Autonomous Agents for Multi-Function Radar Resource Management

Alexander Charlish

A thesis submitted for the degree of

Doctor of Philosophy

of the

University College London.

Electronic and Electrical Engineering

2011

I, Alexander Charlish, confirm that the work presented in this thesis is my own. Where information has been derived from other sources, I confirm that this has been indicated in the thesis.

Signature...

Abstract

The multifunction radar, aided by advances in electronically steered phased array technology, is capable of supporting numerous, differing and potentially conflicting tasks. However, the full potential of the radar system is only realised through its ability to automatically manage and configure the finite resource it has available. This thesis details the novel application of agent systems to this multifunction radar resource management problem. Agent systems are computational societies where the synergy of local interactions between agents produces emergent, global desirable behaviour.

In this thesis the measures and models which can be used to allocate radar resource is explored; this choice of objective function is crucial as it determines which attribute is allocated resource and consequently constitutes a description of the problem to be solved. A variety of task specific and information theoretic measures are derived and compared. It is shown that by utilising as wide a variety of measures and models as possible the radar's multifunction capability is enhanced.

An agent based radar resource manager is developed using the JADE Framework which is used to apply the sequential first price auction and continuous double auctions to the multifunction radar resource management problem. The application of the sequential first price auction leads to the development of the Sequential First Price Auction Resource Management algorithm from which numerous novel conclusions on radar resource management algorithm design are drawn. The application of the continuous double auction leads to the development of the Continuous Double Auction Parameter Selection (CDAPS) algorithm. The CDAPS algorithm improves the current state of the art by producing an improved allocation with low computational burden. The algorithm is shown to give worthwhile improvements in task performance over a conventional rule based approach for the tracking and surveillance functions as well as exhibiting graceful degradation and adaptation to a dynamic environment.

To Mum and Dad

Acknowledgements

Contents

List of Figures

List of Tables

List of Symbols

α	False track confirmation probability
β	True track deletion probability
Γ	Parameter space
γ_k	Parameter selection
$\hat{\gamma}(I_k)$	Policy using information accumulated up to time k
Δ_{az}	Extent of search volume in azimuth
$\Delta(\theta)$	Difference response in angle for two beams
ϵ_T, ϵ_P	True and predicted target angle in elevation
η	Aperture efficiency of array antenna
η_T, η_P	True and predicted target angle in azimuth
Θ	Singer target manoeuvre time
θ_0	Angle of steered beam
θ_B	3dB width of beam
θ_f	Off-boresight measurement angle
θ_{elmin}	Minimum elevation angle of search volume
θ_{elmax}	Maximum elevation angle of search volume
θ_g	Grating lobe angle
θ_s	Beam separation in angle
λ	Wavelength of carrier
μ	False target density
$\tilde{\nu}$	White noise process
ξ	Arbitrary position coordinate
σ	Radar cross section
$\bar{\sigma}$	Mean radar cross section
σ_θ	Measurement thermal noise error standard deviation
σ_m	Measurement error standard deviation
σ_p	Angular estimation error standard deviation
σ_u, σ_v	State estimation error standard deviation in u,v coordinates
$\tilde{\sigma}_\theta$	Off-boresight measurement error standard deviation
$\Sigma(\theta)$	Sum response in angle for two beams

List of Symbols

τ	Pulse width
τ_A	Alert dwell length
τ_C	Confirm dwell length
τ_d	Dwell length
τ_s	Time required to survey a search volume
$\phi(t)$	Phase modulation
$\chi(t_d, f_d)$	Autocorrelation function
$\Psi(t_d, f_d)$	Ambiguity function
Ω	Singer target manoeuvre standard deviation
Ω_s	Search volume in steraradians
A	Coherent signal amplitude
A_A	Set of currently active asks
A_e	Array antenna effective aperture area
A_k	Action at time step k
A_r	Attenuation factor
A_T	Actual antenna aperture area
a_t	Ask announced by agent t
B	Signal bandwidth
B_A	Set of currently active bids
$b_1(\theta), b_2(\theta)$	Angle response of two adjacent beams
b_t	Bid announced by agent t
c	Speed of electromagnetic wave propagation
d^2	Statistical distance squared
d_c	Target closure range between scans
$d_{t_k}(o)$	Delay function for task agent t_k under outcome o
$D = \{d_1, ... d_n\}$	Data vector
$D_{KL}(P||Q)$	Kullback-Leibler divergence between P and Q
D_P	Beam spacing for volume search
E	Signal envelope
$E(\theta)$	Complete array antenna radiation pattern
$E_\alpha(\theta)$	Array antenna radiation pattern with isotropic elements
$E_e(\theta)$	Array antenna element factor
E_r	Received bandpass signal energy
F_k	State transition matrix at time k
F_n	Receiver noise figure
f_c	Frequency of carrier
f_d	Doppler frequency
$f_k(x_k, v_k)$	State transition function

$f(T_A)$	Social choice function for agent set T_A
G	Major axis of uncertainty ellipse
G_t	Transmitter gain
G_r	Receiver gain
g	Validation gate size
$g(t)$	Amplitude modulation
$g_k(\gamma)$	Resource function
$H(X)$	Entropy of random variable X
$H_f(f)$	Frequency response
H_k	Measurement observation matrix at time k
H_N	Target not present hypothesis
H_T	Target present hypothesis
$h(t)$	Impulse response
$h_k(x_k, w_k)$	Observation function
I	Subset of asks
I_0	Modified Bessel function of first kind with zero order
$I(X;Y)$	Mutual information between X and Y
$i_{t_k}(o)$	Information function for agent t_k under outcome o
J	Subset of bids
J_z	Fisher information for measurement z
k	Boltzmann's Constant
$\hat{k}$	k-pricing rule weight
k_c	Complex constant
k_m	Gradient of measurement slope on measurement axis
k_s	Normalised error signal response
l	Length of array antenna
l_d	Resource loading
L	Log Likelihood ratio
LR	Likelihood ratio
L_p	Losses
m	Measurement dimension
N	Mean noise power
n	Number of elements in antenna array
n_1, n_2	Zeros mean unit standard deviation random variables
N_B	Number of detection bins
N_0	Noise spectral density
O	Set of possible outcomes for event
o	Outcome of an event

<center>List of Symbols</center>

P_D	Probability of detection
P_{D0}	Probability of detection at the beam centre
P_{Dc}	Cumulative probability of detection
P_{FA}	Probability of false alarm
$P_{k|k}$	Covariance of state estimate at time k given the measurements up to time k
P_N	Target not present probability density function
P_r	Received power
P_t	Transmit power
P_T	Target present probability density function
PRF	Pulse repetition frequency
p_i	Priority of task i
p_n	Price of offer n
$p*$	Agent's true price valuation
$\tilde{p}$	Transaction price
q_2	Measurement origin uncertainty
$q_k(\gamma)$	Quality function
Q_k	Covariance of target dynamic noise at time k
$\hat{Q}_k$	Quality space
$\tilde{Q}_k$	POMDP Q-value
Q_I	Quantity of ask subset I
Q_J	Quantity of bid subset J
q_n	Quantity of offer n
$\tilde{q}$	Process noise intensity
$\tilde{R}_k(s,a)$	Reward function for action a from state s
R_k	Covariance of measurement noise at time k
R_m	Maximum detection range
R_r	Range resolution
R_t	Range between the target and the antenna
R_u	Maximum unambiguous range
R_0	Range at which signal to noise ratio is unity
R_{85}	Range at which cumulative detection probability is 0.85
R_{90}	Range at which cumulative detection probability is 0.9
r_k	Resource held by task agent k
r_T	Total resource available for set of agents T
S_k	Predicted covariance of track residual
S_{min}	Minimum detectable signal
SNR	Signal to noise ratio
SNR_m	Signal to noise ratio given a detection

List of Symbols

SNR_0	Signal to noise ratio at beam center
s	Antenna element spacing
$s_r(t)$	Received signal
$s_t(t)$	Transmitted signal
T	Threshold
T_1, T_2	Likelihood ratio thresholds
T_S	Total time for search dwell
$T_A = \{t_1, ..., t_k\}$	Set of k task agents
T_0	Temperature of ideal noise source
T_k	Length of one time step
T_s	Effective temperature
t_d	Round trip time from antenna to target to antenna
t_f	Revisit interval
t_k	Task agent k
t_p	Time interval between pulses
U	Variance reduction ratio
$u_{t_k}(o)$	Utility function for agent t_k under outcome o
$u_R(t)$	Complex envelope of the received signal
$u_T(t)$	Complex envelope of the transmit signal
V_k	Validation gate volume
V_I	Total value of ask subset I
V_J	Total value of bid subset J
v_0	Track sharpness
v_k	Target dynamic noise
v_r	Relative radial velocity between target and radar
w_k	Measurement noise
W_k	Filter gain matrix
X_k	State vector
x_k	Target state at time k
$\hat{x}_{k|k}$	State estimate at time k given measurements up to time k
Z_k	Measurement series up to time k
z	Signal intensity
z_k	Measurement at time k
$\hat{z}_k$	Measurement prediction
$\tilde{z}_k$	Measurement residual

List of Acronyms

(CA)-CFAR	(Cell Averaging) - Constant False Alarm Rate
CDA	Continuous Double Auction
CDAPS	Continuous Double Auction Parameter Selection
ECM	Electronic Counter Measures
EDF	Earliest Deadline First
ESA	Electronically Steered Array
FIPA	Foundation for Intelligent Physical Agents
FOV	Field of View
GIF	Greatest Information First
GMTI	Ground Moving Target Indicator
GNN	Global Nearest Neighbour
HPF	Highest Priority First
IMM	Interacting Multiple Model
JADE	Java Agent Development
JPDA	Joint Probabilistic Data Association
KKT	Karush Kuhn Tucker
KLD	Kullback-Leibler Divergence
LFM	Linear Frequency Modulation
LQF	Lowest Quality First
MAS	Multi-Agent System
MDP	Markov Decision Process
MESAR	Multi-Function Electronically Scanned Adaptive Radar
MFR	Multi-Function Radar
MHT	Multi-Hypothesis Tracking
MI	Mutual Information
NCA	Nearly Constant Acceleration
NCV	Nearly Constant Velocity
NYSE	New York Stock Exchange
OARS	Opening Automated Report Service
POMDP	Partially Observable Markov Decision Process

List of Acronyms

(L/M/H)-PRF	(Low/Medium/High) - Pulse Repetition Frequency
PDA	Probabilistic Data Association
QoS	Quality of Service
Q-RAM	Qos Resource Allocation Method
RaDAR	Radio Detection and Ranging
RB-EDF	Rule Based Earliest Deadline First
RBPS	Rule Based Parameter Selection
RCS	Radar Cross Section
RGPO	Range Gate Pull Off
RMSE	Root Mean Squared Error
RM	Resource Management
RRM	Radar Resource Manager
SAR	Synthetic Aperture Radar
ISAR	Inverse Synthetic Aperture Radar
SFPARM	Sequential First Price Auction Resource Management
SNR	Signal to Noise Ratio
SOJ	Stand Off Jammer
STAP	Space Time Adaptive Processing
TWS	Track While Scan

Chapter 1

Introduction

1.1 Motivation

Advances in modern electronic components have driven the commercialisation of electronically steered phased array antenna technology. In contrast to the traditional mechanically scanned antenna, the electronically steered phased array has significantly increased beam agility which allows dynamic allocation of the time-energy resource. This has led to a new generation of multifunction radar systems, where multifunction can be defined as the ability to sequentially execute numerous, differing and potential conflicting tasks which support a variety of different radar functions.

Requirements of multifunction radar according to the maritime, airborne and land domains vary greatly. However, a typical system is required to search a volume for new targets and once detected fuse the information from multiple scans into target tracks. The system may also be required to perform additional functions depending on the application domain such as data link, weapons support, identification or classification. A typical scenario is shown in Fig. 1.1 for a maritime air defence type application. This figure shows the potential operational complexity for the multifunction radar as the finite radar resource is required to be distributed between the wide variety of modes which may need to be deployed. The ultimate performance of the system is dependent on how well the numerous tasks which support the differing modes are able to fulfil the requirements of the system.

Multifunction radars have increasing appeal, which can be attributed to several key benefits:

- *Flexibility* - Flexibility over allocation in space and time, including variable update rates, dwell times and surveillance coverage, tailored to each application or role.

- *Adaptability* - Multifunction radar performance specification can be dynamically adjusted to match the dynamic and uncertain scenario and environment.

- *Efficiency* - Increased efficiency in terms of space, time, energy, production and maintenance effort.

The overall benefit is the potential to vary nearly instantaneously an array of radar parameters to achieve
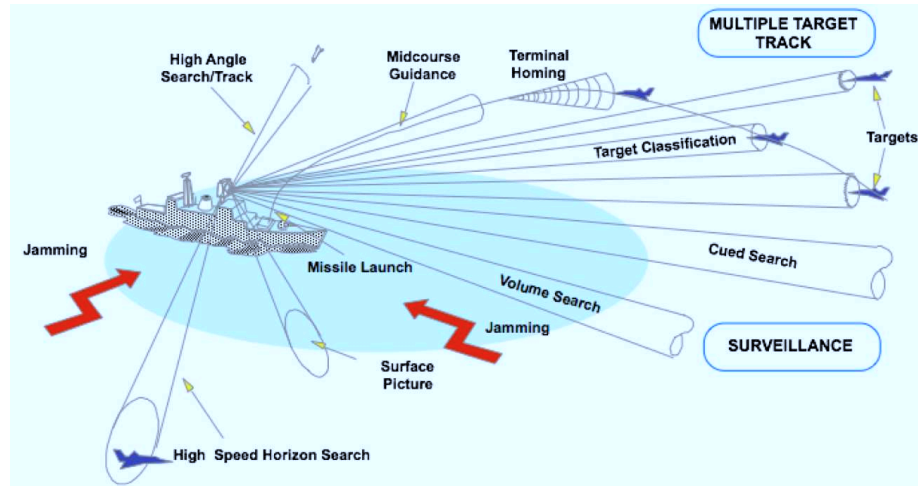
Figure 1.1: Typical maritime air defence multifunction radar scenario [Butler, 1998].

a desired goal. This thesis concentrates on the exploitation of beam agility and development of methods to divide the finite time-energy resource.

The control and configuration of the multifunction radar is beyond the response capability of the human operator and so an automated Radar Resource Manager (RRM), most likely with operator supervision, is required. Consequently, the full potential of the multifunction radar system is only realised through the RRM's ability to automatically allocate and configure the finite resource it has available. In addition, the RRM has access to all the information in the system, which exceeds the information that is able to be displayed to an operator in the loop. The RRM can therefore theoretically achieve superior decision making at a rate faster than the human operator. These factors have created a strong desire to maximise the potential of the hardware by intelligently adapting to dynamic scenarios, environments and missions.

Agent systems are computational societies where the synergy of local interactions between agents produces emergent, globally desirable behaviour. Typically, agent systems are governed by distributed and decentralised mechanisms which are inherently computationally efficient and scalable. The automation of human interaction mechanisms in agents systems, such as economic paradigms, can replicate the ability to achieve robust behaviour in dynamic and uncertain environments. This provides the motivation for their application to multifunction radar resource management.

Economic paradigms and market mechanisms have evolved over centuries in human societies, as efficient, trusted and highly developed methods of distributing goods and commodities. Free markets tend to competitive equilibrium which maximises participant profit and optimises social welfare. This desirable characteristic can be harnessed in resource allocation problems, such as multifunction radar resource management, to produce emergent intelligent and desirable behaviour.

The primary aim of this research has been to investigate for the first time the application of agent systems and economic paradigms to multifunction RRM. This research also had the following secondary aims:

- Provide a thorough review of existing work, to identify where agent techniques can be most beneficially applied.

- Investigate the role of information theory in multifunction radar resource management.

- Explore suitable objective functions and measures which guide the resource allocation.

- Develop agent based resource allocation mechanisms utilising suitable choices of objective functions.

- Create a radar simulator testbed upon which differing agent systems can be applied.

- Demonstrate and quantify enhanced multifunction capability of resulting allocation mechanism.

This thesis is organised according to the layout described in the following section.

1.2 Thesis Layout

This thesis contains eight chapters which detail the aspects of the application of agent systems to multifunction radar resource management.

An overview of the fundamental operation of a multifunction radar system is given in Chapter 2. This includes aspects of radar signal processing, measurement data processing and the electronically steered array antenna. Operational multifunction radar systems are discussed alongside the parameter and mode view of multifunction operation. This collection describes *what* is automatically managed.

Chapter 3 gives an overview of current radar resource management techniques. This includes the affect of parameter selection on performance, a discussion of resource management architectures, methodologies for resource management and approaches to scheduling and prioritisation. Gaps in completed research are identified, to which the research in this thesis is targeted. This chapter details *how* the resource is automatically managed.

An exploration of the critical choice of objective function is given in Chapter 4 for the surveillance and tracking applications. The objective function is crucial for the development of a resource allocation mechanism as it provides the interface to the task function and so determines which attribute of the problem is allocated resource. This includes a discussion of task specific measures, derivation of information theoretic measures and an investigation into the suitability of the derived measures for control.

Agent systems are introduced in Chapter 5 along with some relevant theory for resource allocation problems. The Java Agent Development (JADE) framework is referenced, which is an agent based

extension to the Java platform. Then, development of an agent based multifunction radar resource management testbed using JADE is described. Details of the design and structure of the testbed are also detailed.

The sequential first price sealed bid auction mechanism is applied to the multifunction radar resource management problem in Chapter 6, using the agent based testbed described in Chapter 5. Comparisons are drawn with existing resource management methods using a complex multi-target tracking scenario and with reference to the conclusions from Chapter 4.

In Chapter 7 the continuous double auction mechanism is applied to the multifunction radar resource management problem, which leads into the development of the continuous double auction parameter selection algorithm (CDAPS). Desirable characteristics of the mechanism are demonstrated on multi-target tracking and surveillance scenarios.

Finally the conclusions of the research are presented in Chapter 8 including a discussion of possible future extensions to the work.

1.3 Novel Aspects

The aspects of this work believed to be novel are contained in Chapter. 4-7. Specifically:

- Information theoretic measures for multifunction radar resource management have been derived and developed for estimation and discrimination problems. This has led to an improved understanding of the role of information theoretic measures for multifunction radar resource management and sensor management in general. [Chapter 4, pages 84-90, 94-97]

- The Modified Riccati Equation has been successfully applied to tracking control under significant measurement origin uncertainty. [Chapter 4, pages 93-94]

- An agent based multifunction radar resource management architecture using the JADE framework has been developed. This has provided the basis for a better understanding of agent based resource management architecture designs which allow rapid upgrades and maximum code re-use. [Chapter 5, pages 108-114]

- The sequential first price sealed bid auction mechanism has been applied to multifunction radar resource management including development and analysis of lowest quality first and greatest information first schedulers. This provides a detailed insight into radar resource manager design and selection of appropriate objective functions. [Chapter 6, pages 115-129]

- The continuous double auction mechanism has been applied to multifunction radar resource management leading to the development and assessment of the Continuous Double Auction Parameter Selection (CDAPS) algorithm which generates high performance radar resource management.

[Chapter 7, pages 130-158]

The culmination of these individual aspects constitutes the first application of agent systems to multi-function radar resource management.

1.4 Publications

The following publications are a result of the work in this thesis:

- Charlish, A., Woodbridge, K. and Griffiths, H.: Economic paradigms in cognitive sensor signal processing, *Defence Applications of Signal Processing*, Coolum, Australia, July 2011. (by invitation)

- Charlish, A., Woodbridge, K. and Griffiths, H.: Agent based multifunction radar surveillance control, *IEEE International Radar Conference*, Kansas City, USA, May 2011.

- Ritchie, M., Charlish, A., Woodbridge, K. and Stove, A.: Application of Kullback-Leibler divergence to sea clutter estimation, *IEEE International Radar Conference*, Kansas City, USA, May 2011.

- Charlish, A., Woodbridge, K. and Griffiths, H.: Auction mechanisms in multi-function radar resource management, *Proc. Cognitive Systems with Interactive Sensors*, Crawley, UK, Nov. 2010.

- Charlish, A., Woodbridge, K. and Griffiths, H.: Information theoretic measures for MFR tracking control, *Proc. IEEE International Radar Conference*, pp. 987-992, Washington, USA, May 2010.

Chapter 2

Multifunction Radar

A multifunction radar system is capable of supporting numerous tasks which in turn support differing radar functions. The multi-functionality is primarily enabled by some degree of beam agility, which is predominantly attributable to the use of an electronically steered, phased array antenna. In contrast to a non-agile system where fixed behaviour and hence performance is specified at design time, beam agility allows the performance of the radar system to be adapted during operational deployment. Additionally, as the execution of differing tasks is separable, the signal and data processing applied can be controlled and optimised given the objectives of each specific task.

This chapter describes the theoretical principles of the signal and data processing that can be applied in a multifunction radar. The automatic radar resource manager is required to optimise this processing, which in this thesis is taken as the selection of parameters for all supported tasks, which controls the processing applied. Included in this chapter is fundamental radar theory, the production and processing of radar measurements, and the Electronically Steered Array (ESA) antenna. Finally, an overview of the system's multifunction capability in terms of the variety of task parameters and modes under control is given, alongside examples of operational systems.

2.1   Radar Systems

The Radio Detection And Ranging (RaDAR) system has matured over a period exceeding half a century in a range of civilian and military applications for the ground, airborne and maritime domains. As the name suggests, radar systems provide detection and accurate range measurement of distant or otherwise unobservable objects. The following radar system theory has provided the basis upon which multifunction radar systems are built.

2.1.1   Radar Fundamentals

A radar operates by emitting electromagnetic energy from an antenna, the energy is scattered by the environment, with some of the scattered energy being re-intercepted by the receiving antenna. In the monostatic case, which is assumed throughout this thesis, the transmit and receive antennas are co-

located and potentially a single antenna is used for both transmission and reception. The received signal can be processed to retrieve information on the environment such as the presence and state of a target. Target range can be found by measuring the time taken for a pulse to make the round trip from the antenna, to the target, and back to the antenna. The round trip time (t_d) is proportional to target range (R_t):

$$t_d = \frac{2R_t}{c} \tag{2.1}$$

where c is the speed of the electromagnetic wave propagation. Successive pulses are transmitted at time intervals dictated by the pulse repetition frequency (PRF). Range ambiguities occur, which depend on the PRF, when it is not clear from which of the recently transmitted pulses the received pulse originated. The maximum unambiguous range R_u is proportional to the time interval between pulses t_p:

$$R_u = \frac{ct_p}{2} \tag{2.2}$$

and the time interval between pulses is inversely proportional to the pulse repetition frequency ($t_p = \frac{1}{PRF}$). The radar range resolution (R_r), which is the minimum separation between two targets which are individually resolvable, is inversely proportional to the signal bandwidth B:

$$R_r = \frac{c}{2B} \tag{2.3}$$

For an uncompressed pulse $B = \frac{1}{\tau}$, where τ is the pulse width, in which case the range resolution can be visualised as the two way distance travelled during one pulse duration. Longer pulses allow for an increase in the average transmitted power given a fixed PRF, and so result in a greater detection range. However, as longer pulses result in poorer range resolution, pulse compression is used which increases the bandwidth of the signal to improve range resolution while maintaining the larger energy of a longer pulse.

Assuming a coherent pulse train, target radial velocity, or range-rate, can be found by measuring the Doppler shift on the returned pulses. The doppler shift is proportional to the relative radial velocity v_r between the radar and the target according to:

$$f_d = \frac{2v_r f_c}{c} \tag{2.4}$$

where f_d is the Doppler frequency and f_c is the frequency of the carrier. The spectrum of the pulsed signal contains spikes above and below the carrier frequency at multiples of the pulse repetition frequency. Therefore, doppler ambiguities occur for low PRFs when it is not clear how many multiples of the PRF are contained in the measured doppler shift. The width of each spike determines the doppler resolution,

which is inversely proportional to the duration of the coherent pulse train. High-PRF (HPRF) radar is conventionally defined to give unambiguous doppler measurement, Low-PRF (LPRF) radar is conventionally defined to give unambiguous range measurements and Medium-PRF (MPRF) is conventionally defined to give both ambiguous range and doppler measurement.

Simple analysis of the monostatic radar-target geometry yields the widely used radar range equation which provides an indication of the received power from which the maximum detection range for a given target can be deduced. The geometry is modelled as an antenna radiating power P_t with directional gain G_t, which is intercepted and isotropically reradiated by the target before being re-intercepted by an antenna with receiver gain G_r. Assuming free space with no losses, the received single pulse power P_r can be calculated as the product of three terms [Skolnik, 2008]:

$$P_r = \frac{P_t G_t}{4\pi R_t^2} \cdot \frac{\sigma}{4\pi R_t^2} \cdot \frac{G_r \lambda^2}{4\pi} \qquad (2.5)$$

where σ is the target radar cross section and λ is the wavelength of the carrier. The first term is the power density at range R_t given a transmit power P_t and transmit antenna gain G_t. The second term is the power per unit area at the receiver given a target of radar cross section σ. The final term is the receiver antenna effective area A_e which intercepts the return. Given the minimum detectable signal is S_{min} and incorporating losses L_p this can be rearranged to give the maximum detectable range R_m as:

$$R_m = \sqrt[4]{\frac{P_t G_t G_r \lambda^2 \sigma}{(4\pi)^3 S_{min} L_p}} \qquad (2.6)$$

S_{min} is often limited by thermal noise in the receiver, which can be represented as a termination resistor at the receiver antenna output. This equivalent resistor has an effective temperature T_s (Kelvins) which passes noise with spectral density $N_0 = kT_s$ where $k = 1.38 \times 10^{-23}$ J/K is Boltzmann's constant. The temperature of the noise source is expressed in terms of an ideal source with temperature $T_0 = 290K$, multiplied by a noise factor F_n to account for the non-ideal nature of the receiver. The noise power $N = N_0 B$ where B is the bandwidth of the receiver can be combined with the required signal to noise ratio SNR for detection, to give the minimum detectable signal:

$$S_{min} = kT_0 B F_n SNR \qquad (2.7)$$

where the quantity kT_0 has a convenient round value of 4×10^{-21}. Substitution of Eq. 2.7 into Eq. 2.6 with $SNR = 1$ yields the maximum instrumental range R_0, or the range at which the SNR is unity for an expected radar cross section. Calculation of R_0 allows the radar range equation to be conveniently

expressed as [Blackman and Popoli, 1999]:

$$SNR = \left(\frac{R_0}{R_t}\right)^4 \tag{2.8}$$

giving the SNR as a function of range.

2.1.2 Radar Signal Processing

Given the pulsed operation described in the preceding section, it is required to process the received signal to extract measurement information. It is possible to produce measurement data of target range and radial velocity as well as separable measurements of azimuth and elevation.

2.1.2.1 Range and Radial Velocity

Knowledge of the transmitted signal enables the received signal to be processed to detect and measure a potential time delay and doppler shift which is related to range and radial velocity according to Eq. 2.1 and Eq. 2.4 respectively [Skolnik, 2008]. The transmitted signal $s_t(t)$ is comprised of a sinusoidal carrier, of frequency f_c, which is modulated by a comparatively slowly varying waveform:

$$s_t(t) = g(t)\cos(2\pi f_c t + \phi(t)) \tag{2.9}$$

where $g(t)$ is the amplitude modulation and $\phi(t)$ is the phase modulation of the waveform. This signal is known as a narrow bandpass signal as the signal bandwidth is small compared to the carrier frequency. The complex envelope of this transmit signal $u_T(t)$ is expressed as:

$$u_T(t) = g(t)e^{j\phi(t)} \tag{2.10}$$

which undergoes an additional modulation by the environment, which implants the information to be extracted. Specifically, the transmit signal undergoes a time delay t_d, a shift in frequency f_d and an attenuation in amplitude A_r. The received signal $s_r(t)$ can therefore be expressed as:

$$s_r(t) = A_r g(t - t_d)\cos[2\pi(f_c + f_d)(t - t_d) + \phi(t - t_d)] \tag{2.11}$$

If a target is present a delayed and potentially frequency shifted replica of the complex envelope of the transmit signal $u_T(t - t_d)$ is received. Additionally, the time delay shifts the phase of the received signal by $-2\pi f_c t_d$ due to the oscillations of the carrier and the doppler shift applies a linear phase modulation of $e^{j2\pi f_d(t-t_d)}$. Combining these elements gives the complex envelope of the received signal $u_R(t)$ as:

$$u_R(t) = A_r e^{-j2\pi f_c t_d} u(t - t_d)e^{j2\pi f_d(t-t_d)} \tag{2.12}$$

28

To preserve the phase information the received signal is demodulated by two channels in the receiver. The in-phase (I) channel demodulates by $\cos(2\pi f_c t)$ and the $\pi/2$ out of phase quadrature channel (Q) demodulates by $-\sin(2\pi f_c t)$.

