

Adaptive resource allocation for synthetic aperture radars under resource constraints

Gregory E. Newstadt^a, Edmund G. Zelnio^b, and Alfred O. Hero III^a

^aElectrical Engineering and Computer Science, University of Michigan, Ann Arbor, MI 48109, USA;

^bAir Force Research Laboratory, Wright Patterson Air Force Base, OH, 45433, USA

ABSTRACT

In along-track synthetic aperture radar systems, measurements from multiple phase centers can be used to remove bright stationary clutter in order to detect and estimate moving targets in the scene. The effectiveness of this procedure can be improved by increasing the number of antennas in the system. However, due to computational and communication constraints, it may be prohibitive to use a large number of antennas. In this work, an efficient resource allocation policy is provided to exploit sparsity in the scene, namely that there are few targets relative to the size of the scene. It is shown that even with limited computational resources, one can have significant estimation and computational gains over non-adaptive strategies. Moreover, the performance of the adaptive strategy approaches that of an oracle policy as the number of the stages grows large.

Keywords: adaptive resource allocation, synthetic aperture radar, along-track radar.

1. INTRODUCTION

In along-track synthetic aperture radars (SAR), a single transmit antenna and several receive antennas are placed strategically along the direction of the radar flight path. Due to this arrangement, it is possible to receive measurements of the scene from nearly identical phase center locations within a few milliseconds of each other. Thus, it is possible to coherently (i.e., with phase included) cancel all stationary objects in the scene, including high amplitude clutter and speckle noise.

In practice, noise in the system degrades the signal quality making it difficult to detect targets displaced over high amplitude clutter. One can improve the signal-to-clutter-plus-noise-ratio (SCNR) by increasing the number of antennas in the radar system. However, SAR image formation is a computationally expensive process that scales as $\mathcal{O}(n^2 \log n)$, where n is the number of pixels in the image. In large scenes, real-time processing constraints may limit the number of images that can be formed at any one time. Nevertheless, if the locations containing targets is sparse in comparison to the size of the scene, there may be hope that one can acquire the additional performance gains of a multi-antenna system without incurring a prohibitive computational burden.

This paper provides a first look into using adaptive resource allocation techniques for moving target detection in synthetic aperture radar systems. An algorithm is provided that allocates computational resources in order to jointly detect and estimate targets that are sparsely located within a scene. This paper extends existing work in adaptive resource allocation^{1,2} to efficiently use computational resources in a SAR system. Results on a synthetic dataset demonstrate that it is possible to have nearly 20 dB gains over non-adaptive techniques in mean-squared error, while using less than 60% of the resource budget. Moreover, these results asymptotically approach those of an oracle policy that knows the locations of the targets a priori, as the number of stages grows large.

Further author information: (Send correspondence to Gregory E. Newstadt at newstage@umich.edu)
This document was ... public release statement.

The remainder of this paper is organized as follows. Section 2 provides a brief overview of moving target detection in along-track SAR systems. The problem formulation and adaptive policy is provided in Section 3. In Section 4, the performance of the proposed policy is analyzed. We conclude and point to future work in Section 5.

