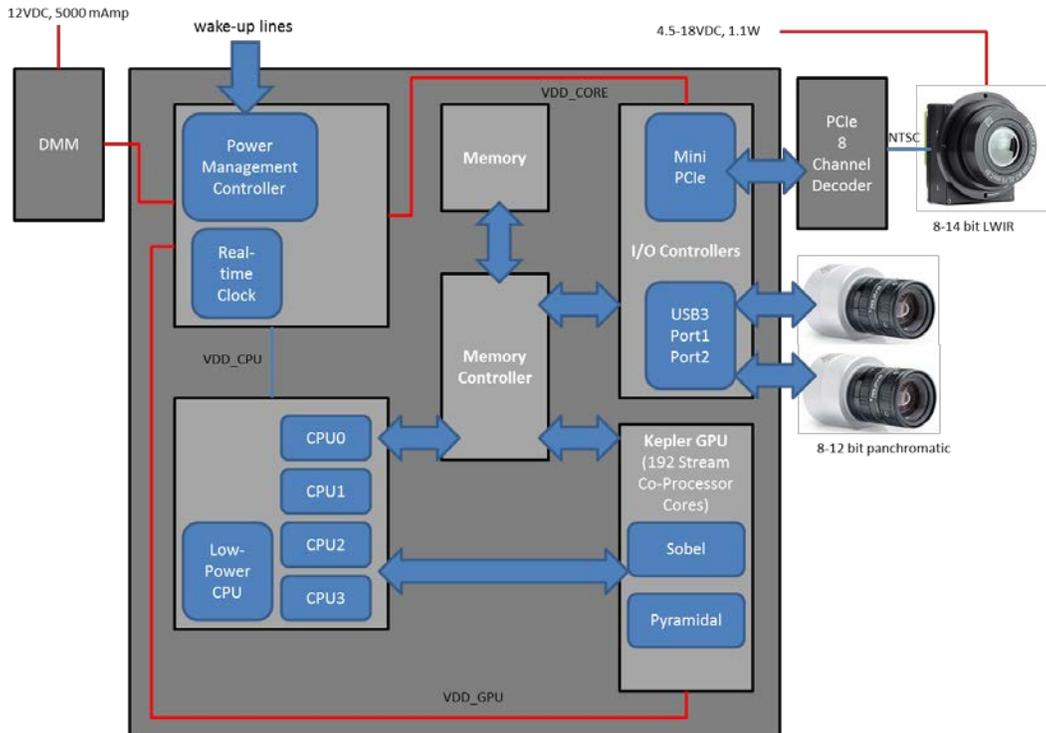


Figure 3. Primary component of interest system level block diagram for NVIDIA Tegra K1 SoC.

Likewise, for the Cyclone V SoC used, we again did not attempt to reduce idle power or power consumed while running the continuous load transforms. Previous work using FPGA SoCs and combined microprocessor and FPGA fabric designs for continuous encode, decode and transform have shown the value of using an FPGA as a co-processor in terms of energy used per pixel [34][35]. The major elements most important to our study are shown in Figure 4 [40]. Again, we chose to ignore the idle power and background power used by the CPU cores and peripherals by characterizing it and subtracting off to estimate the power used by the FPGA fabric for the transform co-processing alone. Total system power used was measured because we will run from batteries and total system power is ultimately of primary interest to our research rather than measuring power consumed by the co-processor alone (breaking down power use more specifically might be of interest in future work). As our project progresses, we do plan to work on reducing idle power (to reduce battery draw down) and power consumed by CPU cores for participation in the sensor network and higher level scene analysis and processing such as saliency metrics for images collected [26].