After demodulation the received signal is filtered by a matched filter [North, 1963] which maximises signal to noise ratio by exploiting the knowledge of the transmit signal. Maximum signal to noise ratio in the presence of white noise is achieved when the filter has a frequency response:

$$H_f(f) = k_c U_T^*(f) e^{-j2\pi f t_0} \tag{2.13}$$

where $*$ denotes the conjugate, k_c is a complex constant and t_0 is a time delay required to maintain a casual impulse response. This matches the frequency response of the filter to the expected spectrum of the signal given the known transmit signal. The corresponding impulse response is expressed:

$$h(t) = k_c u_T^*(t_0 - t) \tag{2.14}$$

which is the conjugate of the transmit signal delayed in time. This matched filter produces a maximum possible output SNR [North, 1963] depending on the received bandpass signal energy E_r and noise power spectrum at the filter input N_0:

$$SNR = \frac{2E_r}{N_0} \tag{2.15}$$

however, when incorrectly matched the maximum SNR is not achieved. The autocorrelation function describes the output of a specific matched filter for varying time delay and doppler shifts and can be expressed as [Skolnik, 2008]:

$$\chi(t_d, f_d) = \int_{\infty}^{\infty} u_T(t) u_T^*(t + t_d) e^{j2\pi f_d t} .dt \tag{2.16}$$

Woodward's [Woodward, 1980] ambiguity function follows as the squared magnitude of the autocorrelation function $\Psi(t_d, f_d) = |\chi(t_d, f_d)|^2$. The autocorrelation and ambiguity function describe the fundamental measurement capability of waveform and matched filter by demonstrating the resolution and sidelobe properties as well as allowing measures such as Fisher information to be extracted. Measurement data for range and range rate is produced from the responses of a bank of matched filters within the unambiguous range and doppler limits.

By matching the filter over the pulse train duration, coherent integration is achieved. The phase coherence ensures that the amplitude and phase of target returns are correlated whereas noise returns are uncorrelated. As such, noise returns cancel and target returns combine which gives an improvement in signal to noise ratio (SNR). Maximum practical coherent integration time is limited by target movement,

as the target returns must be integrated in one filter.

Incoherent integration sums the magnitude of the received signal after envelope detection when the phase information is removed. Noise integrates in the same way as target returns and an improvement in SNR is not achieved. Although incoherent integration is less efficient than coherent integration, it is required to integrate the multiple PRF dwells from a single burst, which are used to mitigate range-Doppler blind zones. Also, by averaging the returns over the integration period the signal is low passed filtered and the fluctuation in the noise amplitude reduced. This improves detection sensitivity as the detection threshold multiplier can be lowered without increasing the false alarm probability.

2.1.2.2   Bearing

Estimates of the target's angular location can be produced with sub-beamwidth accuracy by comparing the signals from two or more beams. This can be achieved using sequential beams, where measurement accuracy is hindered by scintillation errors, or preferably using simultaneous beams on a single (mono) pulse.

Amplitude comparison monopulse interpolates using the difference in amplitude between beams slightly separated in angle. Fig 2.1(a) shows the response of two beams $b_1(\theta)$ and $b_2(\theta)$ with $1°$ beamwidth separated by $0.7°$ and Fig. 2.1(b) shows the sum and difference response of the beams. As the magnitude of the difference depends on the target signal amplitude, the difference response $\Delta(\theta)$ is normalised by the sum of the beams $\Sigma(\theta)$ to give the error signal response:

$$k_s(\theta) = \frac{\Delta(\theta)}{\Sigma(\theta)} = \frac{b_1(\theta) - b_2(\theta)}{b_1(\theta) + b_2(\theta)} \tag{2.17}$$

which is shown in Fig 2.1(c).The gradient of this discrimination slope $k_s'(\theta)$ determines the sensitivity of the measurement which is quantified at the point where the measurement slope crosses the measurement axis $k_m = k_s'(0)$.

Thermal noise creates an error in the monopulse measurements as a function of signal to noise ratio as derived by Barton [2004] which can be modelled by the thermal noise error standard deviation σ_θ:

$$\sigma_\theta = \frac{\theta_B}{k_m \sqrt{2SNR}} \approx \frac{\theta_B}{2\sqrt{SNR}} \tag{2.18}$$

where θ_B is the $3dB$ beamwidth. It is common to assume the measurement error standard deviation σ_m is solely due to thermal noise. When the SNR becomes large the measurement error becomes hardware limited and does not continue to reduce.

Targets which are offset from the beam centre experience a loss of gain relative to the maximum antenna gain which, can be approximately modelled by reducing the target SNR according to a Gaussian

(a) Voltage response of beams b_1 and b_2

(b) Voltage response of sum and difference beams

(c) Monopulse error slope
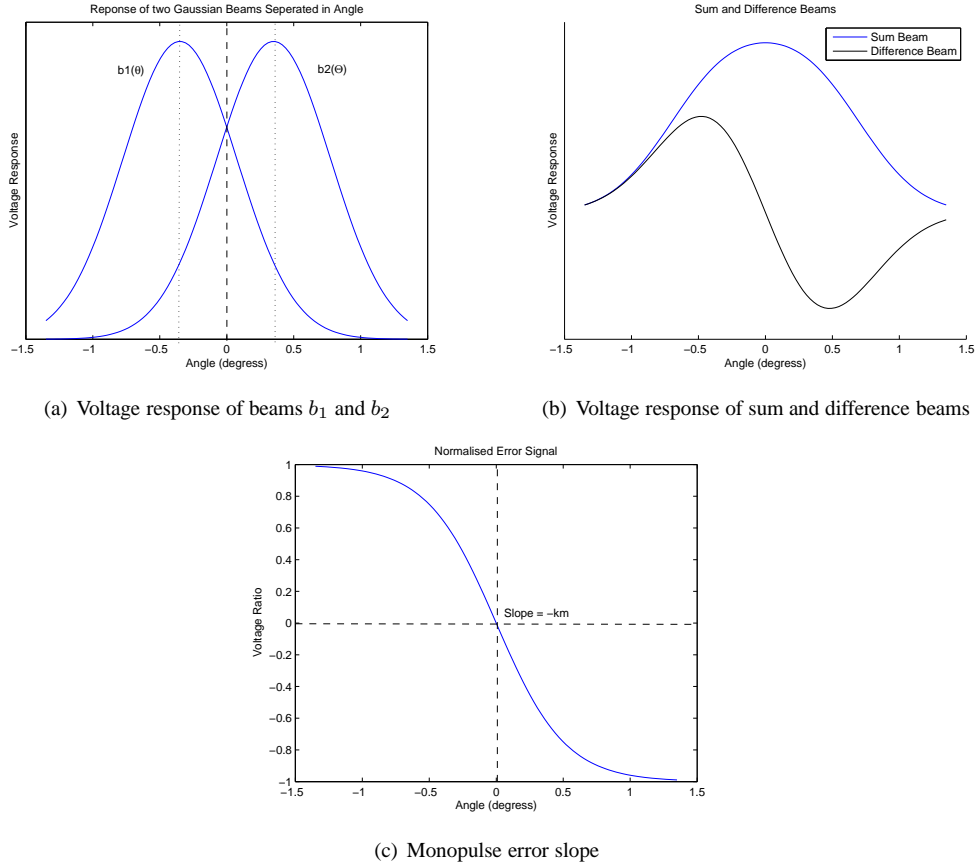
Figure 2.1: Process of monopulse measurement

loss function [Blackman and Popoli, 1999]:

$$SNR = SNR_0 \exp\left[\frac{-C_L[(\eta_T - \eta_P)^2 + (\epsilon_T - \epsilon_P)^2]}{\theta_B^2}\right] \qquad (2.19)$$

where SNR_0 is the beam centre SNR, η_T and η_P are true and predicted azimuth and ϵ_T and ϵ_P are the true and predicted elevations. C_L can be taken as 2.77 which is found by substituting $\frac{SNR}{SNR_0} = 0.5$ when the angle off boresight ($\sqrt{(\eta_T - \eta_P)^2 + (\epsilon_T - \epsilon_P)^2}$) is equal to half the half power beamwidth [Blackman and Popoli, 1999]. This loss in SNR affects the measurement accuracy as defined by Eq. 2.18.

In addition to the loss in SNR, the accuracy of the monopulse measurement degrades as the target is off the centre of the measurement axis, as evident by the reduction in sensitivity visible in Fig. 2.1(c). The off-boresight measurement accuracy $\tilde{\sigma}_\theta$ can be modelled by including a second component of thermal noise error, which causes a scaling of the on-boresight measurement accuracy :

$$\tilde{\sigma}_\theta = \sigma_\theta \sqrt{1 + (k_m \frac{\theta_f}{\theta_B})^2} \qquad (2.20)$$

31

where θ_f is the offset angle.

2.1.3 Electronically Steered Array

The Electronically Steered Array (ESA) [Stimson, 1998; Wirth, 2001; Skolnik, 2008] is able to provide the multifunction radar system's requirement for beam agility. The ESA is an antenna with an array of radiating elements which have controllable phase and amplitude as shown in Fig. 2.2. Modern electronic components allow for the array control to be rapidly applied which enables an agile and flexible beam.
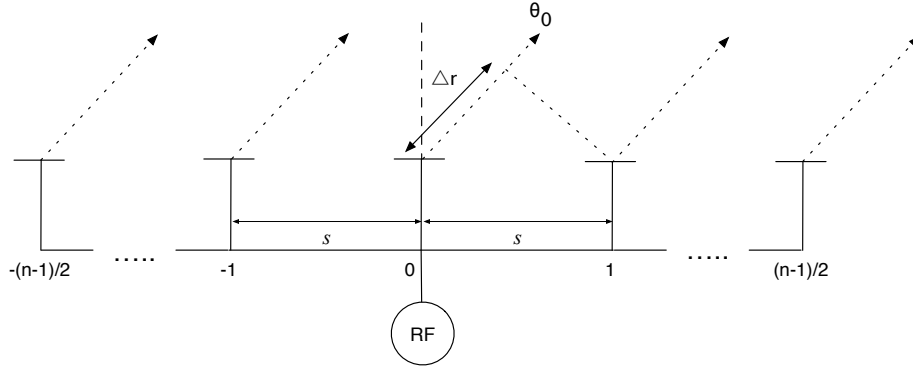
Figure 2.2: Linear electronically steered array, steering at angle θ_0 [Wirth, 2001]

An array of n elements with linear spacing s each isotropically radiating equal amplitude and phase produces a radiation pattern which can be found by summing the vector contributions of all the elements. The subsequent radiation pattern $E_\alpha(\theta)$ is [Skolnik, 2008]:

$$E_\alpha(\theta) = \frac{\sin[n\pi(s/\lambda)\sin\theta]}{n\sin[\pi(s/\lambda)\sin\theta]} \qquad (2.21)$$

and is plotted in Fig. 2.3(a) for 10 and 20 elements with a spacing of $\lambda/2$. The main lobe is clearly identifiable at $\theta = 0$ with additional side lobes. The $3dB$ beamwidth in radians is a function of the wavelength λ and the length of the aperture l in the relevant dimension:

$$\theta_B = \frac{0.886\lambda}{l} \qquad (2.22)$$

which is evident in Fig. 2.3(a) where increasing the number of elements to 20 creates a longer aperture which reduces the beamwidth.

Fig. 2.3(b) shows the radiation pattern for 10 elements with a 1.5λ spacing. Additional main beams called grating lobes can be seen at $\pm 0.2323\pi$. Grating lobes occur, due to spatial under-sampling, at angles θ_g determined by the element spacing in relation to the wavelength:

$$\sin\theta_g = \pm\frac{m\lambda}{s} \qquad (2.23)$$

32

where m is an integer $m = 1, 2, 3.......$

(a) 10 and 20 elements spaced at $\lambda/2$ (b) 10 elements spaced at 1.5λ
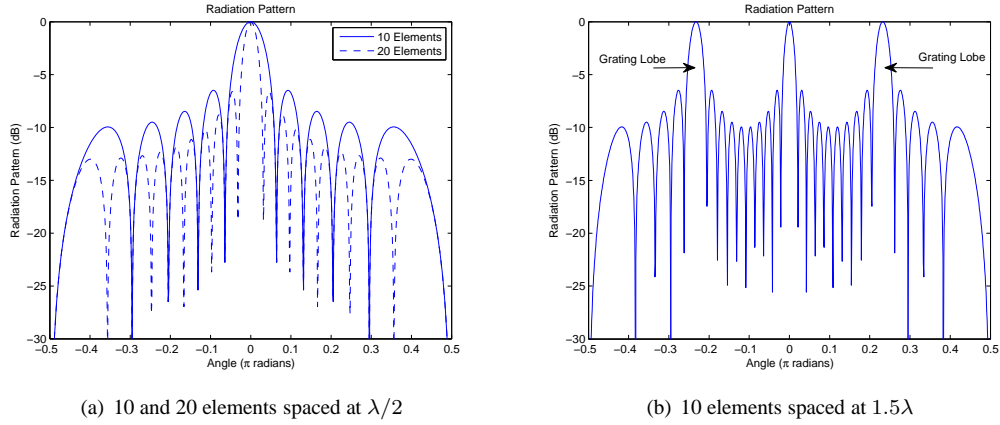
Figure 2.3: Linear array radiation patterns

For a rectangular array, the gain of the main beam is a function of the wavelength and effective aperture area and assuming the same transmit and receive antennas are used:

$$G_t = G_r = 4\pi\frac{A_e}{\lambda^2} \tag{2.24}$$

where the effective aperture area is related to the true area by the aperture efficiency η, i.e $A_e = \eta A_T$ where A_T is the actual aperture area.

In reality each element has a non-isotropic radiation pattern which necessitates the inclusion of the element factor E_e to produce the complete radiation pattern $E(\theta)$:

$$E(\theta) = E_e(\theta)E_\alpha(\theta) = E_e(\theta)\left[\frac{\sin[n\pi(s/\lambda)\sin\theta]}{n\sin[\pi(s/\lambda)\sin\theta]}\right] \tag{2.25}$$

The pointing angle of the main beam can be steered, as shown in Fig. 2.2, by applying linear phase increments $\triangle\varphi$ between each element. It can be seen that the phase difference required to steer the beam at angle θ_0 is the difference in phase over the distance $\triangle r$. As there are 2π radians in one wavelength and $\triangle r = s\sin\theta_0$ the difference in phase between elements required to steer at an angle θ_0 is:

$$\triangle\varphi = \frac{2\pi s\sin\theta_0}{\lambda} \tag{2.26}$$

The linear array radiation pattern in Eq. 2.21 can be adapted to include the effect of beam steering at angle θ_0 to give:

$$E(\theta) = E_e(\theta)\frac{\sin[n\pi(s/\lambda)(\sin\theta - \sin\theta_0)]}{n\sin[\pi(s/\lambda)(\sin\theta - \sin\theta_0)]} \tag{2.27}$$

this radiation pattern is shown in Fig. 2.4 for a beam steered at $\theta_0 = -\pi/4$ in comparison to a linear array

where no beam steering is applied. When scanning off the radar boresight the effective aperture length l is reduced by $\cos\theta_0$ which increases the beamwidth according to Eq. 2.22. However, mutual coupling and the non-isotropic nature of the array elements causes the one way gain to drop off by approximately $cos^{1.5}\theta$ [Sabatini and Tarantino, 1994].
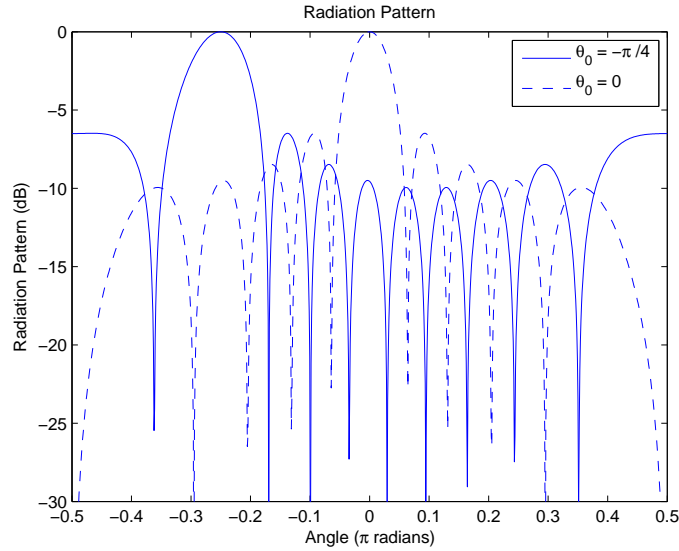
Figure 2.4: Linear array radiation pattern steered at $\theta_0 = -\pi/2$

When radiating elements are closely spaced, energy is coupled between elements which affects the each element's radiation pattern. The magnitude of the coupling depends on the distance between the elements and the distribution pattern of the elements in the array. For an ESA where there are numerous closely spaced elements the effect of mutual coupling can be strong and can result in loss of the main beam, the magnitude of the loss depending on the coherent combinations of the coupling signals between elements in the array.

The spacing between elements is dictated by the desire to avoid grating lobes, which must not appear within the field of view (FOV) when the array is steered to the maximum scan angle. Grating lobes are avoided if:

$$\frac{s}{\lambda} < \frac{1}{1 + |\sin\theta_0|} \tag{2.28}$$

which gives a maximum spacing of $\frac{\lambda}{2}$ for a $\pm 90°$ FOV, 0.536λ for a $\pm 60°$ FOV and 0.586λ for a $\pm 45°$ FOV. Loss of gain due to off boresight scanning typically limits the field of view to $\pm 60°$ or $\pm 45°$ and so several arrays or rotating arrays must be used for full $360°$ coverage.

2.2 Measurement Data Processing

The fundamental radar system described in the preceding section produces range, doppler and angle measurement data. However, the multifunction radar system must apply significant data processing

before it can be presented to the operator in a meaningful way. This processing involves the automatic detection and tracking of targets within the measurement data. The multifunction radar system is required to optimise the data processing given the objectives of each individual task.

2.2.1 Detection

Detection is a binary hypothesis testing problem to differentiate between the target present hypothesis H_T and target not present hypothesis H_N. The Neyman-Pearson lemma defines the optimal decision region for a fixed probability of false alarm P_{FA} as a threshold T on likelihood ratio LR for data vector $D = \{d_1, ..., d_n\}$:

$$LR(d_1, ..., d_n) = \frac{p(d_1, ..., d_n|H_T)}{p(d_1, ..., d_n|H_N)} \mathop{\gtrless}_{\substack{>^{H_T} \\ <_{H_N}}} T \tag{2.29}$$

where T is chosen so that $p(d_1, ..., d_n > T|H_N) = P_{FA}$. The optimality condition ensures the probability of detection P_D is maximised for the fixed false alarm probability.

When the receiver is dominated by thermal noise, the target not present hypothesis relates to an output of the I and Q channels according to a complex Gaussian probability density function [Ward et al., 2006]. The corresponding envelope of the signal $E = \sqrt{E_I^2 + E_Q^2}$, which is the output from a linear envelope detector, is characterised by a Rayleigh probability density function:

$$P_N(E) = \frac{2E}{\bar{z}_n} \exp\left(\frac{-E^2}{\bar{z}_n}\right) \tag{2.30}$$

where $\bar{z}_n$ is the mean noise intensity. The target present hypothesis can be assumed to be a coherent signal embedded in the thermal noise, which produces a signal with envelope characterised by a Rician probability density function:

$$P_T(E|A) = \frac{2E}{\bar{z}_n} \exp\left(-\frac{E^2 + A^2}{\bar{z}_n}\right) I_0\left(\frac{2EA}{\bar{z}_n}\right) \tag{2.31}$$

where A is the amplitude of the signal and I_0 is the modified Bessel function in the first kind with zero order. Fig. 2.5 shows the probability density functions for the envelope of thermal noise and target plus thermal noise. It can be shown through the Neyman-Pearson lemma that thresholding on the envelope of the measurement data is optimal for large signals [Skolnik, 2008]. An example of such a threshold on the signal envelope is marked in Fig. 2.5.

Alternatively, a square law envelope detector can be used which produces an output proportional to the intensity $z = E^2$ of the signal. In this case the target not present hypothesis, which corresponds to thermal noise, has an exponential probability density function:

$$P_N(z) = \frac{1}{\bar{z}_n} \exp\left(\frac{-z}{\bar{z}_n}\right) \tag{2.32}$$
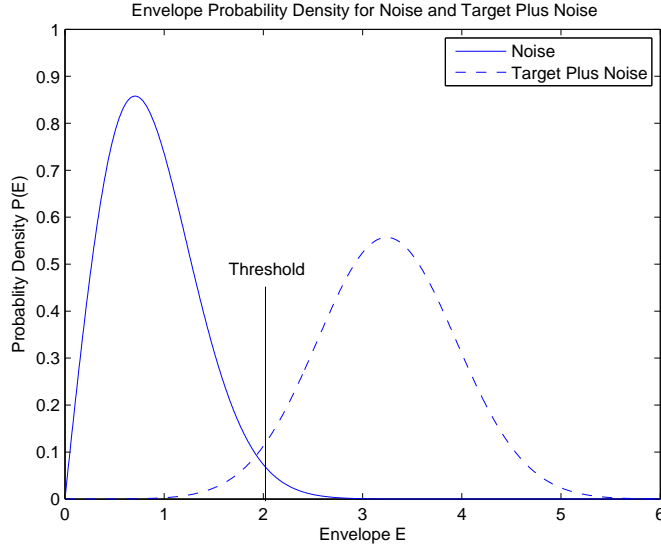
Figure 2.5: Probability density functions for the envelope of noise and target plus noise

and the target present hypothesis, which can be assumed as a coherent signal embedded in thermal noise, is characterised by:

$$P_T(z|A) = \frac{1}{\bar{z}_n} \exp\left(-\frac{z + A^2}{\bar{z}_n}\right) I_0 \left(\frac{2A\sqrt{z}}{\bar{z}_n}\right) \tag{2.33}$$

It can be shown through the Neyman-Pearson lemma that thresholding on the intensity of the measurement data is optimal for small signals. However, the linear and square law detectors exhibit similar detection performance.

2.2.1.1   Calculation of Detection Probability

The probability of detection and false alarm for threshold T can be seen with reference to Fig. 2.5 as the integral of the respective probability density functions above the threshold. Marcum [Marcum, 1947, 1948] investigated this statistical nature of radar measurement data and produced functions for calculating the probability of detection and probability of false alarm for a number of incoherently integrated pulses.

Marcum gave the probability of detection of a single normalised pulse as:

$$P_D(A, T) = \int_T^\infty e^{-(z+A^2)} I_0 \left(2\sqrt{z}A\right) .dz = Q(\sqrt{2}A, \sqrt{2T}) \tag{2.34}$$

$$P_{FA} = e^{-T} \tag{2.35}$$

where Q is Marcum's Q-function.

2.2.1.2 Target Fluctuation

Swerling extended Marcum's work to consider the non-fluctuating or Swerling 0 case and four different cases of fluctuating target radar cross section and hence signal to noise ratio. In case one and two the target is modelled as a number of independent scatters, no one of which is dominant, and is used to describe large complex targets. The radar cross section fluctuations follow a Rayleigh, or chi-squared with two degrees of freedom, probability density function:

$$w(\sigma, \bar{\sigma}) = \frac{1}{\bar{\sigma}} \exp\left(\frac{-\sigma}{\bar{\sigma}}\right) \tag{2.36}$$

where $\bar{\sigma}$ is the mean radar cross section. For case one the fluctuations occur between scans and for case two the fluctuations occur between pulses. The probability of detection for cases one and two can be calculated as a function of the false alarm probability and the SNR:

$$P_D = P_{FA}^{1/(1+SNR)} \tag{2.37}$$

Cases three and four model the target as a single large dominant scatterer surrounded by a number of smaller scatters, which is assumed to be characterised by a Rician probability density function:

$$w(\sigma, \bar{\sigma}) = \frac{4\sigma}{\bar{\sigma}^2} \exp\left(\frac{-2\sigma}{\bar{\sigma}}\right) \tag{2.38}$$

for case three the fluctuations occur between scans and for case four the fluctuations occur between pulses. These four cases are summarised in Table 2.1.

Table 2.1: Swerling target fluctuation models

	Scan to Scan	Pulse to Pulse
Many Small	Case 1	Case 2
One Large	Case 3	Case 4

2.2.1.3 False Alarm Control

In reality a global homogenous background is rarely faced and so the use of a fixed global threshold would produce local regions of excessive false alarms which overload the data processor and tracker. To avoid this it is necessary to estimate the statistics of the local background to apply a dynamic threshold. Typically a model of the probability density of the background is known, with potentially unknown parameters. A constant false alarm rate detector uses a set of local background reference cells to estimate the unknown parameters of the model. The test and reference cells are separated by a number of guard cells to ensure the target is not present in the reference.