2. MOVING TARGET DETECTION ALGORITHMS IN SAR SYSTEMS

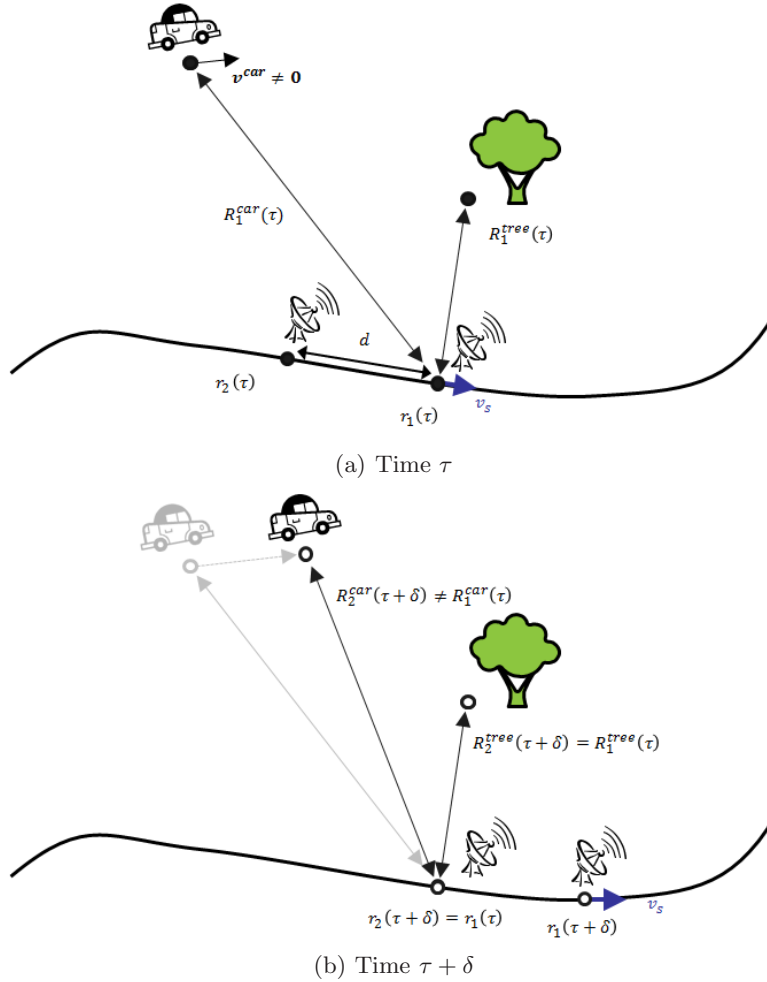


Figure 1. Illustration of an along-track radar with two phase centers. Note that in (b), the stationary targets have identical range and can be effectively canceled. Moving targets exhibit a phase difference and can be localized.

Figure 1 shows an illustration of an along-track system with two phase centers. In the figure, there is a moving target, whose range to the phase center changes from time τ to time $(\tau + \delta)$, and a stationary target, whose range remains constant at both time. This system relies on the fact that the stationary background is identical at both time points. In practice, the phase histories are corrupted by antenna and pulse misalignment, thermal noise, specular glints, and other factors. Nevertheless, many synthetic aperture radar moving target detection systems rely on the fact that there is a high degree of correlation of the stationary “clutter”. This includes displaced phase center array processing (DPCA), along-track interferometry (ATI), and space-time adaptive processing (STAP). Both DPCA and ATI are extremely efficient algorithms that exploit the coherence of the stationary background, though empirical evidence has

demonstrated that each of these algorithms are sensitive to different types of clutter. Deming³ shows that it is possible to combine the benefits of both algorithms to effectively detect targets. Moreover, Deming provides a theoretical analysis of ATI and DPCA that demonstrates that radar measurements from distinct antennas in an along-track system that contain moving targets will be proportional up to a complex phase factor, namely the “interferometric phase”. This phase is related to the target’s radial velocity, which can be used to aid in both target detection and tracking.

STAP extends traditional moving target detection algorithms to the case where there are more than two phase centers. Ender⁴ provides a SAR-focused STAP algorithm that models the stationary clutter as being embedded in a one-dimensional subspace. In a K -antenna system, a single phase center is used to estimate the clutter and the remaining $K - 1$ phase centers are used to estimate the moving component. However, STAP relies on the ability to estimate the clutter distribution in homogeneous target-free data. Moreover, all three algorithms (DPCA, ATI, and STAP) are quite sensitive to calibration errors and specular noise.