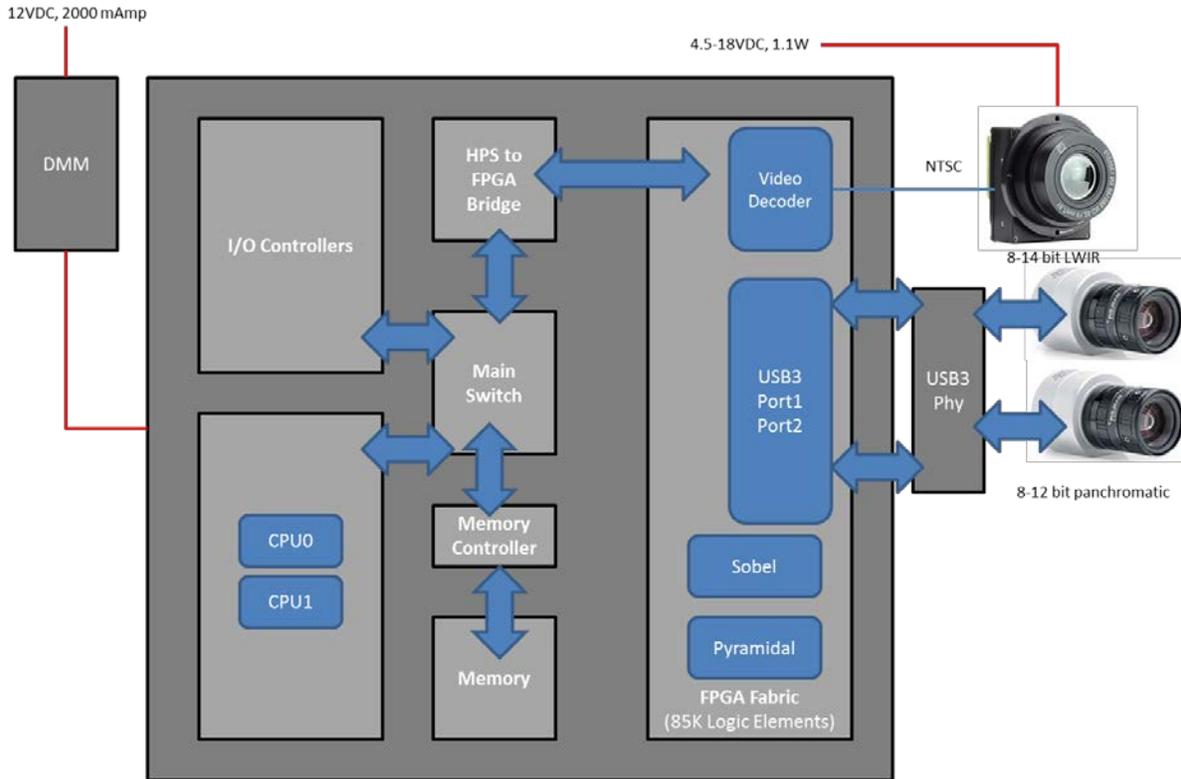


Figure 4. Primary component of interest system level block diagram for Altera Cyclone V SoC FPGA (SCSEMA5F31C6N).

FPGA Co-processor Continuous Power Use by Transform

Using the Cyclone V SoC to offload a Sobel gradient edge transform, the standby power draw (from CPU cores and all other sources) was 4.9 Watts on average. Attempts were not made to disable non-essential power loads. The test was completed with the SoC in a power state typical of running Linux and sequencing transforms from a memory-resident grayscale pixel image. For the FPGA, the power used for continuous operation is constant, but is a function of the size of the image that is presented to the FPGA fabric transform through the HPS to FPGA Bridge and the Main Switch (L3 in Altera technical documentation) [40]. Essentially the power used is a function of the logic elements used by the OpenCL compiler and the synthesis into logic by the Quartus mapping to the fabric and the AMBA AIX bus transfer rate. For the Sobel, the following report was provided on use of the fabric:

ALUTs: 10187
 Registers: 13,561
 Logic utilization: 7,427 / 32,070 (23 %)

Based on the logic utilization, it is clear that creating parallel transforms for tiles of the larger resolution images would be a possible way to increase throughput. This was not done at time of publication, but is now being pursued to verify that the FPGA fabric can scale by tiling and replicating kernels. The constant power, pixel rate per Watt and bus transfer rate indicates that the mapping used relatively constant pixel width for input. As noted in Altera documentation, the mapping and the sizing is automatic regardless of OpenCL kernel directives such as working set size [42]. The power consumption is summarized in Table 2. The advantage of the FPGA fabric compared to the GP-GPU fabric is constant power consumption for continuous transform based on logic utilization and bus transfer rates alone. With tiling of higher resolution images into equivalent lower resolution tiles presented to multiple kernels, it should be possible to scale Sobel on the Cyclone V SoC to more than 8 million pixels/sec per Watt.