For a Rayleigh background it is sufficient to estimate the mean in order to set the correct local

threshold to maintain a specified probability of false alarm. This can be implemented through a Cell-Averaging (CA)-CFAR which is shown in Figure 2.6. In the CA-CFAR a number of reference cells, seperated from the cell under test by guard cells, are used to estimate the mean intensity of the local background. This mean is multiplied by the threshold multiplier, to produce the intensity threshold which must be exceeded to declare the target presence.
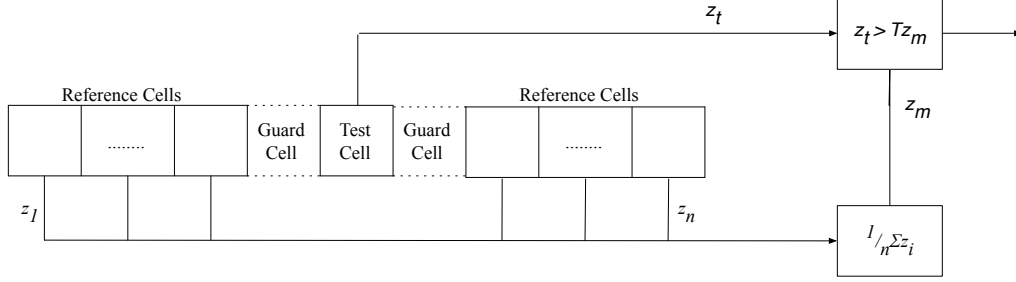
Figure 2.6: Cell averaging constant false alarm rate detector

2.2.2 Tracking

Tracking is the process of fusing sequences of detected measurements to estimate the kinematics of the underlying targets. It is required to estimate the state of the target as a kinematic parameter vector X, such as the position and velocity in cartesian coordinates, i.e. $X = (x, x', y, y')^T$. The kinematic state of the target x_k is assumed to evolve as a potentially non-linear discrete time stochastic system described [Ristic et al., 2004] by the dynamic equation:

$$x_k = f_{k-1}(x_{k-1}, v_{k-1}) \tag{2.39}$$

where f_{k-1} describes the predictable disturbances to motion, x_{k-1} is the previous state and v_{k-1} is a noise sequence which allows for unpredicted disturbances to motion. Measurements z_k are used to estimate the target state x_k which are received corrupted by measurement noise w_k and so are modelled by the measurement equation:

$$z_k = h_k(x_k, w_k) \tag{2.40}$$

where h_k is the observation function and both the process noise v_k and measurement noise w_k are assumed known. Fig. 2.7 shows a block diagram of the discrete time sequential state estimation process, adapted from Bar-Shalom et al. [2001]. Bayes' theorem provides the framework for sequential state estimation, enabling new measurements to be fused with estimates from previous time steps:

$$p(x|Z_k) = \frac{p(z_k|x_k)p(x_k|Z_{k-1})}{p(z_k|Z_{k-1})} \tag{2.41}$$

38

where $p(z_k|x_k)$ is the likelihood function, $p(x_k|Z_{k-1})$ is the state estimate at time $k-1$ and $p(z_k|Z_{k-1})$ is a normalising constant. This can be re-arranged to give the optimal recursive Bayesian estimator [Ristic and Hernandez, 2008]:

$$\text{Prediction: } p(x_{k+1}|Z_k) = \int p(x_{k+1}|x_k)p(x_k|Z_k) \tag{2.42}$$

$$\text{Update: } p(x_{k+1}|Z_{k+1}) = \frac{p(x_{k+1}|x_k)p(x_k|Z_k)}{\int p(z_{k+1}|x_{k+1})p(x_{k+1}|Z_k).dx_{k+1}} \tag{2.43}$$

The state of the underlying system can be estimated as the minimum mean square estimate, which is the conditional mean of the state estimate:

$$\hat{x}_{k|k}^{MMSE} = E[x_k|z_k] = \int x_k.p(x_k|Z_k).dx_k \tag{2.44}$$
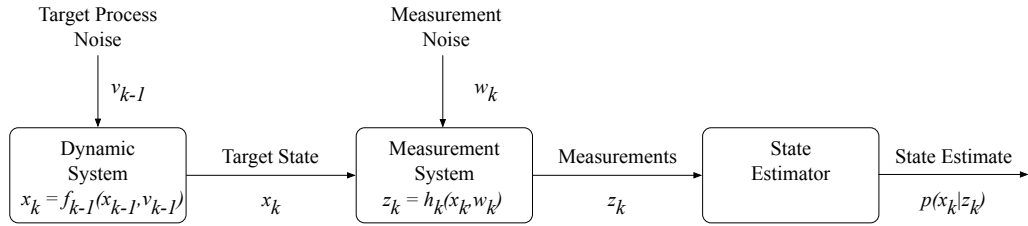
Figure 2.7: Target dynamic system, measurement system and sequential state estimation

This optimal estimator requires the propagation of the entire posterior probability density which is a potentially infinite data vector. As a result it is necessary to find closed form or sub-optimal solutions. The Kalman filter is a closed form solution which can be used if the dynamic and measurement models are linear and the process and measurement noises are Gaussian, with covariances denoted $E[v_k v_k'] = Q_k$ and $E[w_k w_k'] = R_k$. As the state estimate is Gaussian it is completely described by its first two moments, the conditional vector mean Eq. 2.44 and the covariance matrix:

$$P_{k|k} = E[[x_k - \hat{x}_{k|k}][x_k - \hat{x}_{k|k}]'|Z_k] \tag{2.45}$$

The Kalman filter breaks the sequential estimation of the mean and covariance into prediction and update stages as follows:

Stage 1 - Prediction

$$\text{State estimate prediction: } \hat{x}_{k|k-1} = F_{k-1}\hat{x}_{k-1|k-1} \qquad (2.46)$$

$$\text{State covariance prediction: } P_{k|k-1} = F_{k-1}P_{k-1|k-1}F'_{k-1} + Q_{k-1} \qquad (2.47)$$

$$\text{Prediction of next measurement: } \hat{z}_{k|k-1} = H_k\hat{x}_{k|k-1} \qquad (2.48)$$

$$\text{Innovation covariance calculation: } S_k = R_k + H_k P_{k|k-1}H'_k \qquad (2.49)$$

Stage 2 -Update (On receiving measurement z_k)

$$\text{Measurement residual calculation: } \tilde{z}_k = z_k - \hat{z}_{k|k-1} \qquad (2.50)$$

$$\text{Filter Gains Calculation: } W_k = P_{k|k-1}H'_k S_k^{-1} \qquad (2.51)$$

$$\text{State estimate update: } \hat{x}_{k|k} = \hat{x}_{k|k-1} + W_k\tilde{z}_k \qquad (2.52)$$

$$\text{State covariance update: } P_{k|k} = P_{k|k-1} - W_k S_k W'_k \qquad (2.53)$$

The linear-Gaussian assumptions at the heart of the Kalman filter rarely hold in reality and so it is often necessary to use sub-optimal methods. Non-linear models can be incorporated using the extended Kalman filter, which approximates the non-linear functions, or the unscented Kalman filter, which approximates the posterior distribution as a Gaussian sum.

2.2.3 Kinematic Models

The implementation of the Kalman filter requires the specification of models for the target dynamics and the measurement system. The measurement system model can readily be derived from Sec 2.1.2, however, the choice of target process noise covariance and dynamic system transition matrix are not so apparent. Various models exist which are suited to differing target dynamics and the subsequent choice can have significant effect on tracking performance.

2.2.3.1 Continuous White Noise Models

Continuous white noise models assume that deviations to predictable motion enter the system as zero mean white noise *process* $\tilde{\nu}(t)$, i.e. $E[\tilde{\nu}(t)] = 0$ and $E[\tilde{\nu}(t)\tilde{\nu}(\tau)] = \tilde{q}\delta(t-\tau)$. The process noise intensity $\tilde{q}$ represents the strength of the deviations from predictable motion. For an arbitrary position coordinate ξ, the white noise can enter the system as a white noise acceleration, $\ddot{\xi}(t) = \tilde{\nu}(t)$ or white noise jerk (derivative of acceleration) $\dddot{\xi}(t) = \tilde{\nu}(t)$. As derived by Bar-Shalom et al. [2001] this gives the following system transition and process noise covariance matrices:

Continuous white noise acceleration

$$\text{Transition matrix: } F_k = \begin{bmatrix} 1 & T_k \\ 0 & 1 \end{bmatrix} \tag{2.54}$$

$$\text{Process Noise Covariance: } Q_k = E[v_k v_k'] = \begin{bmatrix} \frac{T_k^3}{3} & \frac{T_k^2}{2} \\ \frac{T_k^2}{2} & T_k \end{bmatrix} \tilde{q} \tag{2.55}$$

Continuous white noise jerk

$$\text{Transition matrix: } F_k = \begin{bmatrix} 1 & T_k & \frac{1}{2}T_k^2 \\ 0 & 1 & T_k \\ 0 & 0 & 1 \end{bmatrix} \tag{2.56}$$

$$\text{Process Noise Covariance: } Q_k = E[v_k v_k'] = \begin{bmatrix} \frac{T_k^5}{20} & \frac{T_k^4}{8} & \frac{T_k^3}{6} \\ \frac{T_k^4}{8} & \frac{T_k^3}{3} & \frac{T_k^2}{2} \\ \frac{T_k^3}{6} & \frac{T_k^2}{2} & T_k \end{bmatrix} \tilde{q} \tag{2.57}$$

where T_k is the time between time steps. When the process noise intensity $\tilde{q}$ is small these represent a nearly constant velocity (NCV) and nearly constant acceleration (NCA) model respectively.

An alternative family of models called discrete white noise models allow for deviations to motion to enter the system as a zero mean white noise *sequence*. This manifests itself as a constant acceleration or constant jerk over the sampling period, which is uncorrelated to the previous time step. In this work the previous continuous white noise models are preferred as the same amount of process noise enters the system regardless of the length of the sampling interval [Blackman and Popoli, 1999], i.e.:

$$F_k.Q_k(T_k).F_k + Q_k(T_k) = Q_k(2T_k) \tag{2.58}$$

which is a useful property for adaptive update rate tracking.

2.2.3.2   Singer

Singer [Singer, 1970] provides a more realistic model of a correlated acceleration sequence between time steps represented as a Markov process:

$$\dddot{\xi}(k+1) = \rho_m \dddot{\xi}(k) + \sqrt{1 - \rho_m^2}\,\Omega n_1 \tag{2.59}$$

where $\rho_m = e^{-\beta_m T_k}$, $\beta_m = \frac{1}{\Theta}$ and Θ is the target manoeuvre time constant, Ω is the target manoeuvure standard deviation and n_1 is a zero mean unit standard deviation Gaussian distributed random variable. In the limit where the sampling interval is much less than the manoeuvre time constant, the Singer model tends to the continuous white noise jerk model in Eq. 2.56 and Eq. 2.57. In the opposite case where the sampling interval is much greater than the manoeuvre time constant then estimates of the acceleration are not possible and so the Singer model tends to the continuous white noise acceleration model in Eq. 2.54 and Eq. 2.55.

2.2.3.3 Adaptive Filtering

As target dynamics are likely to change over the track duration, it is necessary to implement adaptive filtering methods which change the model of the target dynamic upon manoeuvre to ensure the filter is matched to the current target dynamic.

Reactive Adaptation - The residual vector from the tracking filter can be monitored to detect manoeuvres. If the residual becomes large, as defined by some rule of thumb, then the process noise can be increased to reduce the smoothing applied by the filter and to apply more weight to new measurements.

Variable Dimension Filtering - When the manoeuvre detection logic indicates a manoeuvre, the dimension of the filter state can be changed. For example in periods of benign motion a NCV can be adopted, which can be changed to a NCA upon manoeuvre.

Multiple Model Filtering - Kalman filters with differing models are run in parallel, the residual is monitored to determine the probability of each of the models being correct. The output is each of the filter outputs is merged by the filter probability.

From the available adaptive filtering techniques Interacting Multiple Model (IMM) has emerged as the best performer but with the greatest complexity and computational cost.

2.2.4 Data Association

It was previously assumed that the Kalman filter was updated with a measurement that was known to be from the target in question. In reality the measurement could also have originated from some form of interference or from a different, nearby target. As such data association techniques are used to improve correct measurement to track assignment.

2.2.4.1 Gating

To reduce the complexity of data association a gate is applied to discard unlikely target to track pairings. The gate is centered on the tracks predicted state, and only detections falling within this gate are considered for assignment. Rectangular and ellipsoidal gates can be used, the size of which is determined by the residual vector in the track. An ellipsoid gate specifies a valid association region within the statistical distance d^2:

$$d^2 = \tilde{z_k}' S_k^{-1} \tilde{z_k} \leq g \tag{2.60}$$

where $\tilde{z}$ is the measurement residual from Eq. 2.50 and g is the gate size. The volume V_k of the validation gate is given by:

$$V_k = \pi g^2 |S_k|^{1/2} \tag{2.61}$$

where S_k is the innovation covariance from Eq. 2.49.

2.2.4.2   Data Association Methods

Measurements falling within the validation gate are eligible for track update. Common methods for data association are global nearest neighbour, probabilistic data association and multi-hypothesis tracking, which have increasing complexity and effectiveness.

Global Nearest Neighbour (GNN) - Nearest neighbour assigns an observation to a track whereby the subsequent assignment minimises the statistical distance of all possible observations to that track. GNN performs this process for all tracks in the system and hence minimises the global statistical distance for all observation to track assignments. GNN is the simplest approach to data association but performs poorly in high clutter or dense target scenarios.

Probabilistic Data Association and Joint Probabilistic Data Association - Probabilistic Data Association (PDA) forms hypotheses on all possible observation to track assignments falling in the gate. The probability of each of these being the correct assignment is calculated and the hypotheses are merged, weighted by the respective probabilities. Joint Probabilistic Data Association (JPDA) extends PDA by calculating the global probabilities of all observations and all tracks. PDA performs better in clutter than GNN and JPDA performs better than PDA in multi-target situations. Both have extra computational cost over GNN.

Multi-Hypothesis Tracking - Multi-Hypothesis Tracking (MHT) forms hypotheses for observation to track assignments which are not merged at each scan as in JPDA. Hypotheses are propagated so that future scans resolve the uncertainty in previous time steps. This produces a branching tree of hypotheses, each with a probability of being correct. This tree is managed so that unlikely hypothesis branches are pruned to manage computation. MHT performs better in clutter and dense target regions but at an added computational cost.

2.2.5   Track Management

Track life cycles must be monitored for tracks to be correctly started, terminated and maintained in the tracking system. Stages in the life cycle of a track can include alert, confirmation, initiation, tentative track, established track or deleted track. Additional events may also occur over the duration of the track life such as track splitting and track merging. Policies for handling track life cycles are given in Blackman and Popoli [1999].

Two methods for determining the track status are:

Logic - In logic based track management systems the status of a track is determined by some pre-defined logic. Based on this logic, rules are designed to determine the status of tracks. For example, a simple rule to trigger track initiation may be two detections out of three, or an example condition for track deletion may be three missed detections. More complicated Markov chains can be constructed to define rules for statuses such as tentative, preliminary or confirmed track.

Track Scoring - Track scoring methods calculate the likelihood ratio of the hypothesis that a true target is present against the hypothesis the returns are due to interference [Blackman and Popoli, 1999]:

$$LR = \frac{p(D|H_T)P_0(H_T)}{p(D|H_N)P_0(H_N)} \equiv \frac{P_T}{P_{FA}} \tag{2.62}$$

Where H_T and H_N are the presence of true target and false alarm respectively, given data D. This is discussed further in Sec 4.1.2.4.

2.3 Multifunction Radar Systems

The multifunction radar described in the preceding section is required to control and optimise the numerous tasks which support differing radar functions. This control can be applied as task parameter and mode selection which is described in this section. Example operational systems are also described.

2.3.1 Control Parameters

Optimising the configuration of the multifunction radar for each task involves the selection of a set of radar control parameters. There is a large number of parameter dimensions under control in a typical multifunction radar system, which are listed in Table 2.2

Table 2.2: Multifunction radar task parameters

Parameter	Description	Impact
RF Frequency	Frequency of the carrier	• Choice of frequency motivated by utilising frequency diversity which allows mitigation of interference and environmental losses.

Continued on next page

Table 2.2 - Continued from previous page.

Parameter	Description	Impact
PRF (Hz)	Frequency of pulses in burst	• Increasing PRF increases energy on target up to the maximum allowable duty cycle but also increases eclipsing loss for a fixed pulse width. • Increasing PRF increases unambiguous doppler range but decreases unambiguous range. • Multiple PRFs selected in a burst to mitigate range-Doppler blind zones.
Pulse Width τ (secs)	Width of modulating pulse.	• Increasing pulse width increases energy on target up to the maximum allowable duty cycle but also increases eclipsing loss for a fixed PRF.
Pulse Compression ρ	Ratio of compressed to uncompressed pulse.	• Increasing pulse compression increases the signal bandwidth.
Coherent Integration (secs)	Duration of coherent integration period.	• Increasing coherent integration time improves frequency resolution and SNR. • Practical integration time limited by target movement.
Non-coherent Integration	Number of non-coherent integrations.	• Increasing non-coherent integrations enables multiple dwells per burst, improving detection probability for diverse targets. • Increasing non-coherent intregrations reduces noise amplitude fluctuations which improves detection sensitivity.
Time on Target (secs)	Coherent and non-coherent dwell duration (secs).	• Increasing time on target improves detection performance.

Table 2.2 - Continued from previous page.

Parameter	Description	Impact
Detection Threshold	Threshold value for target present declaration.	• Increasing threshold reduces P_{FA} but also reduces P_D. • Decreasing threshold increases P_D but also increases P_{FA}.
Average Transmit Power P_{av} (W)	Peak power multiplied by duty factor.	• Increasing average power increases SNR. • Assumed maximum peak power used for waveform given the duty factor constraint of the hardware.
Surveillance Pattern	Geometry of beams within surveillance region.	• Increasing beam spacing increases nulls in search pattern but reduces search loading.
Tracking Beam Pointing	Active track update pointing angle.	• Directable to predicted target position for active tracking.
Beamwidth θ_B (rad)	$3dB$ angular width of beam.	• Minimum beam width maximises power aperture product. • Wider beam is less sensitive to track uncertainty.
Task Revisit Interval (t_f - secs)	Time interval between task dwells.	• Smaller revisit interval improves task quality but increases task loading.
Signal Processing	Choice of processing applied	• Correct choice of signal processing (STAP, GMTI, SAR etc.) applied improves performance for specific situations or objectives.
Measurement Data Processing	Choice of filter, manoeuvre model and data association parameters.	• Correct choice impacts quality of information presented to operator, e.g. tracking error.

This choice of parameters can be thought of as the finite radar resource to be optimised. The choice within each parameter dimension as well as the large number of dimensions renders the optimal parameter choice, which relates to effective resource management, a very challenging problem.

The simplest parameter set for an arbitrary task can be taken as a beam pointing direction, dwell

length τ_d and revisit interval t_f. The revisit interval is the time between successive dwells and the dwell length is the time over which a beam position is illuminated and some integration is performed. In this thesis it is assumed that the integration efficiency is ideal, equivalent to coherent integration, which improves the signal to noise ratio. Additionally in this work a function is defined as a purpose or capability of the system, e.g. tracking, a task is defined as the realisation of a function, e.g. tracking of a specific target, and a look is defined as a dwell which supports a task.

2.3.2 Operational Modes

To reduce the parameter search space, operational modes can be defined that contain a smaller range of tuneable parameters. This enables the prior knowledge of years of research and experimental knowledge to be built into the system and prevents the resource manager from unnecessarily rediscovering parameter selection online. The suite of modes depend on the application domain, a broad description of potential modes is given in this subsection.

2.3.2.1 Surveillance

Surveillance tasks survey volumes in space with the aim of discovering new targets or discovering there are no new targets.

Long Range Search - Search of a specified region out to a long range with the aim of maximising the cumulative detection range. Requires longer dwell times to detect at long ranges.

Medium Range Search - Search of a specified region or area of interest with the aim of detecting targets but also producing measurements of good quality kinematic accuracy and so requires waveforms giving reasonable range and doppler measurements.

Self Protect Search - Search of a region with the aim of detecting close in 'pop-up' targets such as a missile breaking the horizon. Requires a rapid revisit interval and high single hit probability of detection.

Track-While-Scan (TWS) associates observations from the same target over multiple surveillance scans to present to the operator as a single track. By using measurements from surveillance scans this allows tracking and surveillance to be performed simultaneously. The TWS algorithm is required to perform the filtering, data association and management operators discussed in Sec. 2.2.2, to correctly associate observations to targets and reject false returns originating from clutter. TWS is a resource efficient method for tracking multiple targets and so active tracking should only be performed when necessary to augment the performance achieved through TWS.

2.3.2.2 Tracking

Tracking functions support the fusing of detections to maintain an estimate of target kinematics in a volume of interest.

Track Update - Standard tracking mode to produce a measurement for an active track. Time on target and revisit interval depend on target kinematics and required track quality.

Track Maintenance - Rapid revisit or search around the targets predicted position following a missed detection.

Track Initiation/Confirmation - Track initialisation mode follows the alert confirm stage by requesting a sequence of rapid revisit measurements with sufficient kinematic accuracy to initialise the tracking filter. Confirmation occurs when returns can be certified as originating from a target dynamic model and not from clutter.

Track Splitting/Merging - Poor resolution may cause multiple targets to be represented by a single track. Hence tracks may merge when two targets become unresolvable or split as the targets become resolved. In either of these events it is required to schedule an additional initialisation period to stabilise the kinematic estimates in the tracks.

2.3.2.3 Situational Assessment

Situational assessment functions are motivated by gathering information to improve future resource allocation, or determining the current state of mission objectives.

Target Identification/Recognition - Recognition and identification of non-cooperative targets.

Target Acquisition - Variety of modes which can be used to acquire targets which are complementary to mission objectives, such as Synthetic Aperture Radar (SAR), Inverse-SAR (ISAR) or Ground Moving Target Indicator (GMTI).

Raid Assessment - High resolution mode to determine the number of closely spaced targets.

Clutter/Propagation Map - Determine the current clutter and electromagnetic propagation conditions, to improve future allocation decisions.

Calibration - Low priority tasks which are performed to ensure the radar is correctly calibrated.

2.3.2.4 Weapons Support

It may also be necessary for the radar to provide support to missiles in the form of data uplink and midcourse or terminal guidance. These tasks tend to be highly synchronous with high priorities, as untimely scheduling severely reduces the capability of the weapons system.

The radar resource manager (RRM) must be able to juggle the requirements of these differing functions to maximise the performance of the system given the finite resource and with respect to the mission objectives.

2.3.3 Operational Systems

Operational multifunction radar systems are emerging in the ground, airborne and maritime domains. In this section MESAR is given as an example of a multifunction radar system.

2.3.3.1 MESAR Programme

[Stafford, 2007] provides an overview of the UK Ministry of Defence Multi-Function Electronically Scanned Adaptive Radar (MESAR) programme. The aim of the programme was to produce a naval active array radar for surveillance, fire control and ballistic missile defence. The programme ran in excess of twenty years from 1982 and involved the MESAR1 prototype and MESAR2 pre-production prototype.

The MESAR1 system was created to provide a testbed for developing key areas of the radar functionality. Specifically, it was required to develop digital adaptive beamforming, wide frequency agile bandwidth, digital waveform generation and pulse compression. It was desired to generate a single surveillance beam, or multiple simultaneous beams which would support variable surveillance update rates by sector and adaptive tracking control. It was also required to develop the real time software to control the system.

As a initial prototype the MESAR1 system used an octagonal thinned array for 918 elements randomly populated with 156 transmitter receiver modules. The peak power of each module was 2W which produced a mean output power of less than 100W and an instrumental range of 55km. It had an agile frequency range between 2.7GHz-3.3GHz and a duty cycle of 30%. The MESAR2, shown in Fig. 2.8(a) programme utilised a new antenna with 1264 elements and a module peak power of 10W allow an instrumental range of 400km.

(a) MESAR   (b) SAMPSON

Figure 2.8: MESAR2 and SAMPSON multifunction radars [BAE Systems Insyte].

The success of the programme in developing and demonstrate key techniques led to the production of the SAMPSON radar, shown in Fig. 2.8(b) which is going into operation on the Royal Navy Type 45 Destroyers. In this role SAMPSON is part of the principal anti-air missile system providing weapons support as well as complete long range air picture.

2.4   Summary

Multifunction radars are capable of supporting differing functions by utilising an agile beam and configuring the radar operation for each task. This chapter has presented basic radar theory and background relevant to multifunction operation. The multifunction radar system described in this chapter is controlled by the automated resource manager and hence overall performance is dictated by the resource managers' ability to adapt performance to a dynamic and uncertain environment. This resource management control problem is the subject of the research in this thesis and existing methods to the problem are discussed in the following chapter.

Chapter 3

Multifunction Radar Resource Management

Multifunction radar systems can dynamically adapt performance given changing mission objectives and an uncertain environment. As performance is critically limited by the effectiveness of the automated radar resource manager, the management of the sensor has received widespread attention in the literature. The pertinence of sensor management [Musick and Malhotra, 1994; Ng and Ng, 2000; Holloway, 2001] is frequently stressed and the need for a closed loop system [Finch, 1998] is often identified. It is also widely recognised that adaptive radar control [Pollard, 1990; Powis et al., 1992] is required which enables processed received measurements to be combined with *a priori* knowledge [Guerci and Baranoski, 2006] to dictate the future system behaviour.

This chapter provides a critical review of a range of multifunction radar resource management techniques. This includes assessment of the effect of task parameter selection on performance. General methodologies and architectures for resource management are also described, as well as specific methods for performing scheduling and priority assignment. This review concludes by highlighting the gaps in the existing work to which the research in this thesis is targeted.

3.1 Parameter Optimisation

The multifunction radar has a variety of parameters under control as detailed in Sec. 2.3.1. The process of optimising task parameter selection requires knowledge of how task parameter selection affects performance. This section details the conclusions of such studies for the surveillance and tracking functions.

3.1.1 Surveillance

In contrast to a mechanically scanned surveillance radar where beam position energy and sampling rate are fixed, an MFR utilising an ESA can adapt the energy management and sampling rate across the surveillance region. Specifically, operational parameters which characterise surveillance performance are the beam pattern, revisit interval, beam spacing, energy which is proportional to dwell length, and detection threshold. Relevant overviews of the interplay between these parameters are given in Billetter [1989]; Sabatini and Tarantino [1994].

Beam agility enables sequential detection techniques in surveillance, such as alert-confirm [Dana and Moraitis, 1981; Trunk et al., 1995]. Alert-confirm implements a two stage detection policy where a lowered detection threshold acts as a first detection stage to produces alerts, which is followed by a secondary confirmation dwell. The time interval between alert and confirm stages is kept short to ensure a highly correlated radar cross section. When using alert-confirm the total time for the search dwell can be estimated as:

$$T_S = \tau_A + P_{FA} N_B \tau_c \tag{3.1}$$

where τ_A and τ_C are the alert and confirm times respectively and N_B is the number of detection bins. Analysis in Dana and Moraitis [1981] indicates that a correlated confirmation dwell has a 5-6dB improvement in SNR compared to the equivalent non-cued dwell. It is worth noting the detection improvement associated with sequential detection is only beneficial when an independent confirm dwell is achievable, such as in thermal noise. Correlated false returns which are encountered in numerous clutter environments reduce the effectiveness of this method.
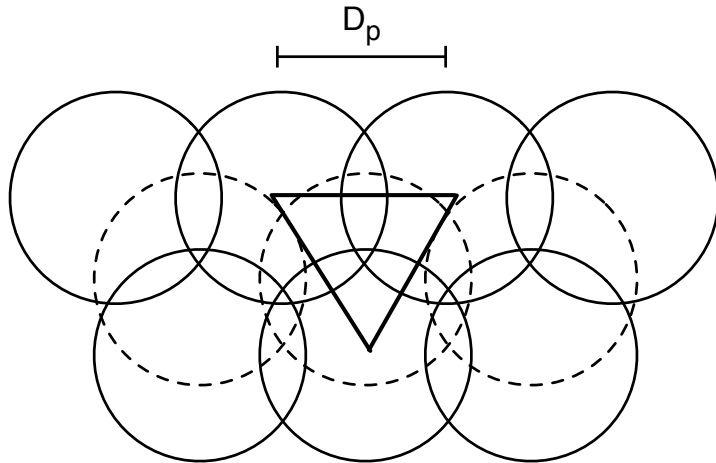
Figure 3.1: Interleaved search beam pattern

Beam shaping loss, which occurs due to the target being offset from the centre of the beam, can cause nulls in the detection probability across the surveillance region. It is common to use a triangular search pattern which is shown by the solid lines in Fig. 3.1. This search pattern offsets the beam centres on adjacent search bars so that the beam centres form a triangle and the potential nulls in the surveillance region are reduced. To reduce the severity of the potential null, the beam pattern can be interlaced [Billam, 1997], so that the beam centres on the next scan are directed at the previous nulls. Using this method the next beam positions are shown in Fig. 3.1 by the dashed line. The triangular search pattern is parameterised by the beam spacing D_p which separates each beam by angle θ_S, which is taken in terms

of the beamwidth:

$$\theta_S = Dp\theta_B \tag{3.2}$$

A wider beam spacing requires fewer beams to search a fixed space, which reduces the time for the search. However, increased energy per beam would be required to offset the loss of gain through beam shaping losses. Analysis by Blackman and Popoli [1999] and Billam [1992] indicates a suitable beam spacing around $D_p = 0.75$ with general insensitivity to the selection of D_p in the region between 0.6 and 1.0.