In recent work,⁵ a generalized SAR image model was proposed that combined the physical, kinematic, and statistical properties of SAR images into a unified Bayesian structure. The model exploited the plethora of structure that is available in SAR images constructed from multi-antenna and multi-pass data collections. Moreover, the hierarchical Bayesian model (a) infers the parameters of the model directly from the data with very little tuning of algorithm parameters, and (b) provides a Monte Carlo representation of the posterior distribution. This algorithm has the additional benefit that it provides an explicit characterization of the clutter distribution as well as the probability that a target occupies a given region within the scene.

The algorithms discussed above can all be used effectively for SAR moving target detection. In this paper, a simplified model is used to demonstrate the benefit of employing adaptive resource allocation to efficiently use limited resources. In particular, it is assumed that (a) the phase centers are well-calibrated and (b) the clutter distribution is known a priori. In practice, both of these could be accomplished by using the SAR hierarchical Bayesian model to estimate these parameters from previous passes of the radar. Within the i -th spatial region, the clutter distribution in an along-track system with K antennas is modeled as complex-valued multivariate Normal with parameters

$$\mathbf{c}_i \sim \mathcal{CN}(0, \gamma_i^2 [(1 - \rho_i)\mathbf{I}_{K \times K} + \rho_i \mathbf{1}_K \mathbf{1}_K^T]) \in \mathbb{C}^K, \quad (1)$$

where γ_i^2 denotes the antenna (channel) variance, and ρ_i denotes the antenna (channel) coherence. Note that both γ_i^2 and ρ_i will vary spatially (i.e., buildings versus roads) and temporally (as the radar platform changes spatial location).

3. PROBLEM FORMULATION

Limited resources can be efficiently used by focusing effort in certain areas, while reducing it in others. Thus, an agile system must be available to dynamically allocate the available resources. Note that these resources come in many forms, such as radar dwell time, monetary resources, or computational time. In this work, we focus on the latter case, though this work is generalizable to other agile systems.

3.1 Scene segmentation and measurement model

This work assumes that the scene of interest is large and can be divided into many smaller subregions that may or may not contain targets. In particular, it is assumed that the number of subregions containing targets is small as compared to the size of the scene. An indicator variable for the existence of a target in the i -th region is defined as

$$I_i = \begin{cases} 1, & \text{target is in the } i\text{-th subregion} \\ 0, & \text{else} \end{cases}, \quad i = 1, 2, \dots, N, \quad (2)$$

where N is the number of subregions. A sparseness prior on these indicators is imposed so that

$$\Pr(I_i = 1) \ll 1 \quad (3)$$

In an along-track SAR system with K_{MAX} antennas, one can choose to use anywhere between 0 and K_{MAX} antennas to form measurements (images) within the i -th subregion. In practice, SAR image formation within subregions of the scene can be done efficiently using digital spotlighting. We define $K_i \in [0, K_{MAX}]$ to be the number of antennas used to create a vector of measurements, which are modeled as

$$\mathbf{y}_i = \mathbf{x}_i I_i + \mathbf{c}_i + \mathbf{n}_i \in \mathbb{C}^{K_i} \quad (4)$$

$$\mathbf{x}_i = \mathbf{h}_{K_i}^\phi \mu_i \in \mathbb{C}^{K_i} \quad (5)$$

where \mathbf{c}_i is the clutter component, \mathbf{n}_i is the noise component, and \mathbf{x}_i is the target component in the i -th subregion, composed of a reflectance μ_i and a interferometric vector defined by

$$\mathbf{h}_{K_i}^\phi = \begin{bmatrix} 1 \\ \exp\{2\pi j\phi\} \\ \vdots \\ \exp\{2\pi j(K_i - 2)\phi\} \\ \exp\{2\pi j(K_i - 1)\phi\} \end{bmatrix} \in \mathbb{C}^{K_i} \quad (6)$$

Note that this model can be generalized so that each of the K_i measurements is a vector quantity (i.e., a group of pixels), but the simpler case is used here to demonstrate utility of resource management.