Table 2. Sobel Continuous Transform Power Consumption by Cyclone V FPGA

Resolution	Transform (Watts)	(Pixel/sec) / Watt	Saturation FPS	Bus transfer rate (MB/sec)
320x240	5.655	2,050,716	151	11.06
640x480	5.700	2,107,284	39.1	11.46
1280x960	5.704	2,143,506	9.95	11.66
2560x1920	5.696	2,157,303	2.50	11.72

Using the Cyclone V SoC to offload a Pyramidal Laplacian transform for resolution up-conversion and Pyramidal Gaussian transform for resolution down-conversion, the standby power draw (from CPU cores and all other sources) was 5.2 Watts for both because both transforms were loaded at the same time, but only one was tested at a time. For the Pyramidal up-conversion from lowest resolution of 320x240 and down-conversion from the highest resolution of 2560x1920, the following report was provided on use of the fabric for both contained in one kernel:

ALUTs: 24456

Registers: 34,062

Logic utilization: 17,721 / 32,070 (55 %) (55 %)

Given higher FPGA utilization, this would make use of the fabric for USB3 and an analog video decoder shown in Figure 4 a challenge to integrate along with the Pyramidal transform as well as other transforms of importance to 3D mapping such as affine transforms for visible camera extrinsic mount differences. The power consumption for Pyramidal transforms (both up and down-conversion) is summarized in Table 3 and 4.

Table 3. Pyramidal Laplacian Resolution Up-Conversion Continuous Transform Power Consumption

Resolution	Transform (Watts)	(Pixel/sec) / Watt	Saturation FPS	Bus transfer rate (MB/sec)
320x240	6.009	889,546	69.6	5.10
640x480	6.013	904,281	17.7	5.19
1280x960	6.038	905,624	4.45	5.21
2560x1920	6.192	889,054	1.12	5.25

Table 4. Pyramidal Gaussian Resolution Down-Conversion Continuous Transform Power Consumption

Resolution	Continuous Transform Power (Watts)	(Pixel/sec) / Watt	Saturation FPS	Bus transfer rate (MB/sec)
320x240	5.968	2,445,040	190	13.92
640x480	6.018	2,399,202	47.0	13.77
1280x960	6.023	2,427,813	11.9	13.95
2560x1920	6.109	2,309,154	2.87	13.45

The constant power, pixel rate per Watt and bus transfer rate indicates that the mapping used a relatively constant pixel width for logic transform input similar to Sobel, but using a majority of available logic. As noted in the Altera documentation, the mapping and the sizing is automatic regardless of OpenCL kernel directives such as working set size.

GP-GPU Co-processor Continuous Power Use by Transform

Using the 192 stream co-processors found on the Tegra K1 SoC Kepler GP-GPU to offload a Sobel gradient edge transform, the standby power draw (from CPU cores and all other sources) was 2.96 Watts on average. Attempts were not made to disable non-essential power loads. The test was completed with the SoC in a power state typical of running Linux and sequencing transforms from a memory-resident grayscale pixel image. For the GP-GPU, the power used for continuous operation is variable. The power consumption is summarized in Table 5 and in Figure 5

Table 5. Sobel Continuous Transform Power Consumption by Tegra K1 GP-GPU

Resolution	Continuous Power at 1Hz (Watts)	Continuous Power at 30Hz (Watts)	(pixels/sec) / Watt @ 1Hz	(pixels/sec) / Watt @ 30Hz	Saturation FPS
320x240	4.241	4.932	18,109	467,153	1624
640x480	4.256	4.984	72,180	1,849,117	840
1280x960	4.266	5.142	288,045	7,169,195	237
2560x1920	4.325	7.326	1,136,462	20,127,764	55

An advantage of the GP-GPU over FPGA fabric is that stream co-processors idle quickly and reduce overall power used.