The dwell energy and search revisit interval, which are inversely related, are the fundamental parameters which are adapted for surveillance in MFR systems [Billam, 1997, 1992]. Increasing the energy of the dwell, which is proportional to the dwell length, increases target detection probability and the single look detection range. However, given a finite search frame time this also increases the revisit interval which increases the target closure between scans and so reduces the cumulative detection range. Given this trade the choice of dwell energy and revisit interval are decided based on the objective of the search function. For example, a self protect search requires a rapid revisit and shorter dwell to detect close proximity pop up targets, whereas a long range search requires a longer dwell and longer revisit interval to detect at greater ranges.

The ESA allows for multiple simultaneous beams to be steered in the surveillance region of interest. By using a broad fan beam on transmit and a cluster of narrow pencil beams on receive, higher angular resolution and hence accuracy can be achieved [Wirth, 2001]. Additionally, multiple beamforming on receive allows for a faster search by enabling a shorter revisit interval. This results from the SNR being increased, in comparison to using a broad fan beam on receive, by a factor of the number of simultaneous beams used as each of the narrow pencil beams has a higher gain. However, the improvement in performance achieved through multiple beamforming is offset by the requirement for multiple receive channels.

At present surveillance is performed according to pre-defined parameters for differing sectors, with little scenario or environmental relevance. As such there is a requirement for resource management techniques which can demonstrate intelligent adaptation to a dynamic and uncertain scenario.

3.1.2 Tracking

In comparison to surveillance there has been considerably more work addressing resource management methods for tracking. Key strands of this work provide methods for adaptively selecting the task revisit interval to maximising the number of targets in track and optimising waveform selection to improve tracking performance. Additionally, the benchmark tests successfully contributed a platform to assess and compare differing tracking and resource management techniques.

3.1.2.1   Adaptive Tracking

Several studies have analysed the trades associated with adaptive update rate selection, with similar conclusions. A standard approach [van Keuk and Blackman, 1993] is to select a revisit interval based on the earliest time after the filter angular prediction error, along the major axis of the uncertainty ellipse G, exceeds a fraction of the beamwidth as shown in Fig. 3.2. The fraction is called the track sharpness and denoted v_0. So, the next revisit time t_{K+1} is chosen according to:

$$G(t_{K+1}|K) = v_0.\theta_B \tag{3.3}$$

Missed detections resulting from a non-unity probability of detection are followed by a revisit scheduled at the minimum revisit time. Choosing the maximum revisit interval that bounds the target uncertainty balances the trade between minimising resource consumption through long revisit intervals whilst minimising looks per update which is a consequence of beam position loss and target uncertainty spread. Assuming a Singer target dynamic model [Singer, 1970], an expression relating the revisit interval and the track prediction error variance is presented, which is used to provide an estimate of the track loading. The analysis indicates that the minimum energy allocation is independent of target range and manoeuvre and can be found through choice of v_0, P_{FA} and SNR_0. The minimum track loading is desirable as it is complementary to the system wide objective of maximising the number of targets in track. This minimum energy only considers target dynamics, and does not consider data association uncertainty or situation assessment which may necessitate parameters to be selected contrary to the suggested track sharpness suggested. The Van Keuk model is discussed at length in Sec. 4.1.2.1.
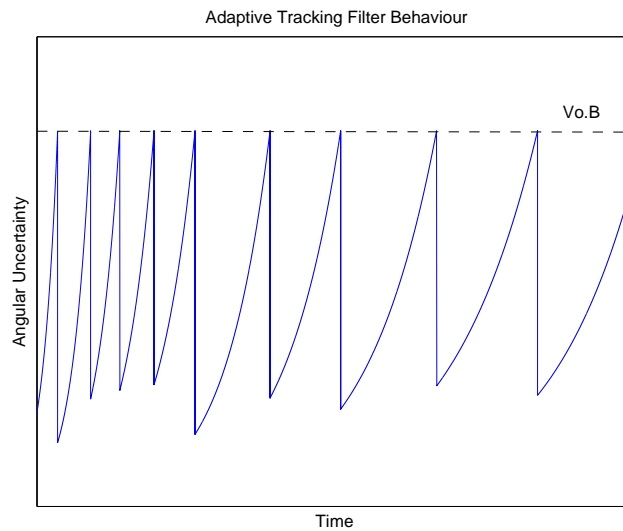
Figure 3.2: Track sharpness adaptive revisit strategy

Gilson [1990] also investigates the power requirement for tracking by comparing differing tracking models for a track-while-scan and fire control radars. It is found that the power requirement for tracking decreases monotonically with revisit interval for a track-while-scan radar whereas a minimum power requirement exists for fire control radar, around $\frac{1}{4}$ of the beamwidth which is similar to the findings of van Keuk and Blackman [1993]. The study also shows that the power requirement for tracking is relatively insensitive to tracking model.

3.1.2.2 Waveform Agile Tracking

In addition to controlling the time-energy budget of the radar it is possible to dynamically adapt the transmitted waveform. Sequential state estimation, which is inherently closed loop, can provide the basis for assessing the current effectiveness of each potential transmit waveform. As different waveforms have different resolution properties, adapting the waveform can reduce the target tracking error and improve target detection by dynamically providing high resolution in the necessary dimension.

The first efforts on intra-pulse waveform agile tracking [Kershaw and Evans, 1994] analysed the effect of waveform agility on tracking performance with one dimensional target motion, unity probability of detection and no clutter. The linearity of the problem permitted the application of a Kalman filter which could be updated with a variety of potential linear FM chirp waveforms. The work successfully produced closed form solutions for the waveform selection which minimised the tracking MSE or the tracking validation gate volume. This work was extended [Kershaw and Evans, 1997] to include the effect of non-unity probability of detection and clutter.

Mutual information, which was first applied to waveform design [Bell, 1993], has also been applied to waveform selection. The mutual information $I(X_k; Z_k)$ between the target state X_k and the waveform dependent measurement Z_k at time k quantifies the reduction in uncertainty in the target state through the measurement. As such, the maximisation of mutual information has successfully been applied as the criterion for waveform selection from fixed libraries [Suvorova et al., 2006; Cochran et al., 2009].

This strand of work on waveform agile tracking has produced interesting conclusions summarised in the review by Sira et al. [2009]. It is commented that intra-pulse waveform modulations which maximise time bandwidth are not necessarily best for tracking and that dynamic waveform selection reduces tracking error, most noticeably in cluttered environments. It is also commented that the use of non-linear chirp waveforms offer significant improvements over linear frequency modulated chirp waveforms. Despite the poor ambiguity properties of non-linear chirp waveforms, the ability to control the nature of the ambiguity and hence choose where resolution is applied improves the tracking performance.

Waveform agile tracking is of significant relevance to the general area of sensor management as it improves sensing efficiency. However, it is equally applicable to non-multifunction systems as it is to multifunction systems. This thesis does not concentrate on waveform agile tracking, instead the focus is

on the allocation of the finite resource between numerous competing tasks.

3.1.2.3 Benchmarks

The benchmark simulations [Blair et al., 1994, 1995, 1998] provided a comparative testbed to assess track and resource allocation performance against manoeuvring targets. The first benchmark [Blair et al., 1994] studied the efficiency of tracking and allocation methods given a fixed SNR, no false alarms and a single target under six different manoeuvre scenarios. Results from the benchmark tests indicated that a Kalman filter, due to variable gain, enabled an increase with a fixed revisit interval over an $\alpha - \beta$ filter [Rhatigan et al., 1994] from 0.85 to 1.0s. However, the key result highlighted the effectiveness of IMM [Daeipour et al., 1994] which increased the revisit interval to 1.3s and 1.5s with a fixed revisit and two and three models respectively, and up to 2.3s with the adaptive revisit strategy described in Sec. 3.1.2.1.

The second benchmark [Blair et al., 1995, 1998; Kirubarajan et al., 1998] extended the problem to include different radar cross sections and the presence of electronic counter measures (ECM) in the form of stand off jamming (SOJ) and range gate pull off (RGPO). Additional flexibility in resource allocation was allowed through the selection of eight different waveforms with varied SNRs and detection thresholds which produced differing false alarms probabilities. Whilst confirming IMM as the best filtering method, it was found that sophisticated data association techniques were required to combat the ECM. Multiple Hypothesis Tracking (MHT) emerged as the best performer, indicating a combined IMM/MHT system to be favourable. It was also found that adaptive tracking in Electronic Counter Measures (ECM) requires more conservative parameter selections than suggested in van Keuk and Blackman [1993] to prevent unacceptably high track loss.

The benchmarks were very successful at comparing the performance of filtering and data association methods. They also demonstrated that when the resource allocation is coupled to the task function, as in adaptive tracking, the performance of the task function can have profound effects on the subsequent resource allocation. For example, if resources are allocated based on the track state covariance which is poorly estimated by the tracker, extra resources are required to compensate and ensure track maintenance. The performance of the adaptive tracking strategy in 3.1.2.1 was consolidated as it was used by all successful methods in the benchmark tests.

3.2   Resource Management Architectures

This section describes the architectures of radar resource managers which are typically made up from combinations of modules providing specific functionality. General sensor management architectures are described by Musick and Malhotra [1994] and Blackman and Popoli [1999] where the emphasis is placed on combining heterogeneous and non-collocated sensors. These general sensor manager architectures are relevant to a radar architecture and it is recognised that the architecture can be centralised, decentralised

or hierarchical. It is also recognised that the architecture will typically require differing levels which partition code cycle times.

Radar resource mangement architectures vary in the literature but generally contain priority assignment, task managers and scheduler modules. It is accepted that a closed 'macro' loop is created which encompasses the resource manager, transmission, the environment and reception, with the potential for additional micro loops within the resource manager. Memory is located within the different modules, capable of storing fixed knowledge or a temporary memory of the current environment or scenario [Haykin, 2006].

A good radar resource management architecture is given by [Miranda et al., 2006] and presented in Fig. 3.3. In this architecture an environmental model is used to generate requests for radar task functions which utilise waveforms from a database. The requests are assigned a priority and formed into a timeline for transmission by the scheduler. The received measurements are used to update the task functions and the environmental model which closes the loop.
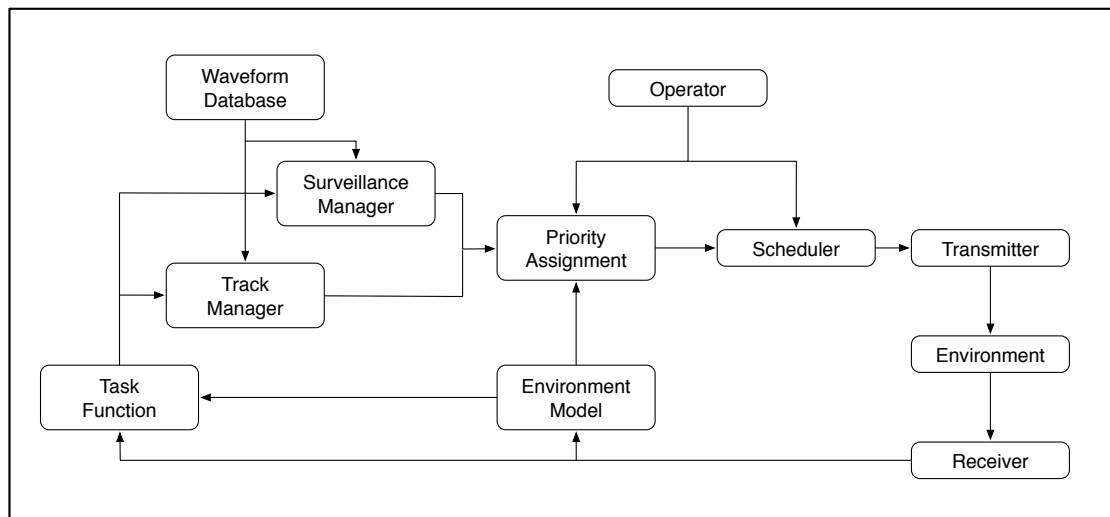
Figure 3.3: Typical multifunction radar resource management architecture [Miranda et al., 2006]

The architecture of the Multi-Function Electronically Scanned Adaptive Radar (MESAR) resource manager [Stafford, 1990] provides insight into an operational RRM. The architecture contains a radar job table, which is a list of jobs prioritised by radar function, a job controller to maintain the job table and waveform selection and scheduler modules. This architecture is able to dynamically schedule tasks and select waveforms with respect to prioritisation in a robust and computationally efficient way.

In the architecture of M3R [Barbaresco et al., 2009], which is also an operational system, the resource management is handled in sequence by the task manager which passes tasks to the dwell manger, which passes dwells to the burst manager, which passes burst to the space time manager for transmission.

The main functionality is provided by the radar task manager and dwell managers which access mission data and a parameter server storing current environment and scenario conditions. Although not explicit it can be assumed that received measurements are processed to update the parameter server and form a closed loop. It can also be assumed that the burst and space time managers become increasingly deterministic due to proximity to the radar front end where the time constraints become increasingly critical. This enables code cycle times to be partitioned which is crucial for efficient online computation.

3.3   Methodologies for Resource Management

This section describes multifunction radar resource management methodologies which broadly fit into the two categories of heuristic or optimisation based. The boundary between resource management methodologies and scheduling is somewhat blurred and often overlap, however, techniques included in this section involve some higher level decision making on resource utilisation.

3.3.1   Rules and Heuristics

Rules and heuristics which guide the resource allocation process are widely implemented in operational systems due to quantifiable task performance under specified conditions and low computational burden. However, they generally suffer from poor and unpredictable performance. Typically, rules are generated which aim to optimise the parameter selection of individual tasks, according to the studies detailed in Sec. 3.1. For example in tracking, it is common to select the revisit interval such that the angular prediction error is maintained beneath a fraction of the beamwidth and a dwell length to maintain a desired SNR [Kirubarajan et al., 1998].

Noyes [1998] provides a description of the rules used to determine track update times for MESAR. Echoing the studies on adaptive tracking, the need to balance short revisit times to ensure the target is close to the predicted position and long revisit times to minimise radar usage is identified. The desired execution time is found as a track accuracy threshold on the updated state covariance and the latest execution time as a function of the predicted state covariance, which are passed to the scheduler. A requirement on the track accuracy is used for the desired revisit interval which reflects the application domain of the MESAR system, whereby the track may need to meet accuracy constraints to cue a weapons system.

Although rules used for tracking control are predominantly based on track accuracy, there are cases when resources should be allocated based on other criteria. For example, Davidson [2007] describes the allocation of resource to aid rapid release in track initiation and Whitewood et al. [2007] details potential improvement for crossing tracks. These alternatives are a consequence of track accuracy being of less importance than track purity and maintenance for surveillance systems. Allocation based on differing criteria has not been widely recognised, except in the recent publication [Song and Musicki, 2010]. A single mechanism which can allocate resource based on multiple differing criteria such as accuracy and

track existence can improve radar functionality and is an aim of this thesis.

In the airborne domain, Bier et al. [1988] describes rules for sensor load management as well as discussion on architecture. In Gillespie et al. [2005] track updates are requested as determined by rules in the tracker and have priority over the search dwells. Heuristics are used to guide the search behaviour given a variable tracking load which produces a dynamic, emergent search behaviour instead of a pre-defined pattern. In overload the searching tasks on the radar boresight are preferred, where there is maximum gain and maximum closing target velocity, while search on the extremes of the radar are severely degraded or dropped.

Vaughan [2001] defines empirical rules of thumb which provides a pragmatic system engineering approach for MFR surveillance control. Five levels of rules are defined for gathering command and environment information, prioritisation of sectors, allocation of time budget, generation of beam management strategy and generation of waveform and signal processing strategy. The crude nature of these control rules is defended by arguing that performance is relatively insensitive to selections close to the optimum. Although this is true to an extent, it is clear that control that is this coarse can be significantly improved upon.

These previous studies use rules and heuristics to optimise the individual task parameters without consideration of the system wide objectives and constraints. As such the locally optimum parameters represent a single desirable point in quality space without the context of the finite resource which places additional responsibility on the scheduler to mediate the access. The scheduler, which has a rapid code cycle time, is only capable of making deterministic decisions often producing poor and unpredictable performance in overload which leads to non-graceful degradation. The system wide resource constraints are tackled in the M3R system [Barbaresco et al., 2009] by providing local and global radar load handling. A radar dwell set is tested for schedulability by considering the summation of the individual dwell loadings l_d, or normalised dwell durations $l_d = \tau_d/t_f$, where τ_d is the coherent dwell duration and t_f is the revisit interval. Resources are mainly balanced between search and track, with TWS able to take on some of the tracking load. Strategies based on time constraint relaxation are used to enable graceful degradation. Additionally, the system functionality is provided by a set of rules which define dynamic search allocation, strategies for robust dynamic tracking allocation and adaptation to the environment through waveform selection.

Rules and heuristics are computationally efficient methods of guiding the resource allocation process. However, individual parameter selection without respect of the system wide resource constraint can provide non-graceful degradation. As such current rules and heuristics produce sub-optimal performance. Despite this, rules and heuristics are widely applied in operational systems.

3.3.2 Optimisation

Optimisation methods aim to minimise or maximise an objective function over some time horizon. The choice of objective function, which is a cost function to minimise or a utility function to maximise, has significant impact. Optimisation methods can potentially produce optimal solutions for a given objective function but are severely hindered by the curse of dimensionality. Sensor management optimisation was first presented by Nash [1977] who used linear programming to determine sensor to track assignments using a cost function which combined both target priorities and track accuracy. Optimisation methods have only seen significant advances recently as computer processing power has increased.

3.3.2.1 Markov Decision Processes

Sensor management is frequently approached as a stochastic control problem where a multistage objective function is optimised using dynamic programming [Washburn et al., 2002]. In stochastic control problems sequential decisions are made to perform varied actions which can generate varied observations. An optimal decision, whose outcome is uncertain, is sought over the time horizon of future stages, given information from previous observations. A Markov Decision Process (MDP) is a type of stochastic control problem where observations provide complete information on the true state of the underlying dynamic system, which is modelled as a Markov process. However, Partially Observable Markov Decision Processes (POMDP) are of more relevance to sensor management, where observations provide only incomplete information on the true state of the underlying dynamic system. In this case the relationship between the observed quantities and the underlying state is modelled statistically as the measurements are acquired.

In a POMDP [Hero et al., 2007] there exists a finite set of possible states X_k and possible actions A_k at each stage k. The pairing of an action with a state produces a single stage reward according to the reward function $\tilde{R}(x, a)$. Decisions to take actions are based on information collected over previous decision stages $I_k = \{x_0, a_0, ..., x_{k-1}, a_{k-1}, x_k\}$. A policy $\hat{\gamma}(I_k)$ provides a mapping from information to an action $\hat{\gamma} : X_k \rightarrow A_k$ depending on the most recent state x_k. The policy dictates a trajectory of actions which produce a total reward summed from each sequential action:

$$\tilde{R} \equiv \tilde{R}_N(x_N) + \sum_{k=0}^{N-1} \tilde{R}_k(x_k, a_k) \tag{3.4}$$

which is demonstrated in Fig. 3.4 where each circles represents a possible state, each dashed line represents a possible action and each solid line represents the action taken. The reward for the whole trajectory is the sum of the individual rewards marked in the figure.

The objective of the POMDP problem is to determine the policy which maximises the total reward. Given a finite state problem with finite stages it is possible to represent a POMDP as an equivalent

MDP which can be tackled using dynamic programming. Dynamic programming is a consequence of Bellman's Principle of Optimality which states that given any starting point on a complete optimal trajectory, the remainder of the complete optimal trajectory is also optimal for the problem starting from that point. This principle enables the optimisation of the complete problem to be decomposed into the choice of optimal actions for each stage. The optimal action at stage k is determined by the Q-value:

$$\tilde{Q}_{H-k}(\pi_k, a) = r(\pi_k, a) + E\left[V^*_{H-k-1}(\pi_{k+1})|\pi_k\right] \tag{3.5}$$

which combines the reward of the current stage, $r(\pi_k, a)$, and the expected reward from future stages given the optimal objective function over future stages V^*_{H-k-1} up to time horizon H, where π_k is the belief state at time k. Exact calculation of the Q-value, called the 'lookahead', is typically intractable and requires approximation.
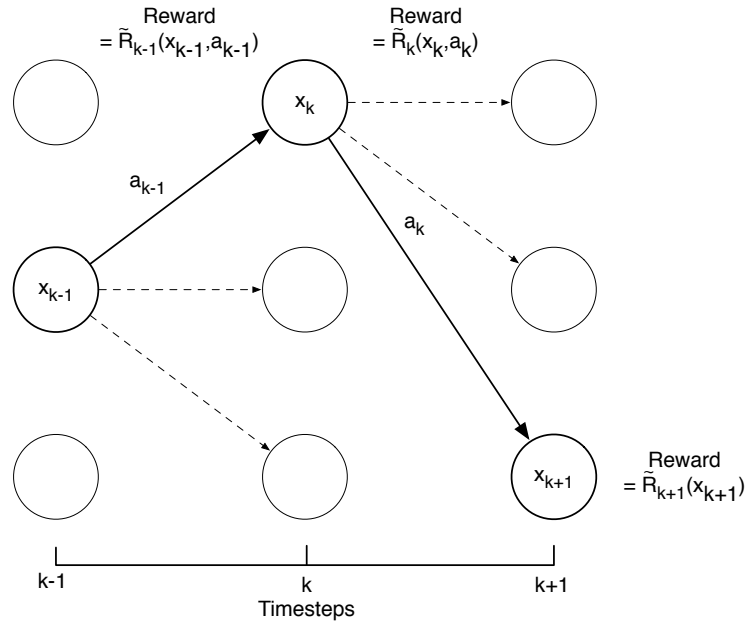
Figure 3.4: Markov decision problem.

The representation of a POMDP as a MDP to by solved be dynamic programming has been applied to classify multiple unknown objects using multi-mode sensor resource [Castanon, 1997]. The combinatorial nature of potential belief states rapidly renders the problem intractable due to difficulty in calculating the Q-value with increasing time horizon. As such efficient methods of approximating the Q-value are required. This is addressed by Castanon [1997] where the action paths are replaced with average resource utilisations which enables the production of near optimal allocations. Various other methods of approximating the Q-value have been proposed Hero et al. [2007], including policy rollout [He and Chong, 2004, 2006] where a base policy is assessed by Monte Carlo simulation. In a different

approach, Blatt and Hero [2006] attempt to solve the POMDP using reinforcement learning. Despite techniques for approximating the Q-value, POMDPs are hindered by the curse of dimensionality and so have not been widely applied on operational systems where fast reaction times are a crucial requirement.

3.3.2.2 Information Theoretic Optimisation

Information theoretic sensor management aims to optimise the information production of the sensor by replacing the optimisation objective function with an information theoretic measure. Different information measures have been proposed for differing sensor management problems. Hintz [1991] and Hintz and McVey [1991] were first to examine the expected change in Shannon entropy with a Kalman filter tracking a target in one dimension. The discrimination gain or Kullback-Leibler divergence has been suggested by Schmaedeke and Kastella [1998] for sensor to target tasking and in Kastella [1997] to optimise detection and classification. The Kullback-Leibler divergence has also been suggested for tracking control [Kreucher et al., 2004, 2005c] combined with the joint multi-target probability density (JMPD)[Kreucher et al., 2005b]. The way in which information measures can be used to estimate the Q-value lookahead in a POMDP is given in Kreucher and Hero [2006].

Kreucher et al. [2005a] present a comparison of task driven and information driven management where it is found that task driven management performs the best for a given task, however, information driven management performs best when multiple competing performance criteria are present. It is therefore suggested that information can provide a 'universal proxy' to represent differing tasks. This is an especially relevant assertion, as such, investigating the role of information in multifunction radar resource management is an aim of this thesis.

3.3.2.3 Q-RAM

Q-RAM (Quality of Service (QoS) Resource Allocation Method) [Ghosh et al., 2003, 2004; J.P. Hansen and Lehoczky, 2004; Hansen et al., 2006] provides a 'quality of service' optimisation method which aims to select parameters to produce a set of best quality tasks given the resource constraint. To this end, Q-RAM models the nature of a dwell and parameter dimensions which define the problem. Each radar search or track dwell is modelled as a transmit power, transmission time, idle interval and reception period. The QoS model is characterised as QoS dimensions, environmental dimensions, operational dimensions and resource dimensions where the QoS dimensions are aspects of task quality, such as position accuracy. Environmental dimensions are aspects which affect performance but are outside of control such as target range and manoeuvrability and operational dimensions are aspects which affect performance but are under control such as the task revisit rate. The resource dimensions are the finite resource to be distributed between tasks which is radar time or loading. A utility function is defined which quantifies the satisfaction associated with each point in the quality space. This utility model is demonstrated in Fig. 3.5 where it can be seen that each operational parameter uses a different resource

loading and produces a different task utility. Parameters along the "concave majorant" [Hansen et al., 2006], where utility per resource is maximised, are preferred.

The goal of the optimisation is to choose operating points for each task which maximise the global utility production given constraints on radar energy and time utilisation. The high configurability of each task combined with the large number of tasks creates a large number of potential set points or operating points for the problem. As such, fast traversal methods are used which exploit the monotonic nature of each parameter dimension to reduce the number of setpoints considered to points on the concave majorant.

The QoS model and differential utility function as an objective function is a very useful contribution of this work as has a strong theoretic justification through the Karush-Kuhn-Tucker conditions. However, despite the fast traversal techniques, the quoted time to compute the allocation for 300 targets in 1 sec is too long for feasible application in an multifunction radar which requires rapid reaction to pop-up targets. This is in part due to unoccupied regions of the concave majorant being unnecessarily and repetitively computed. A continuous mechanism which adjusts the current allocation instead of frequently recomputing the entire allocation could potentially reduce the computation time to be feasible for an operation system. The Q-RAM approach has been extended to include the allocation of multiple resources [Irci et al., 2006].
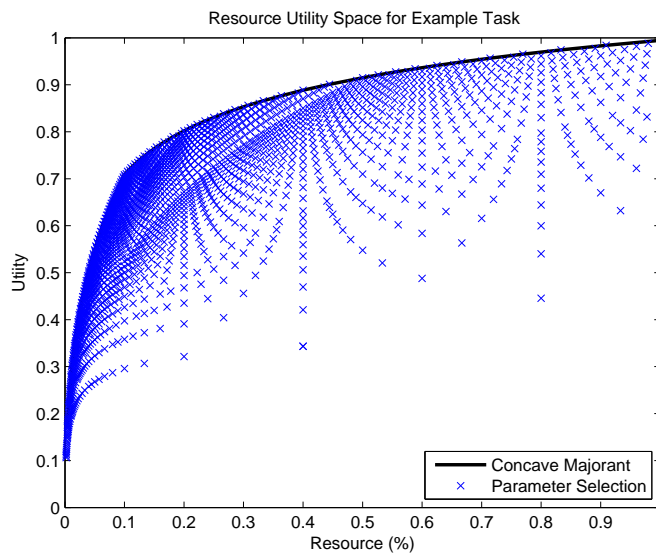
Figure 3.5: Resource utility space for an example task

3.3.2.4   Other Methods

Stromberg and Grahn [1996] describes a minimisation problem solved by dynamic programming which also encompasses scheduling. The problem is broken into the addition of a task to a set of already scheduled tasks and aims to minimise the total scheduled time by using the highest PRF value that

satisfies range and Doppler unambiguity requirement. This approach is described as optimal, however, it is only optimal in the sense of the objective function which is the shortest schedule time. In reality there are more relevant constraints relating to task quality than just the shortest schedule time.

Tracking parameter control has also been considered as a constrained optimisation problem [Zwaga et al., 2003; Zwaga and Driessen, 2005; Boers et al., 2006]. It is identified that the conventional method of tracking parameter control is to select a revisit interval as a fraction of the half beamwidth and a dwell length to maintain a required SNR. It is argued that this does not directly address the minimisation of radar usage for accurate tracking as the separation means there is no dynamic trade off between the dwell length and revisit interval and considering a horizon of one step ahead only provides a locally optimised solution. This is addressed by formulating the problem as a constrained minimisation problem to minimise the tracking task loading subject to a constraint on the updated state covariance, which includes both the revisit interval and the dwell length. This minimisation of a non-linear function subject to a non-linear inequality has no analytic solution, however, numerical solutions are provided which are evaluated by the optimisation toolbox in MATLAB. It is found that using the updated state covariance means an SNR requirement is no longer needed and there is an improvement in radar loading. However, it is recognised that computation is slow and only relevant for offline analysis to compare with alternative online techniques.