In an adaptive setting, multiple stages of measurements are collected, indexed by a time t . A single stage of measurements consists of a vector of measurements $\mathbf{y}_i(t)$ for each subregion $i = 1, 2, \dots, N$, where

$$\mathbf{y}_i(t) = \mathbf{x}_i(t) I_i + \mathbf{c}_i(t) + \mathbf{n}_i(t) \in \mathbb{C}^{K_i(t)} \quad (7)$$

$$\mathbf{x}_i(t) = \mathbf{h}_{K_i(t)}^\phi \mu_i \in \mathbb{C}^{K_i(t)} \quad (8)$$

for $i = 1, 2, \dots, N$. In this formulation, targets are assumed to be static. Hence, there is no time index on μ_i , ϕ_i or I_i . However, the noise and clutter components, as well as the number of measurements per subregion $K_i(t)$, are assumed to vary over time. This model can be generalized to account for time-varying targets in a method similar to Chapter 3 of related work.⁵ The prior distributions on these parameters are given by

$$\begin{aligned} \mu_i &\sim \mathcal{CN}(0, \sigma_\theta^2) \\ \mathbf{c}_i(t) &\sim \mathcal{CN}(0, \Sigma_i^c(t)) \\ \mathbf{n}_i(t) &\sim \mathcal{CN}(0, \Sigma^n) \end{aligned} \quad (9)$$

with

$$\begin{aligned} \Sigma_i^c(t) &= \gamma_i^2(t) [(1 - \rho_i(t)) \mathbf{I}_{K \times K} + \rho_i(t) \mathbf{1}_K \mathbf{1}_K^T] \\ \Sigma^n &= \nu^2 \mathbf{I}_{K \times K} \end{aligned} \quad (10)$$

Note that given the interferometric phase ϕ_i , the posterior distributions of μ_i and I_i given the measurements $\mathbf{Y}_i(t) = \{\mathbf{y}_i(1), \mathbf{y}_i(2), \dots, \mathbf{y}_i(t)\}$ are given as the following closed-form updates

$$\Pr(I_i = 1 | \mathbf{Y}_i(t)) = \frac{f(\mathbf{y}_i(t) | I_i = 1) \Pr(I_i = 1 | \mathbf{Y}_i(t-1))}{f(\mathbf{y}_i(t) | I_i = 1) \Pr(I_i = 1 | \mathbf{Y}_i(t-1)) + f(\mathbf{y}_i(t) | I_i = 0) \Pr(I_i = 0 | \mathbf{Y}_i(t-1))} \quad (11)$$

and

$$\begin{aligned}
\boldsymbol{\eta}_i(t) &= \mathbf{y}_i(t) - \mathbf{h}_{K_i}^\phi(t) \hat{\mu}_i(t-1) \\
\mathbf{S}_i(t) &= [\mathbf{h}_{K_i}^\phi(t)] [\mathbf{h}_{K_i}^\phi(t)]^H \sigma_i^2(t-1) + \boldsymbol{\Sigma}_i^c(t) + \boldsymbol{\Sigma}^n \\
\boldsymbol{\Gamma}_i(t) &= \sigma_i^2(t-1) \cdot [\mathbf{h}_{K_i}^\phi(t)]^H (\mathbf{S}_i(t))^{-1} \\
\hat{\mu}_i(t) &= \hat{\mu}_i(t-1) + \boldsymbol{\Gamma}_i(t) \boldsymbol{\eta}_i(t) \\
\sigma_i^2(t) &= \left[1 - \boldsymbol{\Gamma}_i(t) \mathbf{h}_{K_i}^\phi(t) \right] \sigma_i^2(t-1)
\end{aligned} \tag{12}$$

with initial conditions

$$\begin{aligned}
\hat{\mu}_i(0) &= 0 \\
\sigma_i^2(0) &= \sigma_\theta^2
\end{aligned} \tag{13}$$