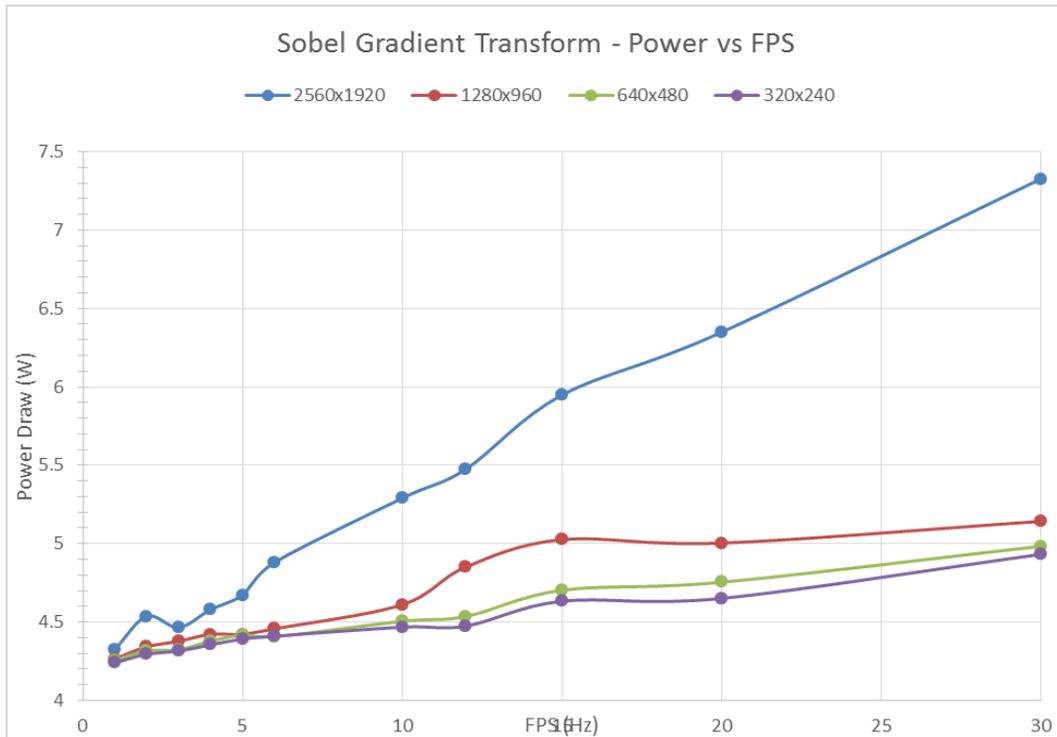


Figure 5. The total power used (including standby of 2.96 Watts) for Sobel transform up to 30Hz.

Using the Tegra K1 SoC Kepler GP-GPU to offload Pyramidal Laplacian resolution up-conversion followed by Pyramidal Gaussian down-conversion in a single test, the results hit saturation for the highest starting and ending

resolution compared to the lowest. This benchmark is different from the FPGA benchmark which loads both up-conversion and down-conversion kernels, but only exercises conversion in one direction at a time. Despite the more rigorous application of conversion both up and down the resolution pyramid, the GP-GPU still scales better than the FPGA. The power consumption is summarized in Table 6 and Figure 6.

Table 6. Pyramidal Up and Down Conversion Continuous Transform Power Consumption by Tegra K1 GP-GPU

Resolution	Continuous Power at 1Hz (Watts)	Continuous Power at 20Hz (Watts)	(pixels/sec) / Watt @ 1Hz	(pixels/sec) / Watt @ 20Hz	Saturation FPS
320x240	4.104	4.824	18,713	477,612	1120
640x480	4.116	5.460	74,636	1,687,912	325
1280x960	4.152	6.864	295,954	5,370,629	82
2560x1920	4.224	13.44	1,163,636	10,971,429	20

The GP-GPU first does an up-conversion transform to the next highest 4:3 aspect ratio and then back through a down-conversion to the original resolution in the same test. For example, from 320x240 down to 160x120 and back to 320x240. So 2560x1920 is taken down to 1280x960 and back to 2560x1920 and so forth for each resolution tested.