3.3.3 Discussion

The translation of task constraints into the time domain through temporal reasoning methods presented in Stromberg [1996]; Stromberg and Grahn [1996] is a somewhat simple and obvious yet widely relevant and applied technique. Translation into the time domain enables computationally efficient control, as the passing of time is the same for all tasks and so it is simple to directly compare differing task specific constraints. This is especially relevant, but not acknowledged in the work, for rotating systems where the limited field of view creates a sequence of schedulable windows.

It is possible to identify characteristics of the two general approaches to radar resource management through the literature. Rules and heuristics are simple to apply, computationally efficient and provide quantifiable performance against specified conditions. Because of these characteristics they have been favoured for application to operational systems. However, individual rules are not able to address the system wide objectives and despite quantifiable performance in specified conditions they can produce unpredictable and poor performance in unspecified conditions. In contrast, optimisation approaches produce optimal or near optimal allocations. However, these methods have a high computational cost and so have not been applied on operational systems. These points are echoed in the comparison between dynamic programming and temporal reasoning by Stromberg and Grahn [1996].

These works indicate that there is a need for resource management mechanisms which are computa-

tionally light for application to operational systems whilst providing near optimal allocations especially when overloaded.

3.4 Scheduling

The scheduler [French, 1981] is responsible for forming the multiple requests from differing multifunction radar tasks into a transmittable timeline. This role requires it to resolve potential conflicts arising from the finite nature of the resource. Scheduling approaches loosely fit into two categories, local optimum or best first, a good overview of these differing types can be found in Blackman and Popoli [1999]. There is overlap between schedulers and the resource management methodologies of Sec. 3.3, the scheduler is differentiated in this thesis as a module which turns requests into a timeline without higher level decision making capability.

3.4.1 Local Optimum

Local optimum or brick packing methods attempt to form the time line by creating a series of allocation frames of fixed duration. Whilst the previous allocation frame is being executed, the next frame is being calculated. This is represented in Fig. 3.6 for the set of schedulable tasks T_A. Given a measure of optimality, an exhaustive search can provide the optimum solution over the time horizon, however, heuristics [Winter and Lupinski, 2006; Winter and Baptiste, 2007] are used to guide the 'packing' of the tasks into the frame. A result of this method is that pop-up tasks which require immediate execution are required to wait up to the duration of the frame which can seriously degrade the reaction time of the radar.
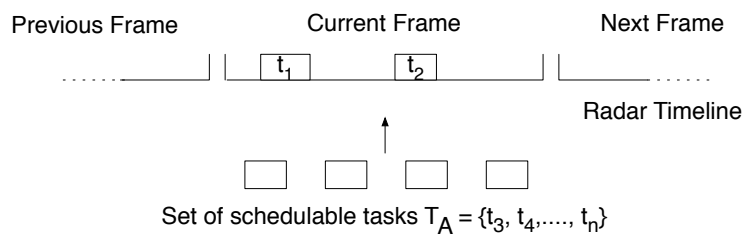
Figure 3.6: Local optimum/brick packing scheduler

A local optimum method is presented by Orman et al. [1996] where five heuristics which guide the task placement in the frame are compared. A job is defined which can perform any function and contain a transmission period, idle interval and a reception period. A conflict of optimality between scheduling delay and utilisation is identified, whereby tracking tasks require scheduling as close to the desired time as possible, however, this can reduce the radar utilisation. As such, both must be considered in assessing the performance of the scheduler. The five heuristics differ in the degree to which tracking tasks can be executed off their desired execution time and whether the tasks can be interleaved by scheduling other

tasks within the idle interval for a task. It is found that adjusting the execution time of the tracking tasks significantly increases the radar utilisation without a drastic effect on tracking performance. The resulting scheduler is therefore suitable for forming a timeline comprised of surveillance and tracking tasks, however, the main drawback is the potentially slow reaction time resulting from the allocation frame.

Izquierdo-Fuente and Casar-Corredera [1994a] details results on a local optimum scheduler which allows tasks to be interleaved. Interleaving of tasks enables the 'dead' time between transmission and reception to be utilised for other tasks on differing carrier frequencies. Although increasing radar utilisation this creates an significant additional scheduling challenge which is tackled in Izquierdo-Fuente and Casar-Corredera [1994b] using a neural network. In reality the extent to which tasks can be interleaved may be limited by the allowable duty factor of the radar antenna, but also by the availability of multiple oscillators which allows phase coherence to be maintained for both tasks simultaneously.

3.4.2 Best First

Best first schedulers sequentially execute the next best task from a set of requests, which are ordered into queues according to some criterion. Example criteria are earliest deadline first (EDF) or highest priority first (HPF), or a combination of both as demonstrated in Fig. 3.7. As the queue can be maintained with low computational burden, best first schedulers are computationally efficient. Also, as tasks are scheduled from the queues sequentially, the full radar time is utilised. However, there is still some delay in scheduling and it may be required to send the schedule to the antenna in an allocation frame to allow the array control to be applied. The characteristic of best first schedulers is that this frame is very short in comparison to local optimum methods.
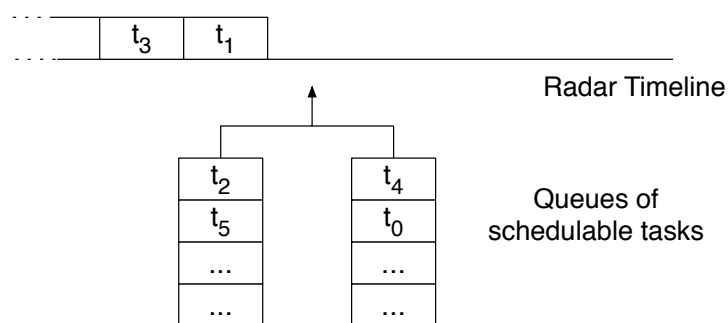
Figure 3.7: Queue/best first scheduler

Huizing and Bloemen [1996] presents a best first scheduling method where differing task requests are ordered into 'branches' according to desired execution time and priority. Additionally, two queues are maintained for normal dwells and terminal guidance dwells. Requests are removed from the branches, according to the highest priority and the earliest deadline and placed in the queues if the current time

is within the dwells transmission window, and if the length of the queue will not exceed the maximum queue length. Finally, tasks are removed from the top of the queues to form the radar timeline. Through simulation of the queuing mechanism the desirable characteristics of modest computation and reaction to pop up tasks is demonstrated. Despite respect paid to priority, requests which expire due to overloading are dropped which leads to ungraceful degradation. It would be preferable to delay tasks or adjust the requests according to the current loading.

Butler [1998] describes a time-balance for radar scheduling which extends that described in Stafford [1990]. In the original algorithm the time-balance represents the amount of time which is owed to each task. However, after finding that this can be implemented using a conceptually simpler approach, the time-balance is changed to represent the earliness or lateness of each task. The desired execution time of the task is determined by a task specific rule. The next task to schedule is chosen as the highest priority task with the largest lateness. From the description provided, scheduling based on time balance appears equivalent to a highest priority-earliest deadline first scheduler. Butler et al. [1997] also investigates scheduling using a rotating phased array, where is found that rotating arrays can offer performance benefits over non-rotating arrays.

Barbato and Giustiniani [1992] presents a simple queue based algorithm which accommodates the variable tracking load by reducing volume search update time. In Stoffel [1994] a highest priority scheduler is compared to a heuristic search scheduler which searches potential non-myopic sequences guided by a heuristic function with costs assigned by fuzzy logic. Under normal loading conditions the two approaches performed similarly, however, in overload the heuristic search was found to perform better. This is attributed to the heuristic search maintaining low priority tasks that improve surveillance and track maintenance which would otherwise be dropped in the highest priority case.

3.4.3 Discussion

From the literature it is possible to identify key aspects for a scheduler:

- Create the smallest possible deviation from requested parameters.

- Deterministic as possible operation for manageable computation.

- Respect task priority.

- Provide graceful degradation in overload.

- Allow rapid reaction to pop-up tasks.

The comparison of the Orman and Butler schedulers given by Miranda et al. [2007a] provides insight on the difference between local optimum and best first schedulers. Through simulation it is demonstrated that the Orman scheduler allows the tracking tasks to be scheduled close to execution

times but at the cost of lower occupancy. The Butler scheduler maximises the radar occupancy, but the tasks have a greater delay which can affect task quality. However, a general assertion [Wintenby and Krishnamurthy, 2006] is that this delay has only a small affect on task quality and so the desirable characteristic of rapid reaction time makes best first or queue based schedulers preferable.

3.5 Priority Assignment

Priority assignment is an essential aspect of resource management which reflects the fact that different tasks have differing importance. The priority is typically a value which represents the task's entitlement to resource relative to other tasks given the current mission and tactical scenario. When the system is underloaded priority has little effect, however, in overload the priority assignment is crucial for graceful degradation.

3.5.1 Function Ranking

Differing functions maintained by the system inherently have differing importance. As such a simple method of priority assignment is to assign a priority value depending on the importance ranking of the respective function. Surveillance is lowly ranked as it has the lowest sensitivity to scheduling delay in comparison to tracking which requires tasks to be scheduled close to desired time. Track initiation usually takes priority over track update as it will only be successful with several frequent updates. Critical functions, such as plot confirmation which requires a rapid revisit for a correlated radar cross section and track maintenance to prevent track loss have high priorities. Generally weapons control functions have the highest priority as they are very sensitive to scheduling delay, and their successful operation is usually closely aligned with survival. A typical priority table which was used by MESAR is shown in Table 3.1. Many similar ranking tables can be found in the literature [Huizing and Bloemen, 1996; Butler, 1998; Gillespie et al., 2005; Orman et al., 1996; Nelander and Stromberg, 1997; Stafford, 1990; Stoffel, 1994].

Table 3.1: Priority structure for MESAR radar resource manager

Priority	Task	
7	Track Maintenance	(Highest Priority)
6	Plot Confirmation	
5	Track Initiation	
4	Track Update	
3	Surveillance	
2	Slow Track Map/Surface Picture	
1	Receiver Calibration	(Lowest Priority)

Although the assumption that differing functions have different priorities is valid, it is limited in the assignment of different priorities to tasks within a function. Situational assessment or mission requirements may dictate that tasks within specific regions or threat directions are of higher priority.

3.5.2 Fuzzy Logic

Fuzzy logic methods improve upon the fixed priority assignment by allowing assignment over a continuous range depending on the specifics of the task. A detailed review of fuzzy logic as well as situational assessment can be found in Blackman and Popoli [1999]. Fuzzy logic has been applied in Vine [2001] that to schedule tasks based on the highest membership of a 'task ready' fuzzy set.

Miranda [2004]; Miranda et al. [2007b] provides an analysis of a fuzzy logic priority assignment system. Fuzzy values are assigned to variables representing attributes of the surveillance sector or target track. Fuzzy if-then rules are applied to determine the priority of the target track or surveillance sector. The method is validated against test trajectories and the fuzzy approach is compared to a fixed priority assignment and a hard logic approach. The hard logic approach uses the same rules as the fuzzy approach but allows for only one rule to be fired at a time. It was found that the hard logic approach and the fixed priority approach were less computationally demanding than the fuzzy approach. However, the fixed approach allowed for no variations in different target or surveillance sector types and the hard logic approach had priority transitions which tended to jump suddenly between values. The fuzzy logic approach showed smooth transitions allowing greater variations in priority. This was as a result of including all possible information into the priority decision-making process. It is asserted that by improving the quality of the priority assignment the resulting allocation is improved. Although this successfully produces a continuous priority value, the degree to which it can be trusted to provide the correct value is uncertain, which is a serious concern for operational systems.

3.5.3 Discussion

In addition to these methods, rules can be applied which could produce more predictable behaviour than the fuzzy logic, however, specific rules for priority assignment applied in real systems is rarely published. In other methods Popoli and Blackman [1987] details an expert system approach and Komorniczak et al. [2000]; Komorniczak and Pietrasinski [2000] utilise neural networks to enable a learning ability for the priority assignment.

Fuzzy logic methods enable a continuous priority assignment, however, the aspect of trust and stability remains a concern for their operational application. As such accurate and trusted priority assignment which considers all aspects of tactical and situational awareness remains a challenge. It is commonly overlooked that the resource management mechanism must effectively manifest the priority into behaviour. The process of prioritisation is not the focus of this work, however, it is highlighted here because the scheduler or resource allocation must efficiently transform prioritisation into behaviour.

3.6   Summary

Radar resource managers implement the automatic and online control of the multifunction radar system. A typical architecture for a radar resource manager involves a number of modules providing specific functionality. Modules include the environmental model and waveform database which are used by the task management modules to generate task requests which are weighted by the priority assignment module and formed into a timeline by the scheduler module. The modules form a closed loop for adaptive control and are separated to partition code cycle times enabling efficient computation. The operator is removed from this closed loop and takes on a supervisory role.

From this review of radar resource management techniques it is possible to identify areas to target research:

- Intelligent surveillance is required which improves upon current fixed surveillance behaviour to react to a dynamic and uncertain environment.

- A single mechanism is required which can effectively allocate based on multiple requirements, such as track accuracy and track existence.

- The resource allocation mechanism is required to be computationally light like rule based methods whilst producing near optimal solutions of the optimisation methods.

- The global finite resource constraint must be considered to enable graceful degradation.

- The role of information theory in multifunction radar resource management can be further investigated and ideally exploited.

- Priority assignment can be better used to dictate system performance by transforming priority into behaviour.

These areas are targeted in the development of the work in this thesis.

Chapter 4

Resource Allocation Measures and Models

Non-multifunction radars are performance tested by relevant measures during the design stage. In contrast, multifunction radars are capable of dynamically adjusting performance online and so require evaluation during run time. The measures which provide the evaluation and the associated models which relate control parameters to performance become an integral part of guiding the allocation of the resource for multifunction systems.

Resource allocation mechanisms and techniques require a single measure to optimise. For optimisation this means the choice of objective function, or for an agent system utilising an economic paradigm this means the choice of utility, which represents the currency in the system. By describing the problem to be solved, function measures and models act as an interface between the task functions and the resource allocation mechanism. Clearly the calibre of the measures and models critically limits the quality of the decision making process.

This chapter discusses existing and explores new measures which can be used by the resource allocation mechanisms which are developed in the following chapters. Sec. 4.1 describes task specific measures and the methods which can be used to model them. In Sec. 4.2 information theoretic measures are explored and applied to estimation and discrimination problems, which are at the heart of surveillance and tracking functions. In Sec. 4.3 these measures are analysed in terms of their suitability for radar resource management. Finally, in Sec. 4.4, the concept of utility as a single common measure is introduced.

4.1 Task Specific Measures

There are a number of performance measures which are specific to the objective each function is aiming to achieve. These measures are numerous and incomparable between functions, which creates difficulty for the control of multiple functions. Given a number of task specific measures, models are required to estimate the relationship between task parameters and performance. The aim of this section is to explore task specific measures and models which can be incorporated into the resource allocation mechanism.

4.1.1   Surveillance

The surveillance function has the purpose of detecting targets to track or providing measurements for existing tracks. Surveillance performance measures are based around detection performance, detection range or track acquisition performance.

4.1.1.1   Loading

An essential measure of surveillance task performance is the resource loading it is currently exerting on the radar system. Intuitively, the resource loading l_d for a dwell length τ_d and revisit interval t_f is:

$$l_d = \frac{\tau_d}{t_f} \tag{4.1}$$

which can be expressed as a percentage or as a power. Resource loading as a function of dwell and revisit times, using Eq. 4.1, is shown in Fig. 4.1 where it can be seen that greater loading occurs at longer dwell and shorter revisit times. The loading is the cost at which a certain performance level is achieved.
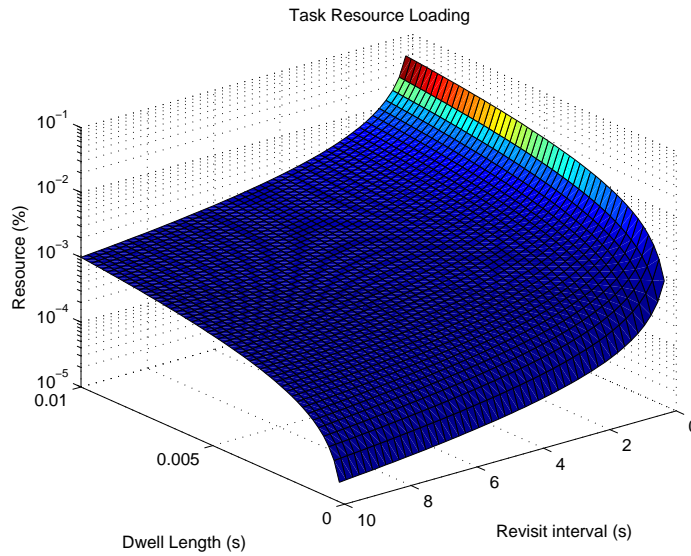
Figure 4.1: Radar loading for dwell length and revisit interval parameters

4.1.1.2   Single Look Probability of Detection

The single look detection probability quantifies the probability of a detection occurring on a look containing coherent or incoherent integration. As demonstrated in Sec. 2.2.1.2 the single look probability of detection for a Swerling 1 target can be modelled as a function of signal to noise ratio SNR, and probability of false alarm P_{FA}:

$$P_D = P_{FA}^{1/(1+SNR)} \tag{4.2}$$

The SNR as a function of range can be calculated from the standard range equation [Skolnik, 2008], including a two way loss of gain by a factor of $\cos^3\theta$ which results from off-boresight beam scanning as discussed in Sec. 2.1.3. An example of the single look probability is shown as a function of range in Fig. 4.3.

4.1.1.3   Detection Range

The single look probability of detection can be used to calculate the range at which a specified detection probability is achieved. For example, the range R_{50} can be defined as the range where a certain target is detected with a probability of 0.5. This can be calculated using the standard radar range equation Eq. 2.8 and assuming a Swerling 1 target to give a probability of detection according to Eq. 4.2. In addition the instrument range R_0 is defined as the range at which the SNR is unity, or zero dB. Both the instrumental range and R_{50} are marked in Fig. 4.3 assuming an instrumental range $R_0 = 200km$, $10ms$ coherent dwell with probability of false alarm 10^{-5}.

Fig. 4.2 shows an example of the range at which the single look detection probability exceeds 0.8 as a function of the dwell length. This measure can be useful for controlling the self protect search function where pop-up targets require a high probability of detection.
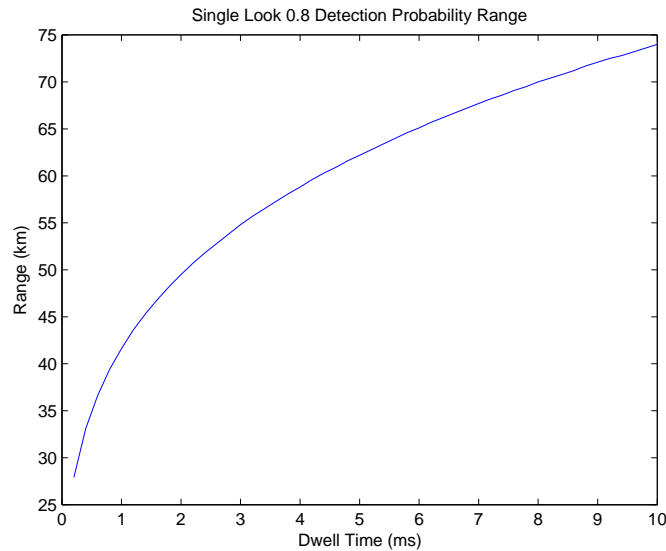
Figure 4.2: Single look 0.8 detection probability range as a function of coherent dwell length

4.1.1.4   Closure Range

The closure range is the distance a target travels towards the radar on an assumed trajectory between successive scans. The closure range d_c is a function of the radial velocity v_r and the revisit interval t_f:

$$d_c = v_r . t_f \tag{4.3}$$

this measure is useful for assessing performance of the self protect search function where it is desirable to minimise the closure range for pop-up targets.

4.1.1.5   Cumulative Detection Range

The cumulative detection range is the range at which a target is detected over successive scans with a specified cumulative probability. Typical values for the cumulative detection range are 0.9, denoted R_{90} or 0.85, denoted R_{85}.

The cumulative detection probability can be modelled by assuming a radar cross section σ and radial velocity v_r. The cumulative detection probability P_{Dc} can be calculated from a sequence of i single look probabilities of detection P_D [Blackman and Popoli, 1999]:

$$P_{Dc} = 1 - \prod_{j=0}^{i}(1 - P_{Dj}) \tag{4.4}$$

The range for each measurement in the sequence in Eq. 4.4 is found by decrementing the measurement range by the closure range. The target can be modelled as arriving at any range, for long range surveillance this could be the instrumental range R_0 or for a self protect surveillance this could be the range at which the target breaks the horizon which is approximately $25 - 30km$. The single look and cumulative detection probability is shown as a function of range in Fig 4.3 where the R_{90} is marked. The cumulative detection range is produced using the previous assumptions, a $10s$ revisit interval, $300ms^{-1}$ target radial velocity and $1m^2$ radar cross section.
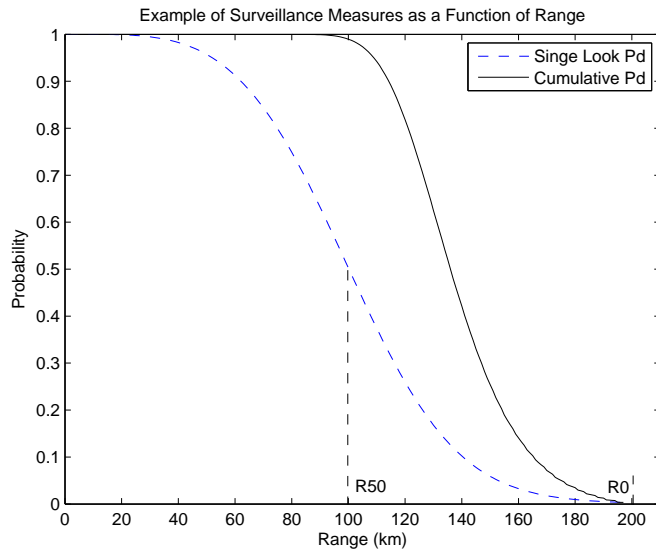
Figure 4.3: Single look and cumulative detection probabilities as a function of range

The cumulative detection range as a function of the dwell length and revisit interval parameters is shown in Fig 4.4. The cumulative detection range is useful for assessing performance of the long range

search function where it is desirable to detect targets at long ranges.
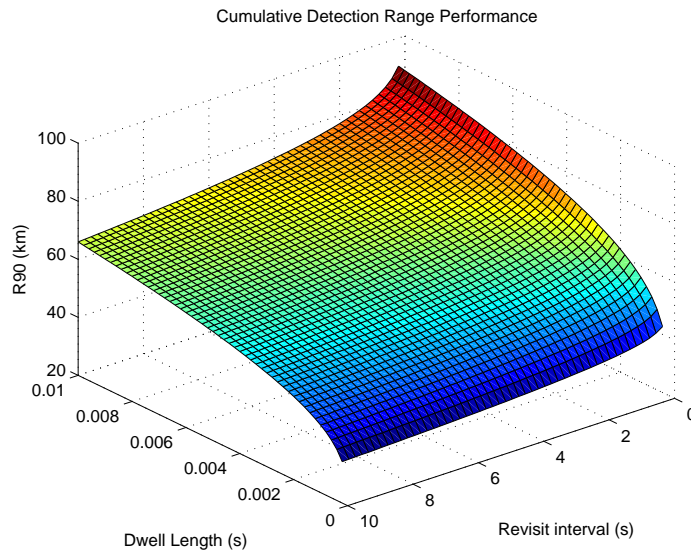
Figure 4.4: Cumulative detection range as a function of dwell length and revisit interval parameters

4.1.1.6   Track Acquisition

As detections are a precursor to acquiring tracks, it may be of more interest to measure the acquisition probability of the tracking system. In a similar way to the cumulative detection range, the range at which a target track is confirmed without further deletion can be defined. Specific target geometry and radar cross section is assumed, and the model of acquisition range depends on the track initiation method. Track acquisition range is a useful performance measure for surveillance where tracks are maintained through track-while-scan. Similarly, additionally measures based on the kinematic accuracy of the observations produced by the surveillance scan can be used for assessing the performance of track-while-scan.

4.1.2   Tracking

Tracking measures are most commonly related to the predicted state estimation error which is extracted from the tracking filter. However, there are additional measures such as the likelihood a set of measurements originated from a target or the probability that a target exists, which can be utilised to provide additional functionality for the multifunction system.

4.1.2.1   Track Loading

The amount of loading an active track exerts on the multifunction radar is a critical measure of task performance. Minimising the individual track loading is a local objective which is complementary to the global objective of maximising the number of targets in track. The loading of a tracking task can be calculated as in Eq. 4.1 and as with surveillance is the cost associated with a certain performance level.

4.1.2.2   Predicted State Estimation Error

The measure of the current accuracy of the track can be taken as the predicted estimation error extracted from the tracking filter. General sensor management approaches, which often model measurements in cartesian coordinates, use the determinant or the trace of the predicted estimation error. However for radar applications where measurements are received in cartesian coordinates, it is common to express track accuracy in terms of the angular predicted estimation error along the major axis defined by the state uncertainty ellipse. It is usually assumed that the track accuracy is correct, which means the filter is perfectly matched to the target dynamic. In reality this would not be the case and the mismatch can potentially undermine the resulting allocation.

The choice of coordinate system used to measure accuracy can impact the allocation when used as an objective function. A measurement uncertainty ellipse resulting from a specified range and angle accuracy is constant in polar coordinates for all range. However, the measurement uncertainty ellipse increases in cartesian coordinates with range. Although measuring uncertainty in cartesian coordinates is equally valid, and highlights the measurement characteristic of the radar, it could potentially add preference to closer targets where the cross range distance is smaller.

The predicted angular estimation error and the trace and determinant of the estimation error in cartesian coordinates are shown in Fig. 4.5 for a fixed update interval of $0.8s$ and a target on the radar boresight at $60km$, which has an instrumental range $R_0 = 200km$. It can be seen that between updates the target dynamic noise modelled by the filter, which is a continuous white noise jerk model with process noise intensity 10 causes the uncertainty to grow. The uncertainty is reduced when an update occurs, however, the magnitude of the reduction in uncertainty depends upon the instantaneous SNR, which in turns depends upon the target location within the beam. The estimation error is at its greatest at the start of the simulation before the filter reaches steady state. All the accuracy measures have similar characteristics, however, they have differing units and magnitudes. The black dotted lines represent the measurement accuracy standard deviation. The ratio between the measurement accuracy standard deviation and the estimation error standard deviation is known as the variance reduction ratio.

To allocate resource to a track it is necessary to model the relationship between the predicted estimation error and the parameters selected. Van Keuk provides an approximate method to achieve this, alternatively covariance analysis can be used to model at a higher fidelity. These models are derived, discussed and compared here at length as they form the basis for the simulations and the allocation models used in the following chapters.