Note that $\sigma_i^2(t)$ actually has the form:

$$\begin{aligned}
\sigma_i^2(t) &= \sigma_i^2(t-1) - \boldsymbol{\Gamma}_i(t) \mathbf{h}_{K_i}^\phi(t) \sigma_i^2(t-1) \\
&= \sigma_i^2(t-1) - \sigma_i^2(t-1) [\mathbf{h}_{K_i}^\phi(t)]^H (\mathbf{S}_i(t))^{-1} [\mathbf{h}_{K_i}^\phi(t)] \sigma_i^2(t-1)
\end{aligned} \tag{14}$$

Using various matrix inverse identities, it can be shown that

$$\sigma_i^2(t) = \left[\frac{1}{\sigma_i^2(t-1)} + \frac{1}{\alpha_i(t)} \left(K_i(t) + \frac{\beta_i(t) K_i^2(t) |\text{sinc} \frac{\phi_i K_i(t)}{2\pi}|^2}{K_i(t) \beta_i(t) + \alpha_i(t)} \right) \right]^{-1} \tag{15}$$

where

$$\begin{aligned}
\alpha_i(t) &= \gamma_i^2(t) (1 - \rho_i(t)) + \nu^2 \\
\beta_i(t) &= \gamma_i^2(t) \rho_i(t)
\end{aligned} \tag{16}$$

Moreover, Equation (15) has a simple upper bounded that is tight for large interferometric phase ϕ_i (where the sinc term is approximately zero), given by

$$\begin{aligned}
\sigma_i^2(t) &= \left[\frac{1}{\sigma_i^2(t-1)} + \frac{1}{\alpha_i(t)} \left(K_i(t) + \frac{\beta_i(t) K_i^2(t) |\text{sinc} \frac{\phi_i K_i(t)}{2\pi}|^2}{K_i(t) \beta_i(t) + \alpha_i(t)} \right) \right]^{-1} \\
&\leq \left[\frac{1}{\sigma_i^2(t-1)} + \frac{K_i(t)}{\alpha_i(t)} \right]^{-1}
\end{aligned} \tag{17}$$

This shows that the posterior variance of μ_i is inversely proportional to the number of measurements $K_i(t)$, though the benefit of each measurement will depend on the channel variance $\gamma_i^2(t)$, the channel coherence $\rho_i(t)$ and the noise variance ν^2 (all through the $\alpha_i(t)$ term). Note that when ϕ_i is also random, then these will not be the exact update equations of the posterior distribution, so one would need to use a richer posterior estimation technique such as an extended Kalman filter or particle filter.

3.2 Computational resource constraints and cost function

We formulate an optimization problem by considering a resource budget on the total number of measurements that can be used at any particular stage:

$$\sum_{i=1}^N K_i(t) \leq K_{budget}(t), \quad t = 1, 2, \dots, T \tag{18}$$

where N is the number of subregions, T is the number of stages of measurements and $K_{budget}(t)$ is the known budget per stage. For an extension to the case where the per-stage budget is also a design parameter, see the work of Wei and Hero.⁶ The optimization problem is then of the form

$$\min_{\{K_i(t)\}_{i,t}} J(\{K_i(t)\}_{i,t}) \quad s.t. \quad \sum_{i=1}^N K_i(t) \leq K_{budget}(t), \quad t = 1, 2, \dots, T \quad (19)$$

where $J(\{K_i(t)\}_{i,t})$ is a cost function that depends on our design criteria. In this work, we choose to optimize the mean-squared error within the region of interest, defined as

$$J(\{K_i(t)\}_{i,t}) = \mathbb{E} \left[\sum_{t=1}^T \sum_{i=1}^N I_i (\hat{\mu}_i(t) - \mu_i)^2 \right] \quad (20)$$

Moreover, this cost is optimized in a greedy fashion, so that decisions at the t -th stage depend on the greedy cost function:

$$J_t(\{K_i(t)\}_i) = \mathbb{E} \left[\sum_{i=1}^N I_i (\hat{\mu}_i(t) - \mu_i)^2 \right] \quad (21)$$