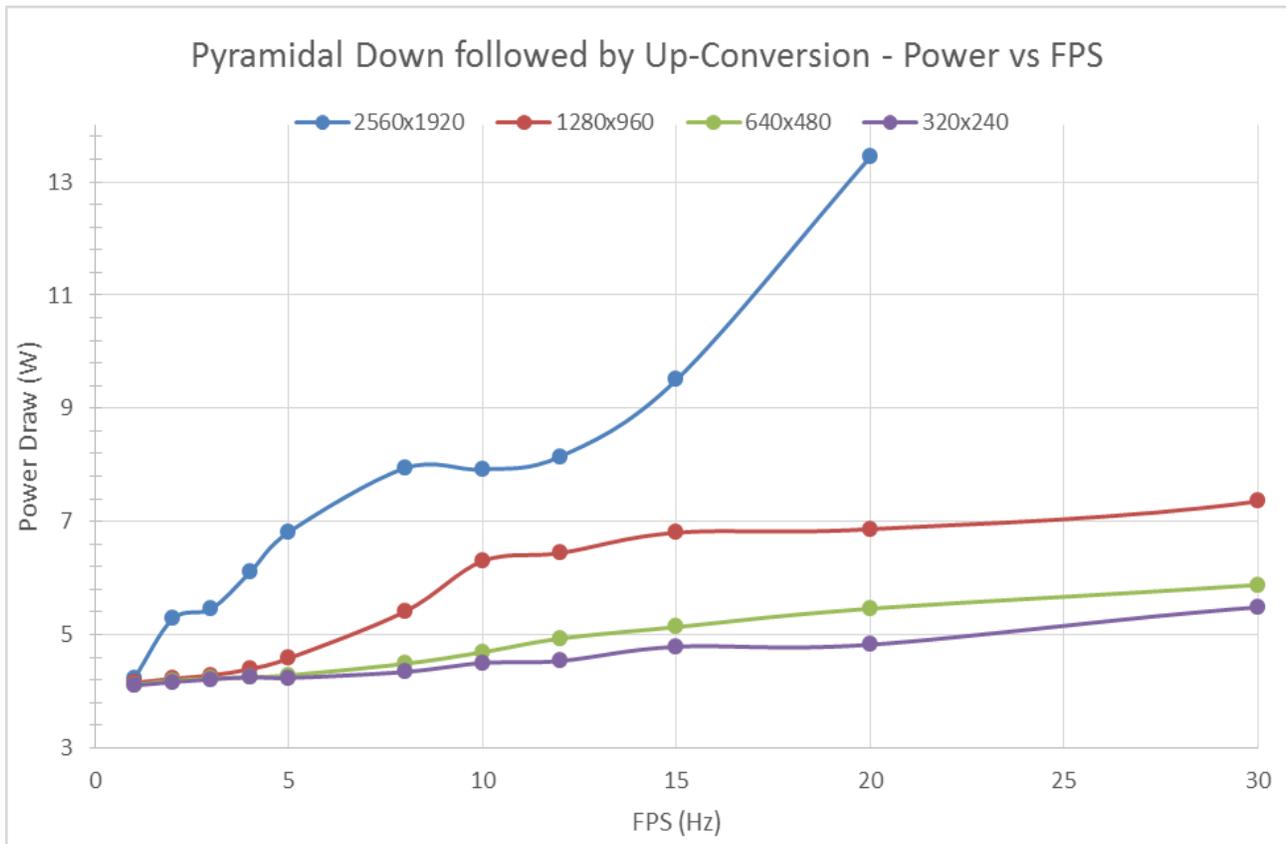


Figure 6. The total power (standby of 2.96 Watts) for Pyramidal up and down conversion up to 30Hz or saturation.

Comparative Co-processor Continuous Power Use by GP-GPU and FPGA

General observations about power use by the FPGA compared to the GP-GPU:

1. The FPGA in the Cyclone V SoC has an order of magnitude lower throughput compared to the Tegra K1 (for example, 151 frames/sec for Sobel 320x240 for the FPGA compared to 1624 for the GP-GPU).
2. The FPGA has zero power gain as a function of frame rate and the GP-GPU has a mostly linear power gain for simple transforms like the Sobel, but much less linear for more complex transforms like Pyramidal.
3. The GP-GPU power gain is a function not only of frame rate, but also based on resolution and has significant gain as a function of both, whereas the FPGA power draw is not affected by frame rate at all (but does saturate) and only by resolution to a small degree.
4. The GP-GPU has much higher peak power consumed than the FPGA.

Overall, the Tegra K1 meets the SDMSI requirements better than the Cyclone V FPGA because it provides better scaling up to the desired frame rate of 30Hz for the transforms of interest that are in common to both 3D mapping and heterogeneous fusion. While it requires more power and has more power gain than the FPGA, for the SDMSI application, it meets all requirements up to 1280x960. With work on reduction of the idle and standby power consumed by the GP-GPU it would seem feasible that 2560x1920 could be achieved at 30Hz based on only needing to up-convert or down-convert one of the two images in fusion.