Van Keuk Model

Van Keuk models the relationship between track loading and predicted estimation error under the condition that track updates are scheduled at times when the predicted estimation error is equal to a fraction of the beamwidth, known as track sharpness v_0. Assuming Singer target dynamics, Van Keuk

(a) Standard deviation of angular predicted estimation error

(b) Determinant of predicted estimation error

(c) Trace of predicted estimation error
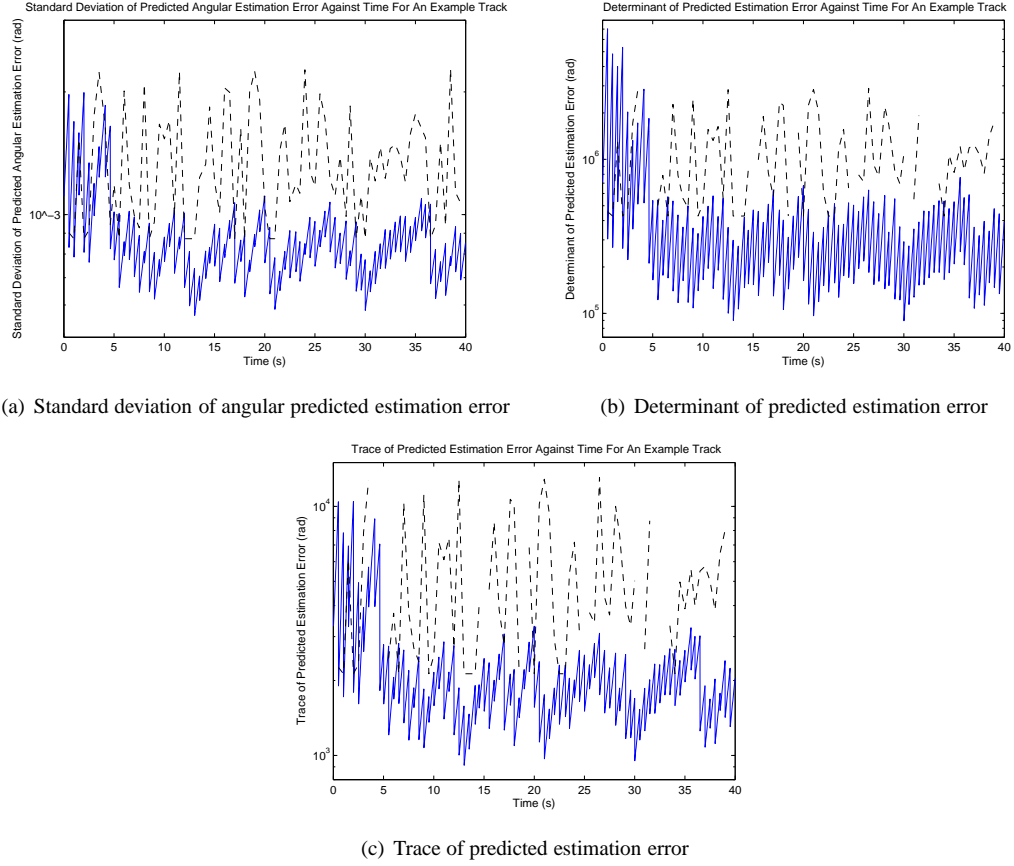
Figure 4.5: Measures of predicted track estimation error

provides an empirical approximation which relates the revisit interval to the predicted angular estimation error at the time of the track update:

$$T \approx 0.4 \left(\frac{R_t \sigma_m \sqrt{\Theta}}{\Omega} \right)^{0.4} \frac{U^{2.4}}{1 + \frac{1}{2}U^2} \tag{4.5}$$

where Ω and Θ are the manoeuvre standard deviation and time respectively. U is the variance reduction ratio, which for a track sharpness v_0 is equal to:

$$U = \frac{\theta_B v_0}{\sigma_m} \tag{4.6}$$

the measurement error standard deviation σ_m is modelled as in Sec. 2.1.2.2. Van Keuk also estimates that at each update the search strategy required to produce a detection results in an expected number of beam positions n which is related to the predicted angular estimation error according to:

$$E[n] = \frac{1}{P_{D0}}(1 + (\alpha v_0^2)^2)^{0.5} \tag{4.7}$$

77

where:

$$\alpha \approx 1 + 14\sqrt{\left(\frac{|\ln P_{FA}|}{SNR_0}\right)} \tag{4.8}$$

and P_{D0} is the probability of detection on the beam centre. The tracking load as a function of SNR and track sharpness v_0 parameters is shown in Fig. 4.6 assuming a $240km$ instrumental range with a nominal dwell of $20ms$ and target manoeuvre standard deviation $10m^2$ and $60s$ time constant. It can be seen that a minimum tracking load occurs around 0.15 of the beamwidth with a broad minimum for a wide range of SNR. The performance of the track as a function of dwell and revisit parameters is shown in Fig. 4.7 under the same assumptions. The tooth like structure of the surfaces in this figure are an artefact resulting from the finite number of sampling points used to generate the surfaces in Matlab. It can be seen that angular estimation error is reduced for small revisit intervals and large SNR, with the greatest dependence on the revisit interval. This model is insightful but highly empirical and so further models are useful which more directly model the tracking process.

(a) Van Keuk model

(b) Adapted Van Keuk model

(c) Covariance analysis

(d) Approximate covariance analysis
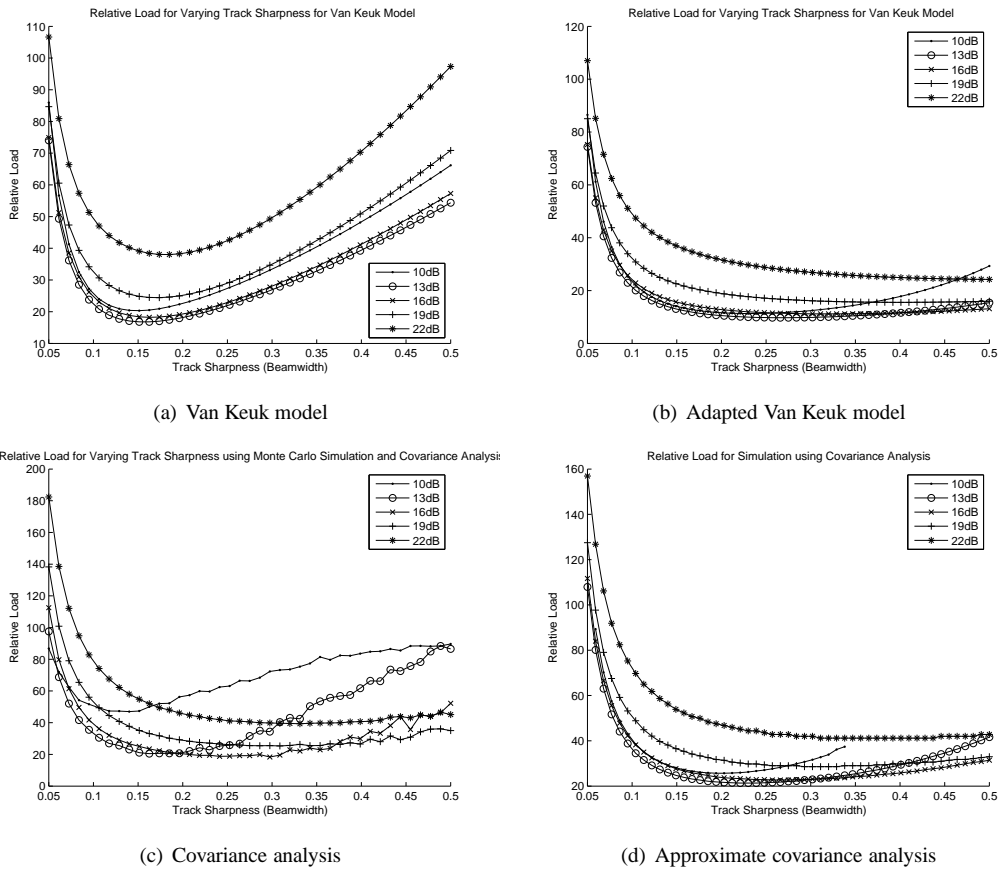
Figure 4.6: Tracking loading using various models for differing signal to noise ratios

Adapted Van Keuk Model

The original Van Keuk model is adapted by Blackman [1986] to account for non-unity probability

(a) Van Keuk model

(b) Adapted Van Keuk model

(c) Covariance analysis

(d) Approximate covariance analysis
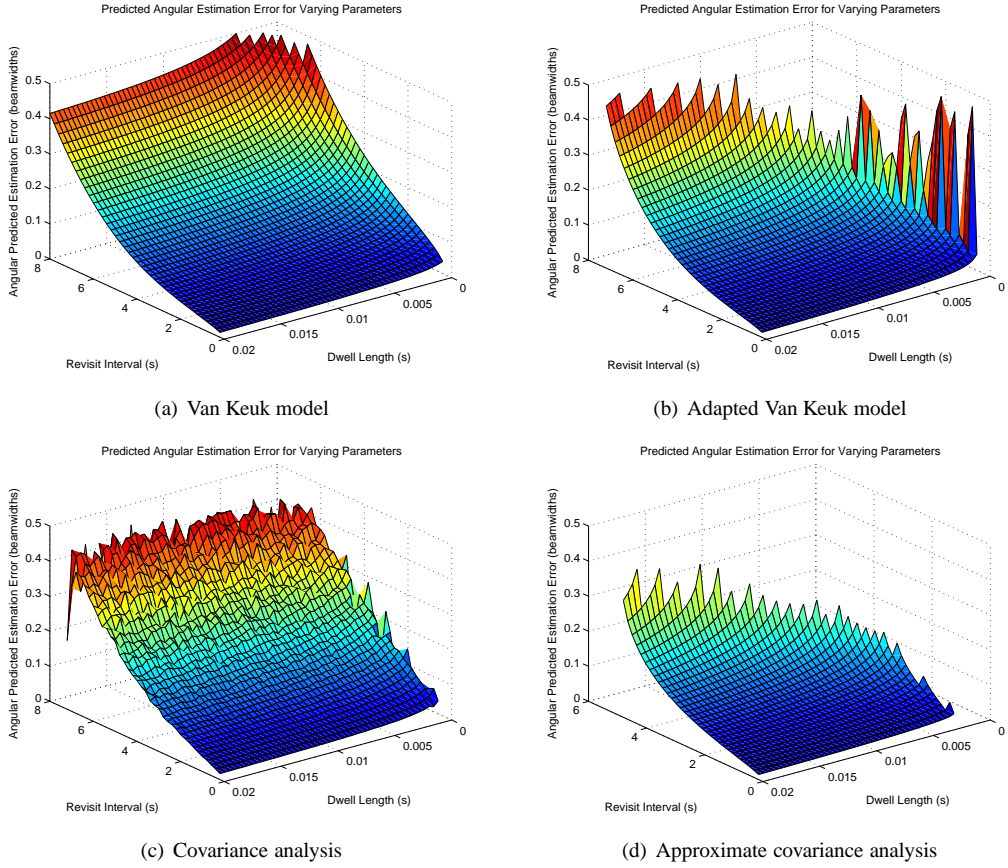
Figure 4.7: Tracking performance using various models for differing signal to noise ratios

of detection:

$$T \approx \frac{0.4}{P_D} \left(\frac{R_t \sigma_m \sqrt{\Theta}}{\Omega} \right)^{0.4} \frac{U^{2.4}}{1 + \frac{1}{2}U^2} \tag{4.9}$$

The probability of detection is a function of the boresight signal to noise ratio and the current target state. Assuming the tracker is ideally matched to the target dynamic, the target state will be distributed according to the predicted state estimation error. The subsequent loss in signal to noise ratio due to the target being offset from the centre of the beam can be modelled as a Gaussian loss function as in Eq. 2.19 from Sec. 2.1.2.2. The expected signal to noise ratio loss can be found by integrating over the target's predicted position [Blackman and Popoli, 1999] to give:

$$SNR_{\theta\phi} = SNR_0 \frac{\theta_B}{\sqrt{\theta_B^2 + 2C_L\sigma_v^2}} \frac{\theta_B}{\sqrt{\theta_B^2 + 2C_L\sigma_u^2}} \tag{4.10}$$

where $C_L = 2.77$ and σ_u and σ_v are the standard deviations of state estimation error in azimuth and elevation. This reduced SNR can be used to calculate the probability of detection according to Eq. 4.2. The expected SNR given a detection has occurred can be approximated as:

$$SNRm = SNR - \log(P_{FA}) \tag{4.11}$$

79

which can be used in Eq. 2.18 to estimate the measurement accuracy. The expected angular offset related to target uncertainty can be approximated as:

$$E[\theta_{u,v}] = 2\pi\sigma_{u,v} \tag{4.12}$$

which can be used in Eq. 2.20 to approximate the off beam centre measurement accuracy.

The tracking load resulting from varying track sharpness and boresight SNR is shown in Fig. 4.6 for comparison with the original Van Keuk model using the same assumed parameters. It can be seen that the original Van Keuk model has a higher increase in tracking load for greater track sharpness. The difference is a result of the empirical approximation of the expected number of beams required in Eq. 4.7, it can be assumed that this approximation includes the effect of miss-association where as the adapted model does not. However, the relevance of this difference is somewhat moot for a resource allocation model as it is undesirable to operate in this region, due to a higher accuracy being achievable with a lower track sharpness and hence a reduction in loading. The performance of the track as a function of dwell and revisit parameters is shown in Fig. 4.7. Similarly to the Van Keuk model it can be seen that angular estimation error is reduced for small revisit intervals and large SNR, with the greatest dependence on the revisit interval. However in this model long revisit intervals are often lost because revisits are not assumed on a missed detection.

Covariance Analysis

Covariance analysis can be used through Monte Carlo simulation to analyse the prediction estimation error in the track for comparison to the Van Keuk models. This has been applied using a Kalman filter with a continuous white noise jerk dynamic model which is equivalent to the limiting Singer model where the manoeuvre time is much greater than the sampling time [Blackman and Popoli, 1999]. When updates are executed a target position is generated according to the target state which is assumed matched to the track estimation error. The SNR resulting from the target position being offset from the centre of the beam can be calculated using Eq. 2.19. Assuming Swerling 1 fluctuations, an instantaneous SNR can be generated as a sample from the signal envelope [Blackman, 1986]:

$$SNRm = 0.5 * (((A + n_1)^2) + n_2^2); \tag{4.13}$$

where n_1 and n_2 are standard normal variables, $A = \sqrt{2SNR}$. An update is performed if the instantaneous SNR exceeds a threshold determined by the false alarm probability, i.e $T = -\log(P_{FA})$. If a successful detection occurs then the instantaneous SNR is used to calculate the measurement accuracy of the subsequent update. If a detection does not occur a revisit is scheduled after 0.1s.

The tracking load for track sharpness and boresight signal to noise ratio is shown in Fig 4.6 for

comparison to the Van Keuk models. It can be seen that the adapted Van Keuk model underestimates the rise in the tracking load as the track sharpness increases. However, this Monte Carlo simulation demonstrated that this region is highly unstable with large track loss which reiterates that operating in the region is undesirable. Although including the probability of false alarm, this model does not include the effect of false returns. The performance of the track as a function of dwell and revisit parameters is shown in Fig. 4.7 and found to be similar to the Van Keuk model.

Approximate Covariance Analysis

Instead of using Monte Carlo simulation, the random elements in the previous covariance analysis simulation can be replaced with their expectations. The same Kalman filter can be used, however, the expected signal to noise ratio can be used instead of generating samples of the envelope and the expected offset in Eq. 4.12 can be used instead of the random offsets.

This approach can be extended to include the effect of measurement origin uncertainty which necessitates data association. The track estimation error depends on the measurement sequence and can only be evaluated through numerical simulation, however, it can be estimated through the modified Riccati equation, as first derived in Fortmann et al. [1985], which replaces the random elements in the covariance update equation Eq. 2.53 with their expectations, to give:

$$P_{k|k} = P_{k|k-1} - q_2 W_k S_k W_k'$$
(4.14)

where the scalar q_2, which takes values between 0 and 1 to represent the measurement origin uncertainty, is a function of P_D and the clutter density μ. The calculation of q_2 is non-trivial and requires numerical integration, however Kershaw and Evans [1996] give an analytic approximation of q_2:

$$q_2 = \frac{0.997 P_D}{1 + 0.37 P_D^{-1.37} V_k \mu}$$
(4.15)

which is of sufficient accuracy and allows online computation.

The tracking load resulting from this model is shown for comparison in Fig 4.6. It can be seen that it is in close agreement with the adapted Van Keuk model. The performance of the track as a function of dwell and revisit parameters is shown in Fig. 4.7 and found to be similar to the adapted Van Keuk model whereby revisits on missed detections are not scheduled.

The Van Keuk models are useful for resource allocation as they give a simple and computationally light relation between parameters and performance. However, with the increase in computational power, more exact covariance analysis could be useful for online resource allocation, and the modified Riccatti equation offers an enhanced performance assessment in cluttered regions.

4.1.2.3   Root Mean Squared Error

The root mean squared error (RMSE) can be used in simulations to assess the absolute performance of the state estimate when the true kinematics are known. The root mean squared error for a position vector $X = [xy]$ is equal to:

$$\hat{X} = \sqrt{\frac{1}{N}\Sigma_{i=1}^{N}(\tilde{x}^2 + \tilde{y}^2)} \tag{4.16}$$

The root mean squared of zero mean random variable is the standard deviation, hence the RMSE when the filter is perfectly matched to the target dynamic is the root of the trace of the state estimation error covariance. As the true target kinematics are not known in reality, RMSE is only useful for assessing performance in a simulation.

4.1.2.4   Likelihood Ratio

The likelihood of a measurement belonging to a target or to clutter can be measured using the measurement likelihood ratio or more commonly the log likelihood ratio. The measurement likelihood ratio is expressed as [Blackman and Popoli, 1999]:

$$LR = \frac{p(D|H_T)P_0(H_T)}{p(D|H_N)P_0(H_N)} \equiv \frac{P_T}{P_{FA}} \tag{4.17}$$

where H_T and H_N are the probability of presence of true target and false alarm respectively, given data vector D. This can be modelled by assuming the target returns have a Gaussian distribution and clutter returns are distributed uniformly in the track validation gate with a density of μ. Changes in the measurement log likelihood ratio can be computed as:

$$\Delta L(k) = \begin{cases} \ln[1 - P_D]; & \text{no detection} \\ \triangle L_U(k); & \text{detection on scan } k \end{cases} \tag{4.18}$$

For detection only data the measurement log likelihood ratio $\triangle L_U(k)$, is given by:

$$\triangle L_U = \ln\left[\frac{P_D}{(2\pi)^{M/2}\mu\sqrt{|S_k|}}\right] - \frac{d^2}{2} \tag{4.19}$$

where M is measurement dimension, μ is the false target density S_k is the residual covariance matrix and d^2 is the normalised statistical distance for the measurement.

The likelihood of a set of measurements being due to a target can be found by combining the individual measurement likelihood ratios. This can be expressed as a recursive formula:

$$L(k) = L(k-1) + \Delta L(k) \tag{4.20}$$

Thresholding on the likelihood ratio associated with a set of measurements can be used to determine track confirmation or deletion status. Thresholds are used for track deletion and confirmation given by:

$$T_2 = \ln\left[\frac{1-\beta}{\alpha}\right], \quad T_1 = \ln\left[\frac{\beta}{1-\alpha}\right] \tag{4.21}$$

where α is the false track confirmation probability and β is the true track deletion probability. Fig. 4.8 shows an example of the value of the likelihood ratio as successive measurements are received during track initiation. As this is a true target the likelihood ratio can be seen to increase by an amount dependent on the measurement residual. This is produced assuming a target at $60km$ on the radar boresight with a $1°$ beamwidth, instrumental range $R_0 = 240km$ with a nominal coherent dwell of $20ms$ and a $0.5x10^{-6}/m^2$ false target density. The target dynamic is assumed to evolve according to a continuous white noise jerk process model with process noise $\tilde{q} = 3.33$.

(a) Likelihood ratio   (b) Track existence
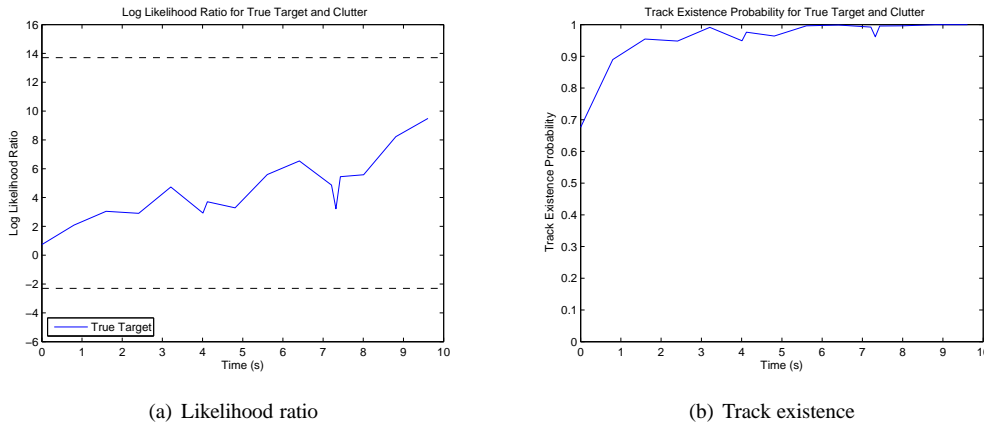
Figure 4.8: Likelihood ratio and track existence examples for track initiation

The likelihood ratio is useful for the performance assessment of track initiation, which can be controlled to reduce the number of updates required to release track. The number of updates required to release the track during initiation is shown in Fig. 4.9. The tooth like structure of the surfaces in this figure are an artefact resulting from the sampling points used to generate the surfaces in Matlab. It can be seen that for a given false target density, assuming an independent measurement, the number of updates required is lowest for short revisit intervals as the validation volume is smaller. In reality clutter returns are not uniformly distributed within the track validation gate and have spatial and temporal coherence which can reduce the effectiveness of this method.

4.1.2.5 Track Existence

The probability of track existence as a measure of quality is introduced by Musicki et al. [1994] where it is built into the probabilistic data association framework. The probability that the target exists $p(x)$ and the probability that the target does not exist $p(\hat{x}) = 1 - P(x_k)$ is modelled as Markov process which

transitions between stages according to a Markov chain:

$$P(x_k) = p_{11}P(x_{k-1}) + p_{21}(1 - P(x_{k-1})) \tag{4.22}$$

$$1 - P(x_k) = p_{12}P(x_{k-1}) + p_{22}(1 - P(x_{k-1})) \tag{4.23}$$

where p_{11}, p_{21}, p_{12} and p_{22} are the Markov chain coefficients where $p_{11} + p_{12} = p_{21} + p_{22}$. The enables the recursive calculation for the track existence probability:

$$P(x_k|Z^k) = \frac{1 - \delta_k}{1 - \delta_k P(x_k|Z^{k-1})} P(x_k|Z^{k-1}) \tag{4.24}$$

where δ_k is related to the likelihood ratio as $1 - \delta_k = LR$. This is a useful measure for tracks where the purity of tracks are of a higher importance than the accuracy which is often the case in surveillance systems. Fig 4.8 shows an example of the track existence probability using the same assumptions as for the likelihood ratio. It can be seen that it is also a valid measure which is similar to the likelihood ratio.

(a) SPRT
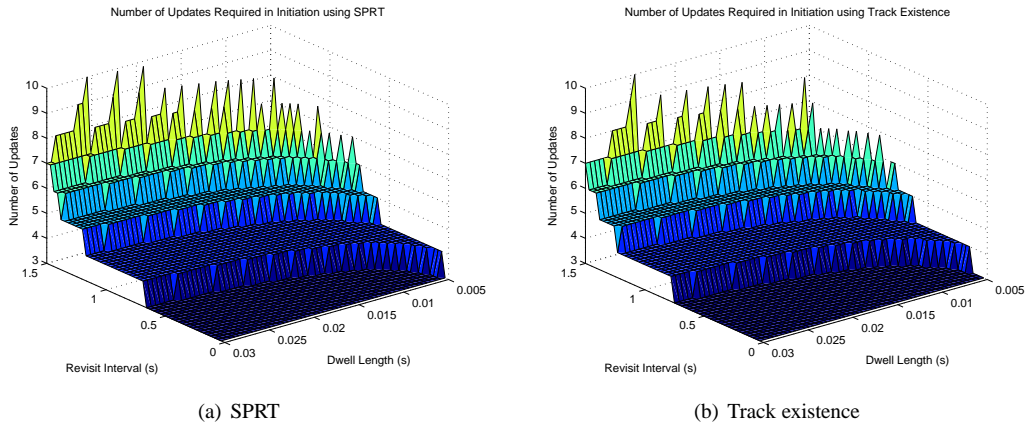
(b) Track existence

Figure 4.9: Number of updates required to release track

Track existence is a useful performance measure for track initiation and could be useful for allocating resource in clutter regions. This is shown in Fig. 4.9 where similar conclusions can be drawn as for the likelihood ratio.

4.2 Information Theoretic Measures

Measures based on information theory differ from task specific measures as they are surrogate functions being independent of the specifics of each task. It has been suggested that information theoretic measures can provide a 'universal proxy' [Kreucher et al., 2005a]. As information theoretic measures were identified as potentially beneficial in Sec. 3.3.2.2 an aim of this thesis is to investigate their role in multifunction radar resource management. This section derives relevant information theoretic measures and describes their application to sequential estimation and detection.

4.2.1 General Derivations

This subsection derives and highlights the interplay between relevant information theoretic measures. Detailed discussion on these measures are found in the texts by Cover and Thomas [2006] and Hero et al. [2007].

4.2.1.1 Fisher Information

Fisher information, J, is the mean curvature of measurement log likelihood function $\ln p(Z|x)$ and quantifies the amount of information that can be extracted from the measurement:

$$J_z = E[[\nabla_x \ln p(Z|x)][\nabla_x \ln p(Z|x)]'] \tag{4.25}$$

The inverse of the Fisher Information Matrix is related to the Cramer-Rao lower bound which bounds the variance of the subsequent maximum likelihood estimates $\hat{x}$ of x, and hence the track accuracy:

$$E[[\hat{x}(Z) - x][\hat{x}(Z) - x]'] \geq J_z^{-1} \tag{4.26}$$

In Trees [2001] the likelihood function is related to the radar ambiguity function, the curvature of which is related to the SNR and the signal bandwidth. In angle the Fisher information is a function of SNR and beamwidth. Maximisation of the Fisher information as an objective function means choosing between measurement likelihood functions and so it is useful for selecting a sensor or sensor mode.

4.2.1.2 Differential Entropy

Differential Entropy is a measure of the uncertainty of a continuous random variable. Employing a measure of uncertainty is logical as it is the role of the sensor to reduce uncertainty about the environment. Given a random variable X and its probability density function $p(x)$, differential entropy is defined as:

$$H(X) = - \int p(x) \log p(x).dx \tag{4.27}$$

It is also useful to measure the entropy of the random variable X conditioned on the variable Z, given the density function $p(z)$ and the conditional density function $p(x|z)$:

$$H(X|Z) = - \int p(z).dz \int p(x|z) \log p(x|z).dx \tag{4.28}$$

Conditional entropy is the expectation of the entropy of the conditional probability density function $p(x|z)$ with respect to Z.

4.2.1.3 Kullback-Leibler Divergence

Kullback-Leibler divergence (KLD) is a measure of discrimination between two distributions. Given the random variable X and two probability density of functions of X, $p(x)$ and $q(x)$, the Kullback-Leibler divergence is defined as:

$$D_{KL}(P||Q) = \int p(x) \log \frac{p(x)}{q(x)}.dx \qquad (4.29)$$

Kullback-Leibler divergence is always non negative and equals zero when $p = q$.