For extensions to non-greedy optimization techniques, see Chapter 3 in related work.⁵ This cost function is optimized by letting

$$\hat{\mu}_i(t) = \mathbb{E} [\mu_i | \mathbf{Y}(t-1)], \quad (22)$$

i.e., the posterior mean. Plugging Equation (22) into Equation (21) yields

$$\begin{aligned} J_t(\{K_i(t)\}_i) &= \sum_{i=1}^N \mathbb{E}(I_i | \mathbf{Y}(t-1)) \sigma_i^2(t) \\ &\leq \sum_{i=1}^N \Pr(I_i = 1 | \mathbf{Y}_i(t-1)) \left[\frac{1}{\sigma_i^2(t-1)} + \frac{K_i(t)}{\alpha_i(t)} \right]^{-1} \end{aligned} \quad (23)$$

where the last term is provided by the upper bound in Equation (17).

3.3 Adaptive resource allocation policy

The optimization problem in this paper follows a form very similar to related work¹ and has an approximate closed-form solution as follows. First, define

$$w_i(t) = \sqrt{\Pr(I_i(t) | \mathbf{Y}_i(t-1))} \quad (24)$$

$$c_i(t) = 1/\sigma_i^2(t) \quad (25)$$

$$z_i(t) = w_i(t) \sigma_i^2(t) \quad (26)$$

$$\lambda_i(t) = \frac{K_i(t)}{\alpha_i(t)} \quad (27)$$

and order the z_i 's so that

$$z_{\tau(1)}(t) \leq z_{\tau(2)}(t) \leq \dots \leq z_{\tau(N)}(t). \quad (28)$$

Define the cutoff $k_0 = 0$ if

$$c_{\tau(1)}(t) \sum_{i=1}^N \frac{w_{\tau(i)}(t)}{w_{\tau(1)}(t)} < K_{budget}(t) + \sum_{i=1}^N c_i(t); \quad (29)$$

otherwise, k_0 is the integer k satisfying

$$z_{\tau(k)}(t) \leq \frac{\sum_{i=k+1}^N w_{\tau(i)}(t)}{K_{budget}(t) + \sum_{i=k+1}^N c_{\tau(i)}(t)} < z_{\tau(k+1)}(t) \quad (30)$$

Note that k_0 is proved to exist and be unique.⁷ Then the intermediate solution to the optimization problem is given by

$$\lambda_i(t) = \frac{\lambda_{total}(t) + \sum_{k=k_0+1}^Q c_{\tau(k)}(t)}{\sum_{k=k_0+1}^Q w_{\tau(k)}(t)} w_i(t) - c_i(t) \quad (31)$$

for all $\tau(i) > k_0$ and 0 else. The final solution is given by

$$K_i(t) = \min \{ \text{floor}(\lambda_i(t)\alpha_i(t)), K_{MAX} \}, \quad (32)$$

where it should be noted that $K_i(t)$ must be an integer between 0 and K_{MAX} . Thus, the provided solution is a relaxation to an integer program, and subsequently may not be an optimal solution. However, it is very tractable from a computational standpoint as the complexity is dominated by the sorting algorithm which runs in $\mathcal{O}(N \log N)$ time (as compared to significantly higher computational complexity of integer programs).

4. PERFORMANCE ANALYSIS

This section provides an analysis of the proposed policy under considerable resource constraints. Initially, policies for comparison to the one proposed in Section 3.3 are provided, followed by performance measures used for numerical comparisons.