6. FUTURE WORK

Higher Level Fusion Application Testing

Next steps for the ADAC SDMSI project following benchmarking is to compose applications for fusion (LWIR and visible) as well as 3D mapping (two homogeneous megapixel cameras) and to test with integrated cameras. Presently the project is using five megapixel Basler cameras for 3D visible with USB3 interfaces to the Jetson and USB2 to the DE1-SoC and FLIR (and DRS) 640x480 analog LWIR cameras to interface to both the Jetson and DE1-SoC. Fusion will be tested to combine a resolution transformed visible image with the 640x480 LWIR image with required geometric, lens and field-of-view transformations required. The goal is to test the capability in Anchorage this summer 2016 in a security application at University of Alaska, but to also pursue UAV testing at Embry Riddle Aeronautical University in Prescott.

3D Mapping Application Testing

Follow-on work for 3D mapping mostly involves registration and computation of disparity images and depth maps using two homogenous cameras (avoiding need for Pyramidal conversion at this point). Pyramidal conversion could be useful to reduce resolution and longer term for feature correspondence at lower resolution.

Combined 3D Mapping with Multi-spectral Fusion

Significant value of the SDMSI is seen to come from the ability to share common transforms and to provide fusion, 3D mapping or combined fusion and 3D mapping. Target applications for this include soil erosion mapping from a UAV to stereo map coastal topography and to combine with soil moisture estimates derived from LWIR imaging. A related goal is to fly a camera over the Arctic Ocean or mast mount to collect ocean surface height and color to compare with remote sensing of ocean color from satellite remote sensors. Comparison to satellite remote sensors is planned to include Worldview 2, 3 [17], MODIS [18] and Landsat [19]. These satellite remote sensing systems cover similar 0.39 micron to 14 micron spectral ranges (with more spectral resolution), but with less spatial and temporal resolution.

Image Saliency and Ranking

Image saliency methods have been researched and the plan is to make use of the off-loaded CPU cores (4 on the Jetson and 2 on the DE1-SoC) to process higher level scene algorithms such as saliency and target recognition. The goal of off-loading common transforms to the FPGA or GP-GPU co-processor is largely based upon efficiency and making the CPU available for higher level uses.

7. SUMMARY

In general the FPGA has essentially fixed power consumption for any given transform and image resolution with no gain as a function of frame rate. By comparison the GP-GPU is quite variable in terms of power used as a function of frame rate, by transform type and by resolution, but has much higher throughput at similar power levels.

Overall, software defined instrumentation such as the SDMSI can add value to security and environmental remote sensing applications at low cost by leveraging mostly off-the-shelf optics and detectors ranging from long-wave, short-wave, and near infrared to the visible spectrum if these heterogeneous cameras can be integrated through continuous coordinated image processing. Use of a co-processor such as an FPGA or GP-GPU frees up the main SDMSI processor. The work presented provides a starting point and useful suite of benchmarks that can be used for the ADAC mission objectives for the SDMSI, but can also be of use to a broader range of sensor fusion and smart camera applications in general. As optics, detector technology, computing and standards for integrating these devices progresses, it is expected that many instrument designs will become mostly defined by the algorithms of interest and applications and therefore share much in common with the SDMSI presented.

