Analysis of Computational System Performance in Automatic Target Recognition

Joseph A. O'Sullivan
Mark A. Franklin
Michael D. DeVore
Roger D. Chamberlain

Joseph A. O'Sullivan, Mark A. Franklin, Michael D. DeVore, and Roger D. Chamberlain, "Analysis of Computational System Performance in Automatic Target Recognition," in *Proc. of 4th High Performance Embedded Computing Workshop*, September 2000.

Department of Electrical Engineering
Washington University
St. Louis, Missouri

Analysis of Computational System Performance in Automatic Target Recognition*

Joseph A. O'Sullivan[1], Mark A. Franklin[2], Michael D. DeVore[1], Roger D. Chamberlain[2]
[1]Electronic Systems and Signals Research Laboratory
[2]Computer and Communications Research Center
Department of Electrical Engineering, Washington University, St. Louis MO 63130

1 Introduction

An important issue in the choice of which computing platform to deploy for an embedded application is relating the following three quantities: quality of the results, throughput of the system, and computing resources. Here, we investigate the relationship between these three quantities for an Automatic Target Recognition (ATR) application using Synthetic Aperture Radar (SAR) images. In this embedded application, the three quantities of interest are probability of erroneous target identification, number of images processed per second, and computational power (number and speed of the processors and interconnection network).

An understanding of the relationship that exists between these quantities will better enable system designers to make informed choices when specifying, implementing, and using these systems. One of the degrees of freedom that can be explored is the choice of whether commercial systems are capable of achieving the mission requirements or whether custom solutions are required. We have developed analytic models that relate these quantities, helping designers explore the design space in a more informed way.

2 ATR System Requirements

The problem of automatically identifying an unknown target from a sub-region (chip) within a SAR image can be phrased as the optimization of a function that expresses the likelihood of the observed chip given some combination of parameters [1]. This optimization involves a search for the combination of parameters, including target class, configuration, location, orientation, articulation, and obscuration, that maximize the likelihood function. By some estimates, the search space may need to consist of over 900 million discrete combinations of these parameters in order for an ATR system to deliver useful results. Because the typical operation of an ATR system may involve the analysis of a large number of chips, this space must be searched rapidly for each chip.

The performance of an ATR system, in terms of the probability of erroneously classifying a SAR image of a target, is heavily dependent upon the complexity of the model used to represent such images. Allowing more complex specifications of SAR image models can lead to improved system accuracy. Such increased complexity, however, demands increased computational capability in order to maintain a fixed rate of SAR chip processing. The speed of the available computational machinery determines the level of model complexity that can be supported and therefore determines the overall classification error rate.

Optimization of the likelihood function is a highly parallelizable operation and is easily mapped to multiple processor architectures. Here, we consider an implementation in which each processor is assigned the task of maximizing the likelihood function over some region of the search space. The partitioning of the search space may, for example, be according to target class or some combination of target class and other parameters such as vehicle configuration. Evaluation of the likelihood function itself is a parallelizable operation with operations on each pixel that are independent from and identical to those of every other pixel.

*This work is supported in part by the DARPA VLSI Photonics Program under grant DAAL01-98-C-0074, by the Boeing Foundation, and by the Office of Naval Research grant N00014-98-1-06-06.

3 Computational Models

We have developed analytic models that bring together the following quantities: error probability, system throughput, compute power, and image model complexity. A given probability of error requirement determines a minimum image model complexity, defined as the log of the number of floating point values required to describe each target class, which in turn fixes the relationship between the other two values. The compute power is quantified using system parameters such as the computational speed of the processors, the throughput of the interconnection network, and the throughput of the memory subsystem. The analysis approach is to determine, for any combination of system parameters, the probability of classification error as a function of the rate at which SAR image chips can be processed. This curve describes the system in terms of variables of interest and allows for the direct comparison of alternative hardware designs.

As a simplified example of this kind of analysis, we consider the problem of classifying SAR image chips of unknown targets as one of ten possible target types. We assume that the depression angle from the radar platform is known but that the azimuthal orientation of the target with respect to the platform is not known and that the target is unarticulated and unobscured. DeVore and O'Sullivan [1] present comprehensive experimental results using actual SAR images collected as part of the MSTAR program. Their results yield a dependence between probability of error and database complexity which is approximately proportional to computational complexity. Their curves may be viewed as defining a function

$$Pr(error) = f(C, \alpha_{SAR}), \quad (1)$$

where C is the computational complexity, and α_{SAR} is a vector of parameters describing the recognition problem (number of target classes, depression angle, etc.).

The computational models relate the time to process a single SAR image chip, T_{chip}, computational complexity and system parameters,

$$T_{chip} = g(C, n, \alpha_{comp}), \quad (2)$$

where n is the number of processors, α_{comp} is a vector of parameters describing the embedded hardware (clock rate, throughputs of subsystems, etc.).

These two equations define a parametric curve, parameterized by the computational complexity C. Figure 1 shows the probability of error in classifying a SAR image chip versus the rate at which chips can be processed for 16, 32, and 64 processor systems with all other system parameters held constant. The figure shows the increase in chip processing rate that is attainable with increasing number of processors for a fixed probability of classification error. Alternatively, the figure shows the decrease in classification error rate with increased number of processors for a fixed chip processing rate. The curves in the figure converge at both low and high chip rates and show the most separation in the middle of the graph. At high chip rates, which implies a low image model complexity, the time to classify an image chip is dominated by the time to acquire the chip which is not a function of the number of processors. At low chip rates, only modest improvements in performance are gained through the increased image model complexity.

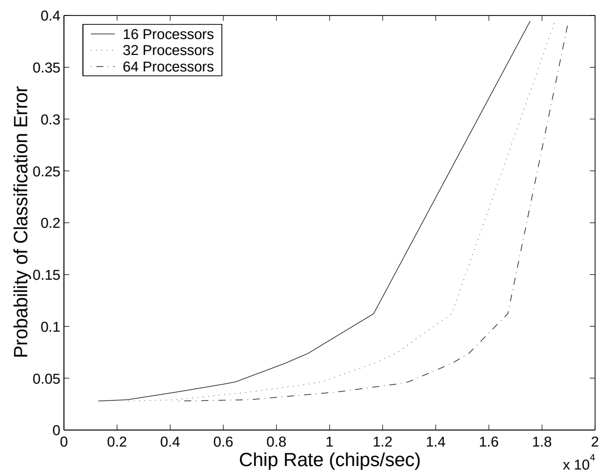

Figure 1: Probability of classification error versus chip processing rate for varying numbers of processors.

References

[1] M. D. DeVore and J. A. O'Sullivan. A performance complexity study of several approaches to automatic target recognition from synthetic aperture radar images. *IEEE Transactions on Aerospace and Electronic Systems*. Submitted for publication.