4.2.1.4 Mutual Information

Mutual information (MI), denoted I, is the Kullback-Leibler divergence between joint and product distributions and is the reduction in uncertainty in the random variable X due to knowledge of Z.

$$I(X;Z) \qquad = D(p(x,z)||p(x)p(z)) \qquad (4.30)$$

$$= \int p(x,z) \log \frac{p(x,z)}{p(x)p(z)}.dx.dz \qquad (4.31)$$

It is straightforward to show that mutual information is the difference between a random variables entropy and conditional entropy:

$$I(X;Z) \quad = H(X) - H(X|Z) \qquad (4.32)$$

$$= H(Z) - H(Z|X) \qquad (4.33)$$

Also, mutual information is the expectation of the Kullback-Leibler divergence between a random variable's probability density function and its conditional probability density function, with respect to Z:

$$I(X;Z) = \int p(z).dz D(p(x|z)||p(x)) \qquad (4.34)$$

These information theoretic measures have roles in estimation problems, which relate to target tracking, and discrimination problems, which relate to surveillance.

4.2.2 Information in Estimation

In sequential state estimation, which is at the heart of target tracking, a sequence of received measurements are combined with prior measurements within a Bayesian framework. The information gain attributable to the measurement can be described in terms of mutual information or Kullback-Leibler divergence. The information gain from the measurement can then be used as a measure for resource allocation.

Mutual information gain of the measurement can be found by calculating the mutual information between the prior state distribution $p(X^k|Z^{k-1})$ and the measurement z_k. For notational convenience the

conditioning on Z^{k-1} is assumed and omitted. A large reduction in uncertainty means the measurement contained a large amount of information on the state, and so the two have large mutual information. Mutual information between the state and measurement can be written, from Eq. 4.31, as:

$$I(X^{k|k}; z_k) = \int p(X^k, z_k) \ln \frac{p(X^k|z_k)}{p(X^k)} .dx.dz \qquad (4.35)$$

Through some algebra this can be rearranged as a difference of the measurement entropy and state conditioned measurement entropy, similar to Equation 4.33:

$$I(X^k; z_k) = H(z_k) - H(z_k|X^k) \qquad (4.36)$$

$$= H(X^k) - H(X^k|z_k) \qquad (4.37)$$

assuming a Kalman filter, the mutual information can be shown to be dependent on the predicted state covariance at time k given the measurement up to time $k-1$ and the measurement noise R_k:

$$I(\hat{x}_k; z_k) = \frac{1}{2} \ln \left(|I + R_k^{-1} P_{k|k-1}| \right) \qquad (4.38)$$

This can be calculated before the measurement is made, due to the entropy of multivariate Gaussians being a function of just their covariance. Intuitively, this tells us that large information is produced from accurate measurements of uncertain targets.
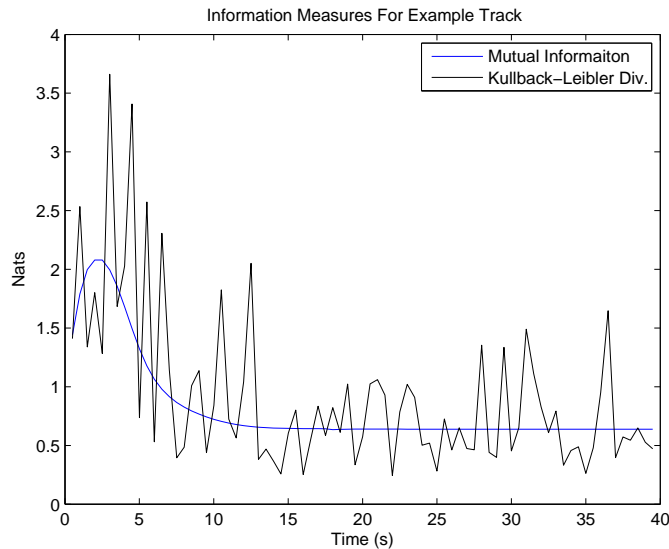
Figure 4.10: Mutual information and Kullback-Leiber divergence

Alternatively the Kullback-Leibler divergence can quantify the information gain of the measurement by determining the divergence between the prior state distribution $p(X^k|Z^{k-1})$ and the posterior

state distribution $p(X^k|Z^k)$. The greater the divergence between distributions, the larger the information gain from the measurement. The divergence between prior and posterior states is:

$$D_{KL}(z_k) \quad = D(p(X^k|Z^k)||p(X^k|Z^{k-1})) \tag{4.39}$$

$$= \int p(X^k|z_k) \ln \frac{p(X^k|z_k)}{p(X^k)}.dx \tag{4.40}$$

Again, in Eq. 4.40 the conditioning on Z^{k-1} is assumed. Assuming a Kalman filter Eq. 4.40 can be rewritten as:

$$
\begin{aligned}
D_{KL}(z_k) \quad &= \tfrac{1}{2}\big(\ln(|S_k R_k^{-1}|) \big) + \tfrac{1}{2}\big(tr(R_k S_k^{-1})\big) \\
&+ \tfrac{1}{2}(W_k \tilde{z}_k)' P_{k|k-1}^{-1}(W_k \tilde{z}_k) - \tfrac{m}{2}
\end{aligned}
\tag{4.41}
$$

where $\tilde{z}_k$ is the measurement residual, the difference between predicted and observed state $\tilde{z}_k = z_k - H\hat{x}_k$, and m is the dimension of the measurement. Calculation of the third term requires the measurement, which is not known before the sensor action.

As stated in Eq. 4.34 the mutual information is the expectation of the Kullback-Leibler divergence. The Kullback-Leibler divergence contains the statistical distance term which assuming the filter is matched to the target dynamic will be distributed according the innovation covariance S_k. Hence the KLD fluctuates around the MI value as shown in Fig. 4.10.

4.2.3 Information in Discrimination

As previously noted, detection is a key element of the surveillance function. There are two ways that the Kullback-Leibler divergence can be recognised in detection, as the expectation of the log likelihood ratio and the loss of statistical power through mis-specifying distributions.

Given the log likelihood ratio which is thresholded in the hypothesis test:

$$L(z) = \log \frac{P_T(z)}{P_N(z)} \tag{4.42}$$

the Kullback-Leibler divergence can be instantly recognised as the expectation of the log likelihood ratio under target presence, and hence describes the ability to discriminate between hypotheses:

$$E_T[L(z)] = D_{KL}(P_T||P_N) \tag{4.43}$$

Similarly, the expectation given a target is not present can be:

$$E_N[L(z)] = -D_{KL}(P_N||P_T) \tag{4.44}$$

These measures of the expectation of the likelihood ratio are shown for Rayleigh noise and a Ricean

target in Fig. 4.11. This measure is somewhat trivial, as it is solely a function of signal to noise ratio.
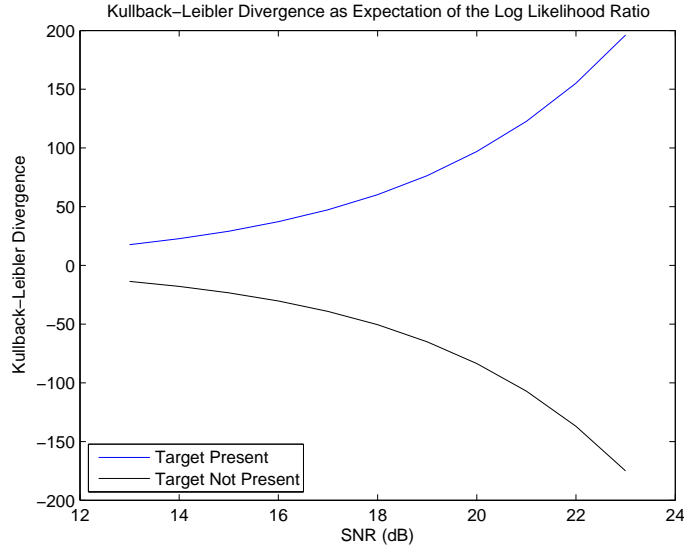
Figure 4.11: Kullback-Leibler divergence as expectation of likelihood ratio

The Kullback-Leibler divergence also has an additional interpretation through the Neyman-Pearson lemma [Eguchi and Copas, 2006], to quantify the loss in miss-specifying the target or background hypotheses. If the null background hypothesis is mis-specified as $\tilde{P_N}(z)$ then an incorrect likelihood ratio is used

$$\tilde{L}(z) = \log \frac{P_T(z)}{\tilde{P}_N(z)} \tag{4.45}$$

An example of a correct ratio formed using a target with 16dB SNR is shown in Fig. 4.12. An incorrect likelihood ratio is also shown, where the mean intensity of the noise is mis-specified. The mis-specification means the log likelihood ratios take different values.

The loss in the detection process from using the incorrect likelihood ratio can be taken as the Kullback-Leibler divergence between correct and incorrect models of the background:

$$\Delta_p = \int_{-\infty}^{\infty} P_N(L > T) - P_N(\hat{L} > T).du \tag{4.46}$$

$$= \int (L - \hat{L}) P_N(z).dz \tag{4.47}$$

$$= D_{KL}(P_N || \tilde{P_N}) \tag{4.48}$$

This loss can be visualised as the difference between the correct and incorrect log likelihood ratios measured through the correct background. However, it is easiest to visual the loss as the difference in probability of false alarm resulting from the true background and the correct and incorrect likelihood ratios, for common likelihood ratio thresholds which is shown in Fig. 4.12(b). Fig. 4.12(c) demonstrates

(a) Correct and incorrect likelihood ratios

(b) Probability of false alarm

(c) Kullback-Leibler divergence

Figure 4.12: Interpretation of discrimination information through Neyman-Pearson lemma

this interpretation by plotting the probability of false of alarm at equal thresholds of the log likelihood ratio for the correct and incorrect likelihood ratio. The previous integral is the area between the two curves.

Interestingly, although a model of a target is required to form the log likelihood ratio, the model of the target has no effect on the final value of the loss in the likelihood ratio test.

The Kullback-Leibler divergence is a useful measure of the loss in the likelihood ratio test by quantifying the magnitude of the deviation in type 1 error (false alarm) from the optimal test described by the Neyman Pearson lemma. As such it can provide a definitive measure of the loss associated with misestimating clutter backgrounds, which could be useful for the allocation of resource.

4.3 Analysis of Measures for Control

Measures aid the run time control of the multifunction radar and facilitate the interface between the task functions and the resource manager. The purpose of deriving the preceding measures was a precursor to their implementation in the allocation mechanisms developed in the following chapters. This section gives a preliminary analysis of the suitability of the measures for the control of radar functions.

4.3.1   Parameter Selection Strategy

This section describes the suitability of each of the previous measures for the control of specific radar functions.

4.3.1.1   Surveillance Function

Presently, surveillance parameters are selected subject to fixed pre-defined rules or modes which are specified during design time. Adaptation in run-time is minimal, however, it is typical to slightly adapt parameters to balance the variable tracking load. An example policy would be to implement a fixed search using 80% loading, leaving 20% for tracking tasks. When the track loading is below 20% the excess resource can be allocated to the lowest elevation search bar which is most severely affected by clutter. Simple calculations can be used to balance the time budget. The search volume in steraradians Ω_s can be calculated as [Morris and Harkness, 1996]:

$$\Omega_s = \Delta_{az}(\sin(\theta_{elMax}) - \sin(\theta_{elMin})) \tag{4.49}$$

where θ_{elMax} and θ_{elMin} are the maximum and minimum angle of the search volume and Δ_{az} is the extent of the search volume in azimuth and the beam space $D_p = 1$ The time to complete a search function τ_s is:

$$\tau_s = \frac{\Omega_s \tau_d}{\Omega_b} \tag{4.50}$$

where Ω_b is the beamwidth in steradians and τ_d is the dwell time.

It is desired that novel resource allocation mechanisms move the parameter selection from design time and into run-time, to increase the adaptation to the environment. To facilitate this, the measures described in Sec 4.1.1 have relevance for run-time evaluation of the following surveillance functions:

- *Long Range Surveillance* - Cumulative detection range or track acquisition range.

- *Medium Range Surveillance* - Track acquisition range or kinematic measurement accuracy.

- *Self Protect Search* - Single look probability of detection, target closing range.

In this case where there is more than one appropriate measure, some combination of the two measures is desirable. Utilising these identified measures combined with an improved resource allocation mechanism will unlock potential by allowing increased adaptation of performance in run-time subject to a dynamic environment.

4.3.1.2   Tracking Function

If the task revisit interval is selected according to a track sharpness setting then the minimum track loading is achieved, regardless of the targets manoeuvre and range. In Fig. 4.13(a) track loading is

plotted for varying target ranges and manoeuvre standard deviations assuming a $240km$ instrumental range for a nominal $10ms$ coherent dwell. It can be seen that the minimum track loading is found at 0.21 regardless of the varying manoeuvres and ranges which makes it a useful measure for target tracking control.

(a) Angular uncertainty
(b) Determinant cartesian
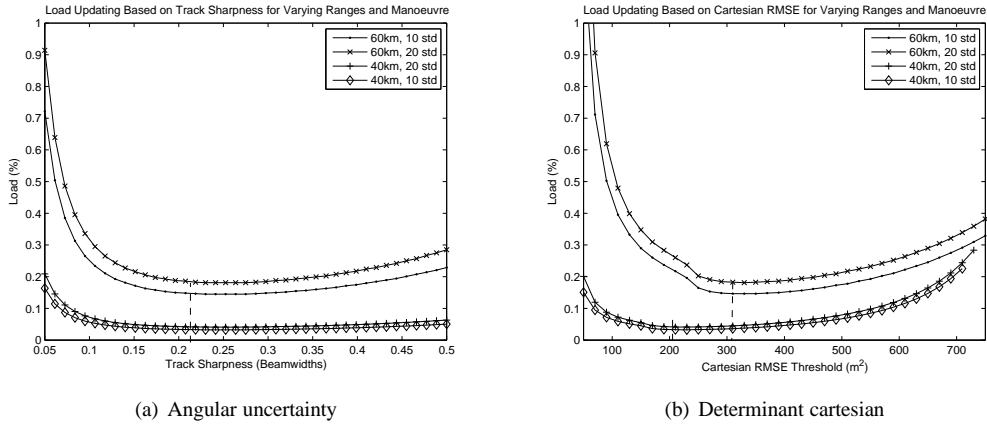
Figure 4.13: Optimal setting of track sharpness

If track updates are requested at intervals which are determined by a limit of the accuracy in cartesian coordinates then there is no common minimum loading track sharpness setting. Fig. 4.13(b) shows the loading for bounds on the trace and determinant for varying ranges and manoeuvres. Also, the determinant or trace in spherical coordinates removes the dependence on range, but is still dependent on range resolution.

In reality, where the filter dynamic model is not perfectly matched to the encountered target dynamic then adaptive sampling based on the incorrect angular uncertainty can lead to track loss on manoeuvre onset. Therefore it is common in operational systems to have fixed updates rates to prevent track loss, which indicates that track continuity can be of greater importance than minimum loading.

Maintaining the angular predicted estimation error beneath a fraction of the beamwidth is typical for tracking control. However, the choice of measure ultimately depends on the requirement of the track. For surveillance applications the track existence may be of greater importance than accuracy and so a novel resource allocation mechanism could use this measure to augment track-while-scan with active updates for troublesome clutter regions or crossing tracks. Additionally in track initiation, the release time can be of more importance.

Strategies for track allocation based on alternative measures are not widely applied. The development of resource allocation mechanisms in the following chapters aims to include a variety of measures. A variety of measures improves the interface between the task and the allocation mechanism, which improves the functionality of the system.

4.3.2 Tracking in Clutter

In clutter regions with a high false target density the number of false returns within the validation gate is larger which increases the difficulty of the data association process. This reduces the mutual information between the measurement and state estimate and increases the probability that false returns are associated to the track which reduces the purity of the track. As described in Sec. 4.1.2.2, the effect of false target density on predicted estimation error can be assessed through the modified Riccatti equation. Fig. 4.14 shows the tracking load for varying clutter density and signal to noise ratio, where it can be seen that greater false target density increases the tracking load required to maintain the track. This is most notable for greater track sharpness setting where the validation gate is larger which encompasses more false returns.
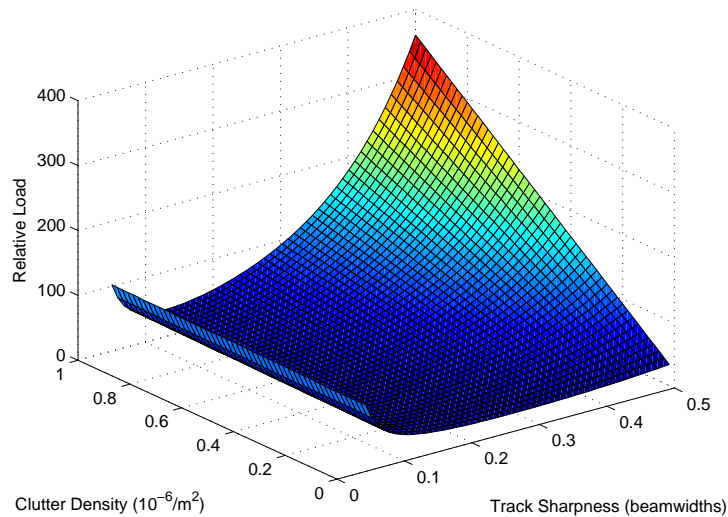
Figure 4.14: Tracking loading using track sharpness method in varying clutter density

It can be seen in Fig. 4.14 that the minimum loading track sharpness setting is not independent of the false target density and so the minimum track loading sharpness is not universal. It is found that the track sharpness setting must be lower than that suggested by Van Keuk to compensate for the measurement origin uncertainty, which echoes the requirement for lower sharpness settings in ECM found during the benchmark tests. Fig. 4.15 shows the minimum loading track sharpness as a function of the false target density where it can be seen the minimum loading decreases. This shows measurement origin uncertainty should be included in the resource allocation model to ensure stable performance in dense false target regions.

In the original paper by Fortmann the data association uncertainty resulting from increased false target density reduces the gain of the Kalman filter. This lessens the reduction in uncertainty of the measurement, which reduces the mutual information gain of the measurement. The mutual information
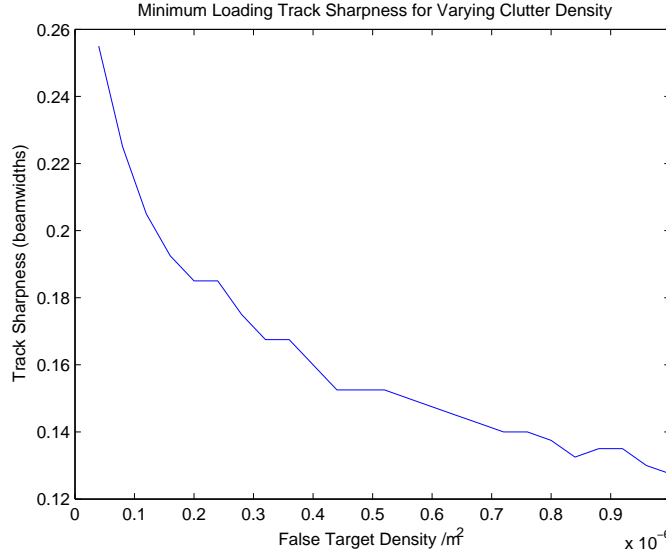
Figure 4.15: Minimum loading track sharpness setting for varying false target density

gain of the measurement given q_2 can be expressed as:

$$I = \frac{1}{2} \ln \left(\frac{|HP_{k|k-1}H^T + R_k|}{|(I - q_2 I)HP_{k|k-1}H^T + R_k|} \right) \tag{4.51}$$

maximising the mutual information gain of the measurement is desirable as the greatest reduction in uncertainty enables the longest time between revisits which reduces the tracking load. Fig. 4.16 shows an example of the mutual information gain of the measurement against the revisit interval time. It can be seen that the mutual information has a maximum at a revisit time which is dependent on the false target density. Hence, mutual information can be used to aid the selection of track revisit intervals, with additional relevance in cluttered environments.

It is desired to use track-while-scan for tracking as many targets as possible, as this is the most efficient use of resource. However, as demonstrated here, targets which are in cluttered regions require update rates which are likely to be faster than that provide by the surveillance scan. Hence, it is desired to augment track-while-scan using additional active track updates when necessary for targets in clutter regions. Enabling this is a consideration for the development of the resource allocation mechanism.

4.3.3 Comparison of Information Theoretic and Task Specific Measures

This subsection investigates the suitability of information theoretic measures in providing a single universal measure to interface into the resource allocation mechanism. To explore the use of information theoretic measures, a simulation has been produced which compares tracking performance measures for choosing updates based on the standard track sharpness method and selecting updates at times when a specified mutual information is produced.

94
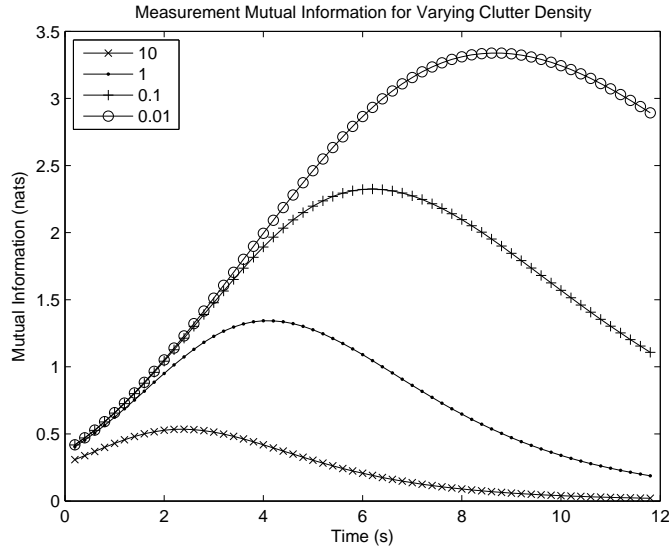
Figure 4.16: Mutual information gain of measurement as function of revisit interval for varying false target density

Loading - Fig 4.17 shows the track loading resulting from updating on a bound of the track sharpness and a bound on mutual information. Fig. 4.17 shows that there is a clear difference in the update rate between track sharpness and mutual information. Track sharpness schedules an update based on the angular estimation error alone, however, mutual information being reduction in uncertainty schedules updates based on the estimation error and the measurement error, which depends on the beamwidth and SNR. This results in a shorter revisit interval for tracks with greater SNR that produce more information. This is not necessarily desirable as it allocates more updates to high SNR tracks than they need for maintenance, and less updates to low SNR targets than they require for maintenance. This is because track sharpness aims to minimise radar load whilst maintaining track, but mutual information aims to maximise information production. So, these approaches have a fundamental difference in what they aim to optimise. This is clear from Fig 4.17 where significant load is allocated to high SNR tracks. It can also be seen that low SNR tracks are dropped using mutual information, as they are unable to produce the information required. This is an additional undesirable characteristic.

As mutual information is the expectation of the Kullback-Leibler divergence, it was found that allocating based on the Kullback-Leibler divergence performed very similarly to mutual information, with fluctuations associated to the stochastic nature of the statistical distance term.

Information Rate - The difference in the optimisation objective resulting from the different measures is further recognised through analysis of the average track mutual information rate in nats, which is information with $\log$ base of e, per second. This is shown in Fig. 4.18 for track sharpness and mutual information. It can be seen that as the bound increases the information rate for the track reduces. It can also be seen that the track sharpness has a similar information rate regardless of SNR, because this

(a) Track sharpness updates

(b) Mutual information updates

(c) Track sharpness load

(d) Mutual information load
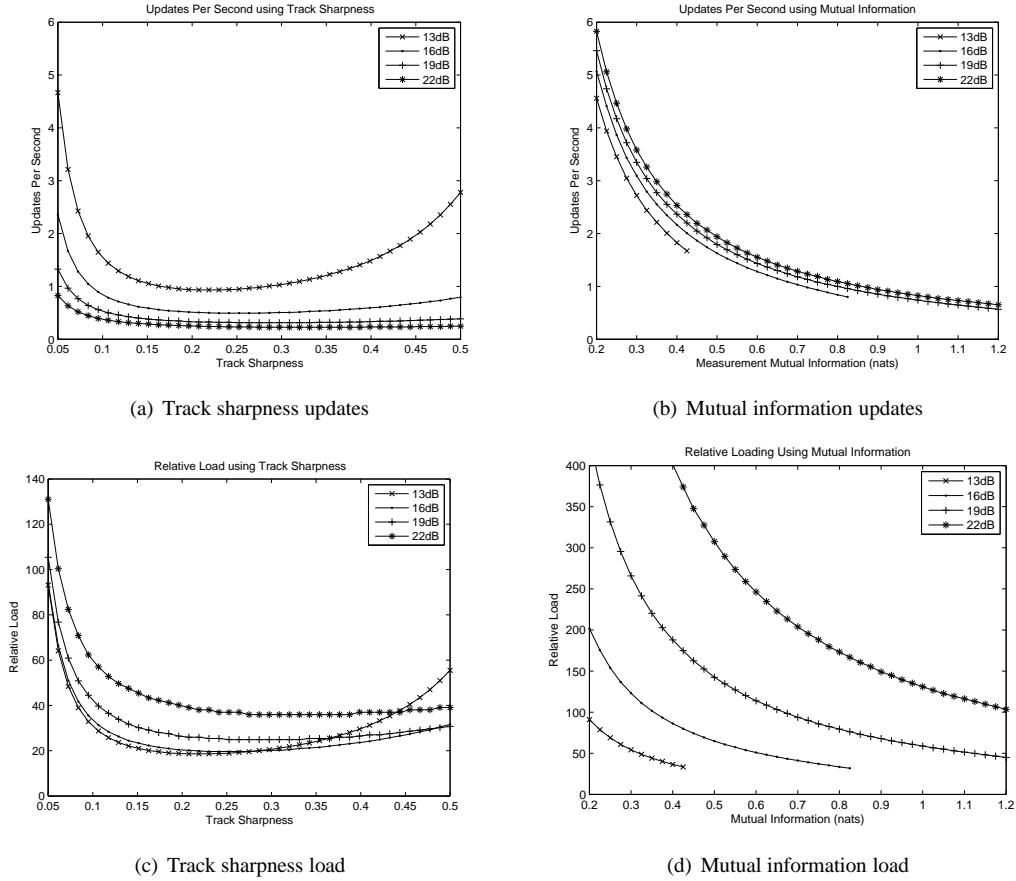
Figure 4.17: Loading for mutual information and track sharpness

is not what this approach optimises. However, it can be seen that the mutual information measure has a logarithmic relationship between the bound and the information rate, highlighting that is what the method optimises. Also, higher SNR produces a higher information rate as higher SNR measurements carry more information.

Root Mean Squared Error - The effect of the different optimisation objective between the measures can be seen in the root mean squared error (RMSE). The root mean squared error for track sharpness and mutual information are shown in Fig. 4.19. The figure shows an increase in the RMSE as the bounds are increased, which is due to the revisit interval increasing with the bound. The information rate and RMSE are directly related and as mutual information optimises this quantity there is a subsequent reduction in the RMSE.

It is common to assess track performance solely on RMSE, which given the reduction shown here would allow the conclusion that mutual information is the superior approach. However, this improvement in the RMSE is potentially unnecessary and comes at the cost of increasing the radar loading, which reduces the number of tracks the system can maintain. This highlights a critical point that not only must a variety of metrics be used to analyse the performance of the allocation, but also that the performance

(a) Track sharpness

(b) Mutual information
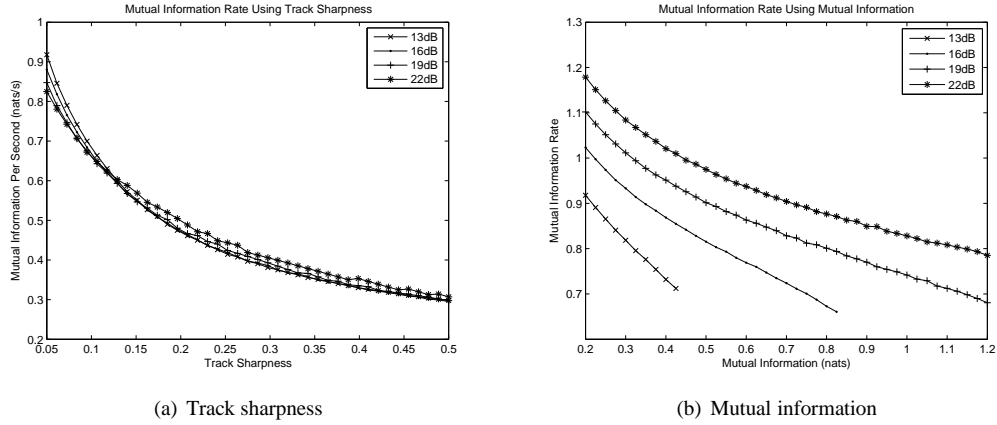
Figure 4.18: Information rate for mutual information and track sharpness

analysis must be in the context of the requirement of the system. In reality, optimisation of the RMSE may be of little concern, while a greater importance is placed on the reduction of track loss, accurate data association and maintenance of the maximum number of targets. It is therefore essential that any mechanism that utilises information theoretic approaches preserves the requirements of the system. For example a target engaged by the weapon system requires high accuracy, whereas a surveillance track may have different requirements, such as track continuity and purity.

