A Streaming Sensor Challenge Problem for Ubiquitous High Performance Computing

Daniel P. Campbell, Daniel A. Cook, Brian P. Mulvaney
Georgia Tech Research Institute / Sensors and Electromagnetic Applications Laboratory
{dan.campbell, dan.cook, brian.mulvaney}@gtri.gatech.edu

Abstract

This presentation will describe a Challenge Problem (CP) for petascale computers that models the requirements of computationally intensive, streaming sensor missions. The Challenge Problem is under development to support DARPA's Ubiquitous High Performance Computing (UHPC) program, and models the requirements of a persistent surveillance mission with high throughput, real-time requirements.

Introduction

DARPA's UHPC program is developing computing technology and systems capable of delivering at least one petaflop of peak performance while drawing no more than 57 KW and occupying no more than one standard server rack of space. UHPC systems must also be modular and downward scalable, in order to meet the needs of a wide range of deployable defense applications. In order to drive the development of architectures that are capable of accomplishing defense-relevant missions, the Georgia Tech Research Institute (GTRI) leads a team[1] developing several scalable CPs for the UHPC program. In order to specify a computing load that is challenging and relevant, while allowing flexibility for innovation, the CPs are defined mathematically rather than by means of an executable benchmark. In most cases, the choice of specific algorithms is left to the discretion of the implementer, so long as the same overall operations are achieved.

The Streaming Sensor Challenge Problem (SSCP) models the requirements of a wide-area, high resolution persistent surveillance mission. The problem is based around radar image formation and analysis for knowledge extraction. The CP represents the processing operations required to transform a stream of inputs from a sensor suite associated with a radar into a set of possible detections of moving objects for further tracking and analysis.

Description

In the mission modeled by this Challenge Problem, an airborne radar system flies a repeated path around a target area to be observed, as illustrated in Figure 1. The radar illuminates the target area with regular, repeated radar pulse. The reflected returns from each pulse are down-converted to a baseband complex signal, pulse-compressed to provide fine range resolution, and then sampled at a number of points in time. The primary sensor inputs for the challenge problem are these pulse-compressed complex return samples; the time of transmission of each pulse; the sampling times for the return from each pulse; and the position of the transmitter and receiver at the transmission time of each pulse.

The computing system must then use the inputs to form a series of Synthetic Aperture Radara (SAR) images of a specified size, by means of a fully general backprojection algorithm. At the implementer's option, each image may be formed as a single full-size image, or by tiling a group of smaller subimages. If tiling is elected, a digital spotlighting step is employed to form a reduced-size data set which is then used to form the subimages using backprojection. The optional tiling can change the total number of operations required to form the image, and also changes the balance of operation types. While other, less computationally demanding, SAR image formation algorithms, such as the Polar Format Algorithm (PFA)[1,2] and omega-k[2,3] are known, they possess shortcomings that make them unsuitable for the wide area surveillance mission modeled by this CP. Such missions require short standoffs and wide beams, violating the tomographic assumption that is required for PFA to be suitable. Omega-k is unsuitable for this scenario, because the algorithm requires motion compensation to a line leading to dramatically increased along-track sampling rates.

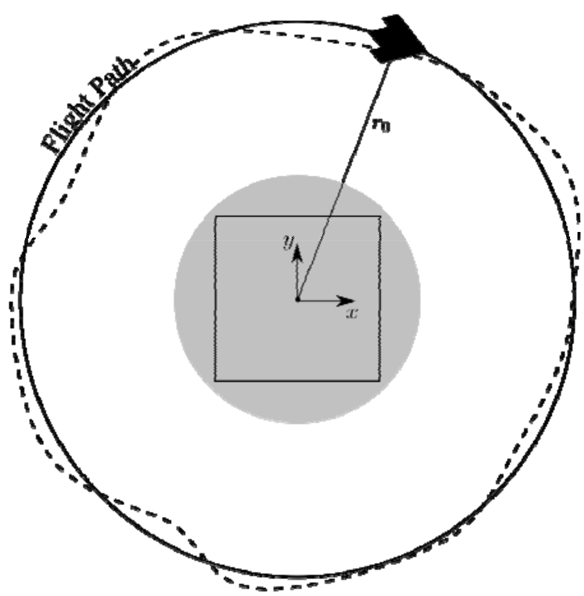

Figure 1: The sensor is flown in an orbit that is nominally circular (solid line) around the scene of interest, as indicated by the dashed line. The circular region in the center represents the region of ground visible to the radar, while the box indicates the area actually imaged.