REFERENCES

- [1] Dominguez, A., Kleissl, J., Luvall, J. C., Rickman, D. L. "High-resolution urban thermal sharpener (HUTS)," *Remote Sensing of Environment*, 115(7), 1772-1780 (2011).
- [2] Hines, G. D., Rahman, Z. U., Jobson, D. J., & Woodell, G. A. (2003, August). Multi-image registration for an enhanced vision system. In *AeroSense 2003* (pp. 231-241). International Society for Optics and Photonics.
- [3] Gyaourova, A., Bebis, G., & Pavlidis, I. (2004). Fusion of infrared and visible images for face recognition. In *Computer Vision-ECCV 2004* (pp. 456-468). Springer Berlin Heidelberg.
- [4] Kriesel, J. M., & Gat, N. (2010, April). True-color night vision (TCNV) fusion system using a VNIR EMCCD and a LWIR microbolometer camera. In *SPIE Defense, Security, and Sensing XIX*, 7697. International Society for Optics and Photonics.
- [5] Cubero-Castan, M., Chanussot, J., Achard, V., Briottet, X., & Shimoni, M. (2015). A physics-based unmixing method to estimate subpixel temperatures on mixed pixels. *Geoscience and Remote Sensing, IEEE Transactions on*, 53(4), 1894-1906.
- [6] Agam, N., Kustas, W. P., Anderson, M. C., Li, F., & Neale, C. M. (2007). A vegetation index based technique for spatial sharpening of thermal imagery. *Remote Sensing of Environment*, 107(4), 545-558.
- [7] Siewert, S. B., Shihadeh, J., Myers, R., Khandhar, J., & Ivanov, V. (2014, May). Low-cost, high-performance and efficiency computational photometer design. In *SPIE Sensing Technology+ Applications*, 9121. International Society for Optics and Photonics.
- [8] Thompson, D. R., Allwood, A. C., Bekker, D. L., Cabrol, N. A., Fuchs, T., & Wagstaff, K. L. (2012, March). TextureCam: Autonomous image analysis for astrobiology survey. In *Lunar and Planetary Science Conference* (Vol. 43, p. 1659).
- [9] Liu, Z. (2010). *Investigations on multi-sensor image system and its surveillance applications*. Universal-Publishers.
- [10] Piella, G. (2003). A general framework for multiresolution image fusion: from pixels to regions. *Information fusion*, 4(4), 259-280.
- [11] Blum, R. S., & Liu, Z. (Eds.). (2005). *Multi-sensor image fusion and its applications*. CRC press.
- [12] Liu, Z., Blasch, E., Xue, Z., Zhao, J., Laganriere, R., & Wu, W. (2012). Objective assessment of multiresolution image fusion algorithms for context enhancement in night vision: a comparative study. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 34(1), 94-109.
- [13] Simone, G., Farina, A., Morabito, F. C., Serpico, S. B., & Bruzzone, L. (2002). Image fusion techniques for remote sensing applications. *Information fusion*, 3(1), 3-15.
- [14] Mitchell, H. B. (2010). *Image fusion: theories, techniques and applications*. Springer Science & Business Media.
- [15] Alparone, L., Aiazzi, B., Baronti, S., Garzelli, A. (2015). *Remote Sensing Image Fusion*. Signal and Image Processing of Earth Observations, CRC Press.
- [16] Szustakowski, M., Ciurapinski, W. M., Zyczkowski, M., Palka, N., Kastek, M., Dulski, R., & Sosnowski, T. (2009, September). Multispectral system for perimeter protection of stationary and moving objects. In *SPIE Europe Security+ Defence 7481*. International Society for Optics and Photonics.
- [17] Apollo Mapping Inc. (<https://apollomapping.com/>), McCarty, B.A., Nelson, K., (2016), "Image Hunter," <https://imagehunter.apollomapping.com/>, Boulder, Colorado, USA.
- [18] National Aeronautics and Space Administration, Moderate Resolution Imaging Spectroradiometer (<http://modis.gsfc.nasa.gov/>), Maccherone, B., Frazier, S. (2016), "Data," <http://modis.gsfc.nasa.gov/data/>, NASA Earth Science Division and NASA Goddard Space Flight Center, Greenbelt, Maryland, USA.