4.1 Policies for comparison

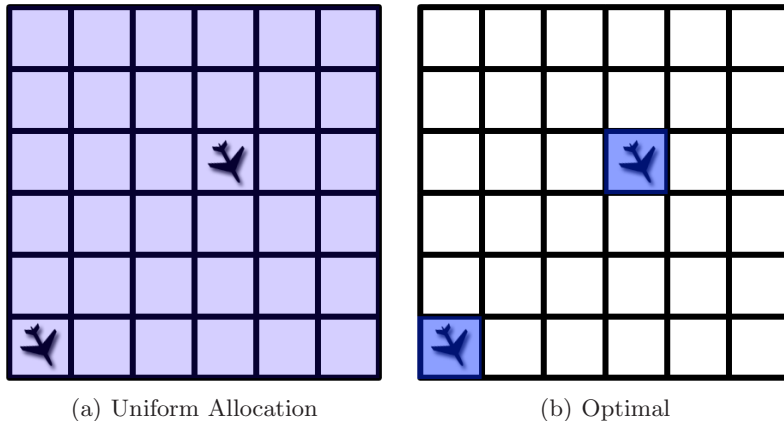


Figure 2. Illustration of an along track radar with two phase centers. Note that in (b), the stationary targets have identical range and can be effectively canceled. Moving targets exhibit a phase difference and can be localized.

We consider two intuitive policies for comparison to the one provided in Section 3.3. On one side of the spectrum, the uniform allocation policy equally distributes resources across all subregions so that

$$K_i^{uniform}(t) = \text{floor} \left(\frac{K_{budget}(t)}{N} \right). \quad (33)$$

Note that any competitive policy should perform better than a uniform allocation. In contrast, the oracle (or omniscient) policy, which knows the locations of targets a priori, provides an upper bound on potential performance of any policy. It can be shown that the oracle policy that optimizes the greedy cost function has an intermediate solution given by

$$\lambda_i^{oracle}(t) = \begin{cases} \frac{K_{budget}}{N_{targets}}, & I_i = 1 \\ 0, & I_i = 0 \end{cases} \quad (34)$$

where $N_{targets}$ is the number of targets in the scene. This yields the oracle policy

$$K_i^{oracle}(t) = \min \{ \text{floor}(\lambda_i^{oracle}(t)\alpha_i(t)), K_{MAX} \} \quad (35)$$

Both the uniform and oracle policies are illustrated graphically in Figure 2, where a darker color illustrates more sensing effort. Note that the oracle policy is unattainable with finite resources, but we hope to approach the performance of this policy through adaptation of computational resources over time.

This work considers two performance measures for numerical comparison: (a) the mean-squared error (MSE) performance as compared to the uniform policy, and (b) the percentage of available resources used. For the former case, the gains in MSE as compared to the uniform policy is defined as

$$G(\mathbf{K}) = -10 \log \frac{MSE(\mathbf{K})}{MSE(\mathbf{K}^{uniform})} \quad (36)$$

where $MSE(\mathbf{K})$ is the average MSE over all stages and trials of the policy defined by \mathbf{K} . Note that one would expect the MSE of a proposed policy to be less than that of the uniform policy (and so G should be positive).

The second performance considered is a product of the fact that the proposed policies may use fewer than the maximum amount of available resources. This is due to using a continuous relaxation to the integer program, as well as having an upper bound on the computational resources $K_i(t) \leq K_{MAX}$. The second measure is defined as the percentage of average resources used

$$H(\mathbf{K}) = \hat{\mathbb{E}} \left[\sum_{t=1}^T \frac{1}{K_{budget}(t)} \sum_{i=1}^N K_i(t) \right] \quad (37)$$

where an empirical expectation is used over independent trials in the simulations.

4.2 Simulated dataset

In this paper, we consider the following setup:

- The number of subregions, $N = 1,000$.
- The prior probability of target existence, $\Pr(I_i = 1) = 0.01$.
- The prior target variance, $\sigma_\theta = 1$.
- The additive noise variance, $\nu = 1/(64^2)$.
- The clutter noise variance, randomly distributed according to Figure 3. Note, most of the clutter noise has lower variance than the target (i.e., negative log-ratio), but there is a significant portion with higher variance.
- The interferometric phase, $\phi_i = \pi/3$, assumed known.