[1] This work was supported by DARPA under contract HR0011-10-C-0145. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressly or implied, of the Defense Advanced Research Projects Agency or the U.S. Government.

Successive full-size images must be constructed at a specified cadence using overlapping subsets of pulses. Those images taken from the same nominal position in the orbit from consecutive orbits must then be registered via a two stage process comprising a global affine transformation (AT) for coarse global registration followed by a thin plate spline warping (TSW) for fine local registration. Once registered, coherent change detection (CCD) between successive images must be applied. CCD must be implemented as a pixel-by-pixel coherence estimate followed by a constant false alarm rate (CFAR) algorithm to identify pixel locations of significant change. These detected change locations form the output tokens for the Streaming Sensor Challenge Problem.

Discussion

The SSCP is defined in terms both of the functionality required, and also the set of parameters that scale the amount of computation, data movement, and resources required to meet the CP requirements. We have defined an initial set of four scenarios with specific settings for these parameters. The parameters, along with the settings for the four initial scenarios are summarized in Table 1, below.

Scenario:	**1**	**2**	**3**	**4**
Ground Area (edge size, m)	610	1,219	2,438	9,753
Image Size (edge size, pixels)	4,000	8,000	16,000	64,000
Pulses per Image	4,800	9,600	19,200	76,800
Samples per Pulse	4,000	8,000	16,000	64,000
Pulses per Second	1084			
Images per second	1			
Affine registration control points	3,629	14,513	58,050	928,799
Thin-spline registration control points	1,452	1,452	1,452	1,452
CCD neighborhood	5x5			
CFAR size	15x15			

Table 1 - SSCP Scenarios

The scenario parameters were selected to provide a challenging problem that requires effective use of resources to complete within the required cadence, while still reflecting the balanced requirements of the mission modeled by the CP. The initial set of scenario parameters, described in Table 1, create a loading that is approximated in Table 2, below.

Because the CP specification allows the implementer the choice of whether to use digital spotlighting (DS) at all, and, if used, the choice of what tile size to create, the actual computational requirements of an implementation of the CP are not fixed. The backprojection image formation (BPIF) step requires fewer operations to form the entire image with smaller sub-image tiles, but the digital spotlight step requires more operations with greater number of tiles.

	1	**2**	**3**	**4**
DS	90.8×10^{9}	1.47×10^{12}	6.27×10^{12}	442×10^{12}
BPIF	326×10^{9}	1.31×10^{12}	10.4×10^{12}	334×10^{12}
AT	1.26×10^{9}	5.04×10^{9}	20.2×10^{9}	323×10^{9}
TSW	469×10^{9}	1.86×10^{12}	7.44×10^{12}	119×10^{12}
CCD	100×10^{9}	400×10^{9}	1.60×10^{12}	25.6×10^{12}
CFAR	900×10^{6}	3.6×10^{9}	14.4×10^{9}	230×10^{9}
Total	990×10^{9}	5.05×10^{12}	25.8×10^{12}	922×10^{12}
Input (bps)	277×10^{6}	555×10^{6}	1.11×10^{9}	4.44×10^{9}

Table 2 - Computational requirements for each SSCP scenario

The operation counts shown in Table 2 correspond to those required for selecting the tile sizes that minimize operation counts. In real systems, operation count is rarely the dominant criteria to impact execution speed. Since the two steps are composed of operations with vastly different memory access patterns and input bandwidth requirements, it is likely that implementers will select the tile sizes based on the balance of capabilities on a particular platform.

This presentation will describe the CP in greater detail, including the specific requirements of each individual step. We will also present additional analysis of the computing resources required to complete the CP for the various scenarios, and discuss the early results of reference implementations of the CP.

References

[1] Jakowatz et al., *Spotlight-Mode Synthetic Aperture Radar: A Signal Processing Approach*, Springer, 1996.

[2] Carrara et al., *Spotlight Synthetic Aperture Radar: Signal Processing Algorithms*, Artech, 1995.

[3] H. J. Callow, *Signal Processing for Synthetic Aperture Sonar Image Enhancement*. PhD thesis, Department of Electrical and Electronic Engineering, University of Canterbury, Christchurch, New Zealand, 2003.