- [19] United States Geological Survey Landsat Missions (<http://landsat.usgs.gov/>), (2016), "Earth Explorer," <http://earthexplorer.usgs.gov/>, United States Department of the Interior, USA.
- [20] Miller, D.W., (July 2015). *2015 NASA Technology Roadmaps, TA4: Robotics and Autonomous Systems*. National Aeronautics and Space Administration (<http://www.nasa.gov>), Office of the Chief Technologist, (<http://www.nasa.gov/offices/oct/home/roadmaps/index.html>).
- [21] Sharma, G., Jurie, F., & Schmid, C. (2012, June). Discriminative spatial saliency for image classification. In *Computer Vision and Pattern Recognition (CVPR), 2012 IEEE Conference on* (pp. 3506-3513). IEEE.
- [22] Toet, A. (2011). Computational versus psychophysical bottom-up image saliency: A comparative evaluation study. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 33(11), 2131-2146.
- [23] Valenti, R., Sebe, N., & Gevers, T. (2009, September). Image saliency by isocentric curvedness and color. In *Computer Vision, 2009 IEEE 12th International Conference on* (pp. 2185-2192). IEEE.
- [24] Wang, M., Konrad, J., Ishwar, P., Jing, K., & Rowley, H. (2011, June). Image saliency: From intrinsic to extrinsic context. In *Computer Vision and Pattern Recognition (CVPR), 2011 IEEE Conference on* (pp. 417-424). IEEE.
- [25] Liu, F., & Gleicher, M. (2006, July). Region enhanced scale-invariant saliency detection. In *Multimedia and Expo, 2006 IEEE International Conference on* (pp. 1477-1480). IEEE.
- [26] Cheng, M. M., Mitra, N. J., Huang, X., & Hu, S. M. (2014). Salient shape: Group saliency in image collections. *The Visual Computer*, 30(4), 443-453.
- [27] Maini, R., & Aggarwal, H. (2009). Study and comparison of various image edge detection techniques. *International journal of image processing (IJIP)*, 3(1), 1-11.
- [28] Duda, R. O., & Hart, P. E. (1972). Use of the Hough transformation to detect lines and curves in pictures. *Communications of the ACM*, 15(1), 11-15.
- [29] Ranchin, T., & Wald, L. (2000). Fusion of high spatial and spectral resolution images: the ARSIS concept and its implementation. *Photogrammetric Engineering and Remote Sensing*, 66(1), 49-61.
- [30] Boyer, K. L., & Kak, A. C. (1988). Structural stereopsis for 3-D vision. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 10(2), 144-166.
- [31] Szeliski, R. (2010). *Computer vision: algorithms and applications*. Springer Science & Business Media.
- [32] Tagliavini, G., Haugou, G., Marongiu, A., & Benini, L. (2015, June). A framework for optimizing OpenVX applications performance on embedded manycore accelerators. In *Proceedings of the 18th International Workshop on Software and Compilers for Embedded Systems* (pp. 125-128). ACM.
- [33] Stokke, K. R., Stensland, H. K., Griwodz, C., & Halvorsen, P. (2015, March). Energy efficient video encoding using the tegra K1 mobile processor. In *Proceedings of the 6th ACM Multimedia Systems Conference* (pp. 81-84). ACM.
- [34] De La Piedra, A., Braeken, A., & Touhafi, A. (2012). Sensor systems based on FPGAs and their applications: A survey. *Sensors*, 12(9), 12235-12264.
- [35] Genovese, M., & Napoli, E. (2014). ASIC and FPGA implementation of the gaussian mixture model algorithm for real-time segmentation of high definition video. *Very Large Scale Integration (VLSI) Systems, IEEE Transactions on*, 22(3), 537-547.
- [36] Eriksen, T., Høye, G., Narheim, B., & Meland, B. J. (2006). Maritime traffic monitoring using a space-based AIS receiver. *Acta Astronautica*, 58(10), 537-549.
- [37] Krapels, C. K., Driggers, C. R. G., & Garcia, C. J. F. (2007). Performance of infrared systems in swimmer detection for maritime security. *Optics express*, 15(19), 12296-12305.
- [38] Hover, G., Mazour, T., Osmer, S., & Nash, L. (1982, September). Evaluation of forward looking infrared (FLIR) as a coast guard SAR sensor. In *OCEANS 82* (pp. 491-495). IEEE.
- [39] Allen, J., & Walsh, B. (2008, May). Enhanced oil spill surveillance, detection and monitoring through the applied technology of unmanned air systems. In *International oil spill conference* (Vol. 2008, No. 1, pp. 113-120). American Petroleum Institute.
- [40] Altera Inc., (November 2015). *Cy 5v4 Cyclone V Hard Processor System Technical Reference Manual*. Altera Cyclone V SoCs (<https://www.altera.com/products/soc/portfolio/cyclone-v-soc/overview.html>), Quartus 15.1.
- [41] NVIDIA Inc., (October 2014). *Technical Reference Manual – NVIDIA Tegra K1 Mobile Processor*. DP-06905-001_v03p.
- [42] Altera Inc., (November 2015). *UG-OCL003, Altera SDK for OpenCL Best Practices*.
- [43] Bradski, G., & Kaehler, A. (2008). *Learning OpenCV: Computer vision with the OpenCV library*. O'Reilly Media, Inc.