Figure 3. Distribution of the ratio of clutter variance to target variance (on a dB scale)

- The number of measurement stages, $T = 30$.
- The average number of measurements per subregion, 2
- The number of independent Monte Carlo trials, $M = 200$.
- The maximum number of antennas, $K_{MAX} = 10, 20, 50$.

The results are provided in Figure 4 with columns representing different along-track radar systems with $K_{MAX} = 10, 20, 50$, respectively. The adaptive policy does significantly better than the uniform allocation as the number of stages increase. Moreover, it approaches the performance of the oracle policy as the number of stages increase. Finally, the optimal/adaptive policies use significantly fewer resources than are maximally available. These computational resources can be used for other tasks such as estimating the interferometric phase, estimating the clutter distribution (i.e., with a Gibbs sampler), or exploring the remainder of the scene (to increase robustness to model mismatch).


5. CONCLUSIONS AND FUTURE WORK

In this paper, we have provided a first glimpse into using adaptive resource allocation in order to efficiently use computational resources in order to simultaneously detect and estimate targets in multiple-antenna along-track SAR systems. We are able to provide significant gains over a non-adaptive uniform strategy, both in terms of estimation error as well as computational savings. This could be used to alleviate the large computational/communication burden associated with large-area along-track systems, as well as provide stronger signal quality for post-processing.


This work was based on a simplified SAR measurement model in order to provide intuition on the ability to use adaptive resource allocation techniques for efficient sensor management. Future work will look into ways to increase the fidelity of the model to real-world sensors, yet maintain the benefits provided in this paper. This would include (a) using the full SAR image model as provided in related work,⁵ (b) using richer filtering/posterior estimation techniques (such as particle filters), and (c) including richer target models that can include dynamic motion, such as transition probabilities and/or population models.

Acknowledgment

This research was partially supported by the Air Force Research Laboratory Sensors Directorate under grant FA8650-07-D-1220-0006.



(a) Gain in mean-squared error compared to uniform policy, $G(\mathbf{K})$



(b) Percentage of resources used, $H(\mathbf{K})$

Figure 4. This figure shows performance of the adaptive, uniform, and oracle policies as measured by MSE and percentage of resources used, for three values of $K_{MAX} = 10, 20, 50$ (from left-to-right, respectively). Note that the adaptive policy approaches the performance of the oracle policy which has greater than 15 dB performance gains over the uniform policy in MSE. Moreover, the adaptive and oracle policies use significantly fewer resources than the uniform allocation. Both of these policies use less than 60% of the available resources for at later stages ($t \approx 30$).

REFERENCES

- [1] E. Bashan, R. Raich, and A. O. Hero III, "Optimal two-stage search for sparse targets using convex criteria," *IEEE Transactions on Signal Processing*, vol. 56, pp. 5389–5402, 2008.
- [2] E. Bashan, G. Newstadt, and A. O. Hero III, "Two-stage multi-scale search for sparse targets," *IEEE Transactions on Signal Processing*, vol. 59, no. 5, pp. 2331–2341, 2011.
- [3] R. Deming, S. MacIntosh, and M. Best, "Three-channel processing for improved geo-location performance in SAR-based GMTI interferometry," in *Proceedings of SPIE*, vol. 8394, 2012, p. 83940F.
- [4] J. H. G. Ender, "Space-time processing for multichannel synthetic aperture radar," *Electronics & Communication Engineering Journal*, vol. 11, no. 1, pp. 29–38, 1999.
- [5] G. Newstadt, "Adaptive sensing techniques for dynamic target tracking and detection with applications to synthetic aperture radars," Ph.D. dissertation, University of Michigan, January 2013.
- [6] D. Wei and A. O. H. III, "Multistage adaptive estimation of sparse signals," in *Proceedings of 2012 IEEE Statistical Signal Processing Workshop*, August 2012.
- [7] E. Bashan, R. Raich, and A. O. H. III, "Adaptive sampling: efficient search schemes under resource constraints," University of Michigan, Communications and Signal Processing Lab., Tech. Rep. 385, October 2007.