(a) Track sharpness

(b) Mutual information
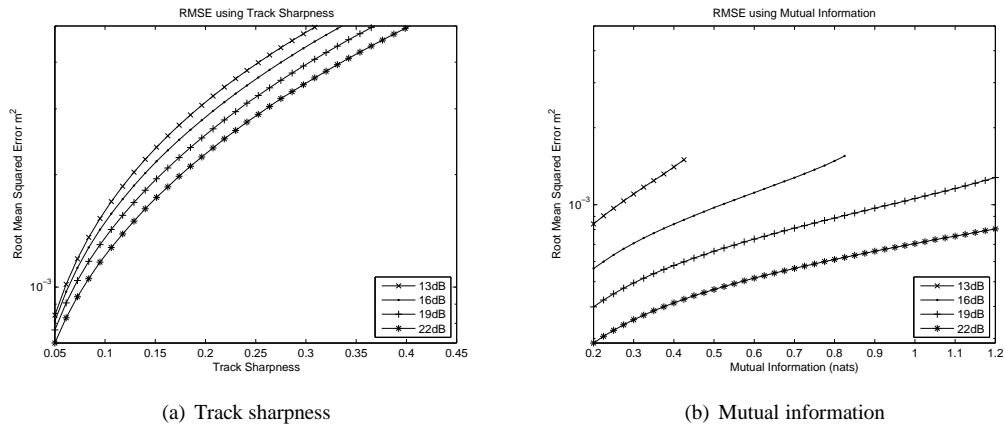
Figure 4.19: RMSE for mutual information and track sharpness

This comparison suggests that there is a subtle difference in the way in which information theoretic measures can be applied. As a surrogate function optimising information improves tracking accuracy and can reduce tracking load and so is relevant to the control of each individual task in isolation. However, in terms of making comparisons and resource allocations between tasks, information is not suitable as optimising information production across all tasks is not the fundamental requirement of the radar. This initial assertion will be explored in the development of the allocation mechanisms in the following chapters.

4.4 Task Utility Functions

The agent based resource allocation mechanism developed in the following chapters requires a single measure to optimise. However, it is clear from the discussion in this chapter that information theory can not fulfil this role and no such single measure exists. In fact, it is found that functionality is improved when resource is allocated based on a variety of measures which adequately describe the potentially complex requirements of each task. A solution to this is to define a mapping from each different quality measure to a common measure called utility.

A utility function u_k can be defined for each task which provides a mapping from quality space $\hat{Q}_k$ to utility:

$$u_k : \hat{Q}_k \to \Re \tag{4.52}$$

This quantifies the satisfaction associated with each point in the tasks relevant performance measure. As the primary quality measure of interest varies between differing radar task types, the utility function provides a single comparable measure.

4.4.1 Linear

A simple utility function is a linear mapping from the relevant performance measure into quality space. For tracking a relevant quality space is the angular estimation error σ_p, and so an example of a utility function is:

$$u_k(\sigma_p) = p_i \begin{cases} 0 & \text{if } \sigma_p > 0.3\theta_B \\ \frac{0.3\theta_B - \sigma_p}{0.15\theta_B} & \text{if } 0.15\theta_B \le \sigma_p \le 0.3\theta_B \\ 1 & \text{if } \sigma_p < 0.15\theta_B \end{cases} \tag{4.53}$$

this utility function is shown in Fig. 4.20. The mapping can be adjusted given the requirements of each task, for example, the accuracy for tracking a target to be engaged can be more accurate than one which is not engaged.

For the long range surveillance function a relevant quality metric is the cumulative detection range and so a mapping similar to Eq. 4.53 can be defined. An example of the utility associated with tracking and surveillance functions is shown in Fig. 4.20.

4.4.2 Logarithmic

It may be more realistic that the satisfaction associated with increases in quality is logarithmic. A similar logarithmic utility function can be defined for other functions such as long range surveillance. An example of the utility associated with tracking and long range surveillance is shown in Fig. 4.20.

The utility function can be a weighted summation of individual tasks measures when more than one is relevant. The choice of utility function can be varied and as it represents the satisfaction associated with each quality metric, complex task requirements can be created.

(a) Tracking function

(b) Surveillance function

Figure 4.20: Example of linear and logarithmic utility functions for tracking and surveillance functions

4.5 Resource Manager Performance Assessment

Performance assessment of multifunction radar resource management algorithms is difficult for a number of reasons. Firstly, it is required to support multiple differing functions and so no single performance metric exists as each function is assessed by differing and disparate measures. Within each function different measures can have different degrees of relevance. For example, track accuracy is important for tracks which require engagement whereas track purity and existence is of more relevance for tracks requiring surveillance. Also, the performance assessment must be in the context of what is currently required from the system, which is likely to change over time. It is desired that the multifunction radar be able to operate in uncertain and dynamic environments which relates to a large variety of possible scenarios. Hence there is no single scenario in which the resource manager should be assessed and if there were, it would not demonstrate how well the resource manager is able to adapt to varying environments. Finally, real data is of limited use as to capture radar data some form of resource management must have already been applied. For a mechanically scanned system this resource management is the mechanical scanning, for an ESA the array face must have been controlled to produce the data. Heavily oversampled data can be of use, but this is rarely available without reducing the realism of the scenario.

As a result of these difficulties the resource management algorithms developed in this thesis are assessed on a variety of appropriate task specific measures. Also in several examples the allocation is assessed in terms of utility as the utility function provides the satisfaction associated with each point in quality space it describes what is required from the system. Maximisation of utility across a system echoes ideas outside of engineering and computer science with direct comparison to Utilitarianism and Jeremy Bentham's *"to achieve the greater good for the most amount of people"*.

4.6   Summary

Measures and models used to guide the allocation of resource critically limit the quality of resulting allocation. A wide variety of task specific measures are relevant, but incomparable between functions which poses a challenge for resource allocation mechanisms.

Information theoretic measures can be applied to the core functions of a multifunction radar to optimise the performance of the tasks. Information theoretic measures as surrogate functions are useful for the optimisation of tasks in isolation, but are not so useful for making higher level allocation decisions as information production is not the aim of the radar system.

As multifunction radar resource management inherently aims to optimise multiple functions, it is desired to use as wide a variety of measures as possible. This chapter has discussed a number of measures, however, these measures are not exhaustive and many more can be considered depending on the requirements of the task. A mechanism which is able to allocate resource to a variety of quality measures is desirable. To convert the differing task specific measures into a single metric for optimisation by an agent system, utility functions can be defined which give the satisfaction associated with each point in task quality space.

Ultimately resource must be allocated against some measure extracted from a model, and how well the measure represents the underlying task can reduce the quality of the allocation. For example, if the target dynamic noise model is poorly matched to the true target dynamic then the resource allocation will be poor. Also some tasks do not lend themselves to be quantified by measures so easily, such as a long duration target recognition task.

When tracking with the presence of false returns, as with clutter, the minimum loading revisit interval is dependent on the false target density. The minimum loading can be found using mutual information, and the modified Riccati equation can be incorporated into covariance analysis to improve tracking resource allocation.

Chapter 5

Agent Systems in Multifunction Radar Resource Management

Agent systems are self-organising computational societies where the synergy of local interaction between agents generates global desirable behaviour. By mimicking human interaction mechanisms, agent systems can provide rapid and intelligent adaptation in uncertain and dynamic environments. Specifically, auctions and markets are suitable for application to RRM as they have evolved in human societies as effective resource allocation mechanisms.

This chapter introduces agent systems and describes the creation of a test bed suitable for developing agent auction mechanisms. The testbed provides agent functionality and generates radar measurement simulation.

5.1   Agent Systems

Agent systems are comprised of multiple computational elements which are able to socially interact by passing messages between each other. This social interaction produces emergent desirable behaviour. The design of an agent system is composed of two aspects, the behaviour of the individuals agents and the design of the mechanisms through which they interact. This section introduces the concept of an agent and the concept of a multi-agent system in which the agents exist.

5.1.1   Intelligent Agents

The term agent is very general but can be characterised quite abstractly in the following way with the support of various texts [Vidal, 2007; Weiss, 1999; Russell and Norvig, 2009; Wooldridge, 2002]. An agent acts on someones' behalf or represents someones' interests, which may or may not be its own. An agent should be able to sense through sensors, and affect, through affecters, the localised region of the environment in which it exists. It should have control of some internal state, which it uses to store localised information on the environment. It should use this information to perform actions, which further change the environment. An agent should have defined goals, and choose actions which bring
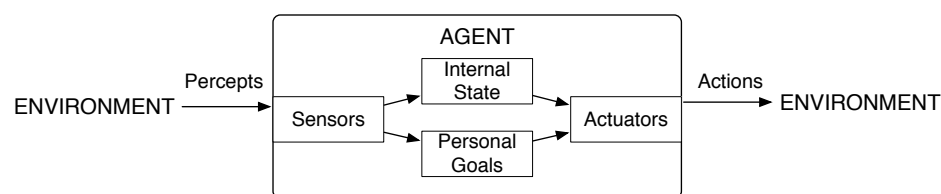
Figure 5.1: Architecture of an agent [Jennings and Wooldridge, 1998]

the agent closer to its goals, given the environmental conditions. This model of an agent is shown in Fig. 5.1.

As this notion of agency covers all types of agents, it is useful to look to a human agent to clarify some of these properties. For example, an estate agent acts on behalf of a property owner wishing to sell a property, however, it is self-motivated as it profits from the sale. The estate agent can sense the state of the market through its interactions with sellers and buyers and may also have access to statistics of the market as a whole. As such, the estate agent has a more complete knowledge of a limited and localised portion of the market, and partial and processed information on the whole market. The estate agent has control over its internal state, as it is able to remember previous interactions which it uses to make decisions. Decisions and subsequent actions complement its goal of selling houses and earning money. Although self-motivated, the estate agent optimises the allocation of property between numerous potential owners which improves the 'social welfare' of the system.

Agent-orientism is the next logical step from object-orientism and so clear differences can be identified. Objects can only be invoked whereas an agent has autonomy over its choice of actions which are invoked subject to the agents goal-directed behaviour. For example a light switch, which is an object, executes a function which is invoked by the operator to control the lights. However, if the light switch were an agent then it would only sense the operators' desire for the lights to be on. Given this it would determine how this sense aligns with its only personal goals; if the goal of the switch agent were to please the user, then it would activate the light and further its personal goal. However, if the goal of the switch agent were to annoy the user then it would not activate the light thereby also furthering its personal goal. This demonstrates how an agent can choose to perform actions, whereas an object is only capable of being invoked. This idea of autonomy allows for agents to interact in ways that an object cannot, creating complex system behaviours.

In addition to the characteristic of autonomy, an agent should possess some of the following key characteristics:

- *Reactive/Adaptive* - An agent must be aware of the environment in which it exists and react to changes in the environment so that its goals continue to be met. This may cause it to adapt its behaviour due to environmental changes, which it should do in a timely fashion.

- *Proactive/Goal directed behaviour* - An agent should formulate goals and exhibit goal-directed behaviour to meet these goals. This means it should not solely react to the environment but be proactive and possibly take the initiative so that its goals are met.

- *Autonomy/Rationality* - An agent should be able to make independent decisions and demonstrate independent behaviour which requires an independent thread of control. An agent should be rational and as such would not perform an action which conflicts with one of its goals or would be likely to leave it in a worse position than before the action.

- *Social Ability* - Agents cohabit environments with many other agents, which can have conflicting or common goals. Therefore an agent must have the ability to socialise with other agents in the environment so as to resolve conflicts or coordinate in a context specific way. Agents can participate in various social mechanisms such as negotiations, auctions, institutions and coalitions.

An agent should be able to demonstrate these key characteristics but the level at which they demonstrate the characteristics can vary widely. This is why the term agent is so general and consequently agents can vary widely. However, at the core of every agent is the concept of autonomy and the resulting independent thread of control.

For completeness, the notion of an agent described here has been very abstract. For the rest of this work and also for agent research in general it is assumed that the agent in question is a computational construct. This construct has its own thread of control, computational ability and some memory. Also, although most agents need to be justified against these characteristics, agents tend to be quite simple and do not need to develop complex individual behaviours or complex demonstrations of these characteristics. The emphasis is on the system behaviour in which the agent is involved.

5.1.2 Multi-Agent Systems

A Multi-Agent System (MAS) is a collection of agents that engage in social interaction. The transition to a MAS is inevitable as there is little that a single agent can achieve. As each agent has a limited, localised knowledge, the data in the system is decentralised and as each agent contributes to the behaviour of the system, the control is distributed. As agents are reactive, rational and autonomous, they can generate desirable system behaviour in environments that are dynamic and uncertain. This visualisation of a multi-agent system is demonstrated in Fig. 5.2.

In classical artificial intelligence the aim is to make one computational construct highly intelligent. Multi-agent systems, however, aim to create intelligent behaviour through the synergy of the system which is generated through relatively un-intelligent interactions between agents. This means that each agent does not need the high levels of intelligence present in classical AI and instead replace this intelligence with the lesser notion of autonomy.
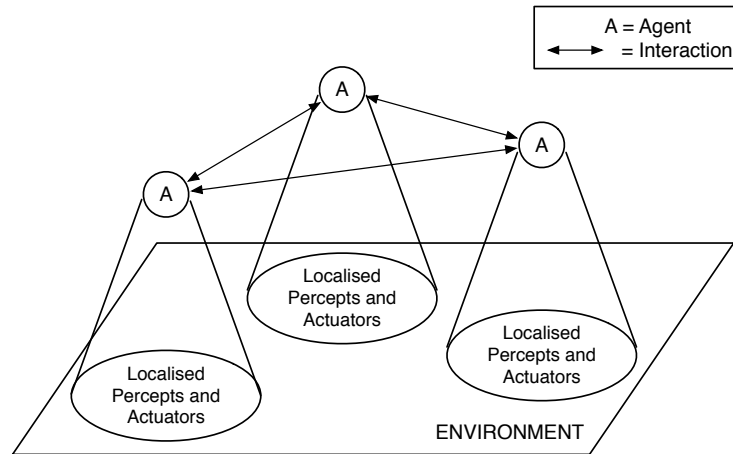
Figure 5.2: A Multi-Agent System (MAS) [Wooldridge, 2002]

Humans frequently engage in multi-agent mechanisms on a regular basis. For example, a road network is a multi-agent space and time multi-constraint optimiser. With a little thought it can be seen that in a road network each vehicle user aims to utilise the shared resource in the best way to meet its goals. If all vehicle users using the road network achieve their goals the problem is optimised and a global 'social welfare' is achieved. This process is very robust against the varying constraints and uncertainty concerning environmental conditions and operability of the network elements. The design of the road network, or mechanism, affects the quality of the global solution, which demonstrates the need for effective mechanism design which governs the agent interaction.

Suitable mechanisms to use in multi agent systems are mechanisms that require low levels of intelligence for each agent. For example auction mechanisms are very efficient at solving resource allocation problems but only require the formation of a valuation and a strategy. Multi-agent systems research covers a wide variety of interactions including trust and reputation, coalitions, institutions, electronic markets, communication and learning. Agents and multi-agent systems combine topics from artificial intelligence, concurrent systems, economics, game theory and social science.

5.2 Mechanism Design

Mechanism design addresses the design of the mechanism through which the individual agents are able interact. The design of the mechanism ultimately determines the behaviour which the agents are required to generate. Specifically, the mechanism should be designed so that it produces the desired outcome based on any preference profile supplied by the agents. Directly relevant to resource allocation problems are auction and market mechanisms, in this case the preferences profiles take the form of valuations of the resource.

5.2.1 Auction Mechanisms

There are a wide variety of auctions which exhibit differing characteristics and are governed by differing protocols. Characteristics include the number of sellers and the number of buyers, whether the auction is for a single good or a combination of goods, whether bids are open or sealed, whether the prices ascend or descend or whether single or multiple units are auctioned.

Common examples of auction mechanisms implemented in human societies are:

- *English Auction* - In the English Auction the bidding price starts at some low value which is incremented after each bidder indicates consent to meet the required bid. The bids are open and the dominant strategy is for an agent to bid without exceeding its private valuation. However, it is worth noting that as the bids are open, information about other agents valuations can be speculated from their bidding habits.

- *Dutch Auction* - The Dutch Auction is similar to the English Auction but uses descending increments. The auction starts with a high price which is lowered until one of the agents indicates it accepts the bid price, hence winning the auction. There is no dominant strategy which can lead to inefficiencies in the allocation.

- *First-Price Sealed-Bid Auction* - Each agent privately submits a bid without knowledge of the other agents' bids. After the bid duration has elapsed the auction clears declaring the highest bidder as the winner. There is no dominant strategy and the agent forms its bid on the basis of any available prior knowledge of the item and the other bidders.

- *Vickery Auction* - The Vickery auction is also known as a second price sealed bid auction. As with the first price sealed bid auction the bids are private, however, it differs as the winning agent placing the highest bid is only required to pay the second highest bid price. This mechanism leads to the desirable dominant strategy whereby each agent bids its true valuation. Hence the auction is *'incentive compatible'* due to the dominant strategy of truth revelation.

To demonstrate the incentive compatibility of the Vickery auction consider the following:

Agent One (t_1) has valuation p_1^* and bids b_1 and Agent Two (t_2) bids b_2

1. If $b_1 < p_1^*$: t_1 risks unnecessarily losing the auction.

2. If $b_1 > p_1^*$ and $b_2 < p_1^*$: t_1 wins and pays less than its valuation.

3. If $b_1 > p_1^*$ and $b_2 > p_1^*$: t_1 wins but pays more than its valuation.

4. If $b_1 = p_1^*$: t_1 never pays more than valuation and maximizes its chance of winning.

In case one t_1 is lowering its chance of winning the auction unnecessarily. In case two t_1 wins and pays less than its valuation. However, it has not benefitted from bidding higher than its valuation and has risked case three occurring. In case three, by both t_1 and t_2 bidding higher than t_1's valuation, t_1 ends up paying more than its valuation which is undesirable. In fact by bidding more than its valuation the only extra auctions t_1 will win are those in which it ends up paying more than its valuation.

This concept of incentive compatibility is important because it means that the agent does not need to model any of the game theoretics about the other agents in the auction. Also if each agent truthfully bids then the good is allocated to the agent who truly values it the highest and global social welfare is maintained.

5.2.2 Social Choice Theory

Social choice theory focuses on the part of the mechanism that transforms a set of preferences from agents into an outcome or outcomes. A social choice function chooses a single outcome and a social choice correspondence chooses a set of outcomes given the preference profiles of participating agents.

Voting mechanisms are typical examples of the application of social choice functions. For example, in a mechanism where agents can vote for one of two candidates, the social choice function transforms the preferences into an outcome which involves counting the votes and declaring the highest polling candidate as the winner which maximises social welfare. However, this simple problem is complicated where there are more than two candidates and agents are permitted to submit preferences over all candidates.

In an auction the auctioneer receives some bids, which represent the agent's preferences, and it must apply a social choice function to transfer these preference profiles into an outcome. In most simple auctions the social choice function is trivial as it is assumed that the auctioneer chooses the bid which maximises the auctioneers potential income. If agent t_k belongs to set of agents T_A and outcome o belongs to the set of possible outcomes O then the social choice function is:

$$f(T_A) = \arg\max_{o \in O} \sum_{t_k \in T_A} u_{t_k}(o) \tag{5.1}$$

where $u_{t_k}(o)$ is the utility production associated with outcome o for agent t_k.

However, just like in voting when more complicated mechanisms are used the social choice function is no longer trivial. Such as mechanisms which involve multiple units, combinations of units and multiple preferences, are developed in this thesis.

5.3 Java Agent Development Framework

The Java Agent Development (JADE) Framework is a software framework which was developed by Telecomm Italia to extend the Java platform enabling development of multi-agent systems. The frame-

work provides the key characteristics required for agency, including concurrency, social ability and behaviours. JADE was developed to be fully compliant with the FIPA agent standards. The Foundation for Intelligent Physical Agents (FIPA) is an IEEE Computer Society standards organisation which promotes the use and interoperation of agent based technology. FIPA specifies software standards for all aspects of agency, including ontologies, communication and platform structure.

The JADE platform consists of a number of containers, in which agents can exist, that can be distributed over several hosts each running one Java application. Agents exist within a container, with each agent ran in an independent thread. In compliance with FIPA, the agent platform is modelled as containing the agents, an agent management system, a directory facilitator and a message transport system as shown in Fig. 5.3. Only a single agent management system exists on a platform, which handles tasks such as agent creation and individual naming. The directory facilitator provides a yellow page service for agents to advertise services and the message transport system handles the passing of messages between agents in the platform, which could be across multiple hosts.
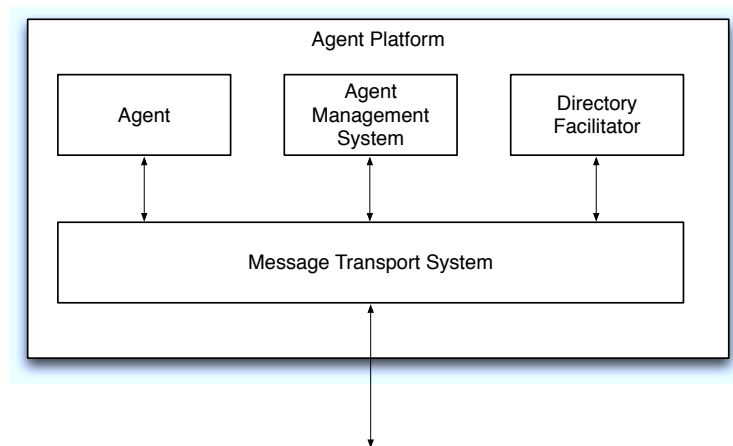
Figure 5.3: Agent platform defined by FIPA

Key functionality provided by JADE is the ability for agents to implement multiple, concurrent behaviours and pass messages. Message objects can be created which contain a set of attributes. These attributes are the sender and receiver ID, the conversation ID, the content of the message and a message performative such as 'request' or 'inform'. The message transport system handles the delivery of messages to the correct agent, and each agent possess a queue of active messages which it can process in order or messages can be extracted from the queue according to specific attributes.

As each agent must be able to implement numerous concurrent behaviours, JADE allows for behaviour subclasses to be defined which can be added or removed from the agent at any time. The scheduling of active behaviours is hidden from the programmer and executed in a round robin sequence. Behaviours can be blocked and await triggering by an event, such as a received message, which prevents

unnecessary CPU utilisation. Various behaviours can be chosen from, such as oneShotBehaviour which executes just once, cyclicBehaviour which executes repeatedly and is useful for processing received messages, wakerBehaviour which executes after a time duration has elapsed or TickerBehaviour which repeatedly executes at a specified interval.

Finally, JADE implements interaction protocols as defined by FIPA which describe the format that different agent interactions or conversations should adhere too. The interaction protocols ensure that all parties involved in the interaction are aware of the current stage of the interaction, which removes the uncertainty associated with agents not responding to any part of an interaction chain. Typical interaction protocols are FIPA-Request, FIPA-Query and FIPA-Recruit.

The JADE Framework extends the Java platform to provide agent functionality which allows development of agent systems. The JADE Framework is used in this work to develop auction mechanisms which are applied to the radar resource management problem.

5.4 Agent Based Resource Management Testbed

An agent based radar resource management testbed has been created in Java using the JADE Framework to enable the development and simulation of the auction mechanisms developed in the following chapters. This involved both creating a testbed environment suitable for developing auction mechanisms, as well as the simulation of radar measurement data using the models from Chapter 4. This section describes the design of the system, including the architecture and the agents and objects used in the software.

5.4.1 System Architecture

The architecture of the complete software system is shown in Fig. 5.4. The architecture was designed to allow integration into the radar resource manager architecture shown in Fig. 3.3. The inheritance structure of the agents allowed for different auction mechanisms to be implemented whilst maximising code reuse. It was also important to design each task agent without the knowledge of the task it represents, again to maximise code resuse.

The software contains a collection of agents and objects, whereby agents extend the notion of an object by possessing autonomy, the ability to program goal-directed behaviours, a social capability and an independent thread of control. The main section of the software, which is the agent based resource manager, contains a number of agents representing radar tasks that are able to participate in varying auction mechanisms. This agent based resource manager generates tasks requests which are passed to a scheduler to be formed into a radar timeline. Upon execution the environment is modelled and simulated measurement data is returned to the task agents to update their respective function. Typical functions implemented in this resource manager are surveillance and tracking functions. The task agents and func-

tions have access to a waveform database, priority assignment and environmental model. The waveform database contains the allowed waveforms which can be selected for transmission, the priority module contains a list of predefined priorities for each task and the environmental model contains environmental information for each task such as estimated target radar cross section or predicted false target density. The TaskAgents, AuctioneerAgents and SchedulerAgents are discussed in Sec. 5.4.2.1, Sec. 5.4.2.2 and Sec. 5.4.2.3 respectively and the Function objects are discussed in Sec. 5.4.3.1.

5.4.2 Agents

Within the software testbed a number of agents have been designed which adhere to the notion of agency. It is required to have agents which take on the role of the auctioneer in a specific auction, as well as agents acting as the auction participants who represent differing tasks. In addition there are agents who facilitate the running of the software.

5.4.2.1 Task Agents

Task agents can represent any of the radar functions from Sec. 2.3.2, such as surveillance or tracking. These agents possess a function object, which could be a tracking or surveillance task, without knowledge of the function it performs. The function object is detailed further in Sec. 5.4.3.1.

The inheritance structure for the task agents is shown in Fig. 5.5. The general *TaskAgent* class inherits agent functionality from the base *Agent* class provided by the JADE Framework. Different agent classes inherit from the general *TaskAgent* class, these agents posses differing behaviours which depend on the auction mechanism within which they are required to operate.

In the inheritance tree shown in Fig. 5.5 three agents are used for differing auction protocols. CDAAgent refers to an agent which engages in a continuous double auction, RBAgent to an agent which selects parameters according to predefined rules and FPAgent which engages in a first price sealed bid auction mechanism. These specific behaviours for the CDAAgent and RBAgent are discussed in the development in Chapters 6 and 7. The behaviours implemented in the *TaskAgent* class are:

- *startNotification* - Responds to a notification of the simulation starting by initialising the agent.

- *receiveSchedNotif* - Responds to a notification of the agent's task being executed by the scheduler by updating the agent and the task given the received data. Abstract class *onSched* is called which enables inheriting agents to execute additional behaviour.

- *ticker* - Monitors the passing of time which enables the recording of data and termination after simulation is complete.

Additionally the methods *updateScheduler* and *takeDown* are defined which update the scheduler with the current operating parameter selection and terminates the agent respectively.

Agent Based Resource Manager

Waveform
Database

AuctioneerAgent

OrderBook

Priority

TaskAgent

Function

....

TaskAgent

Function

Environmental
Model

Scheduler

Tx/Environment/
Rx
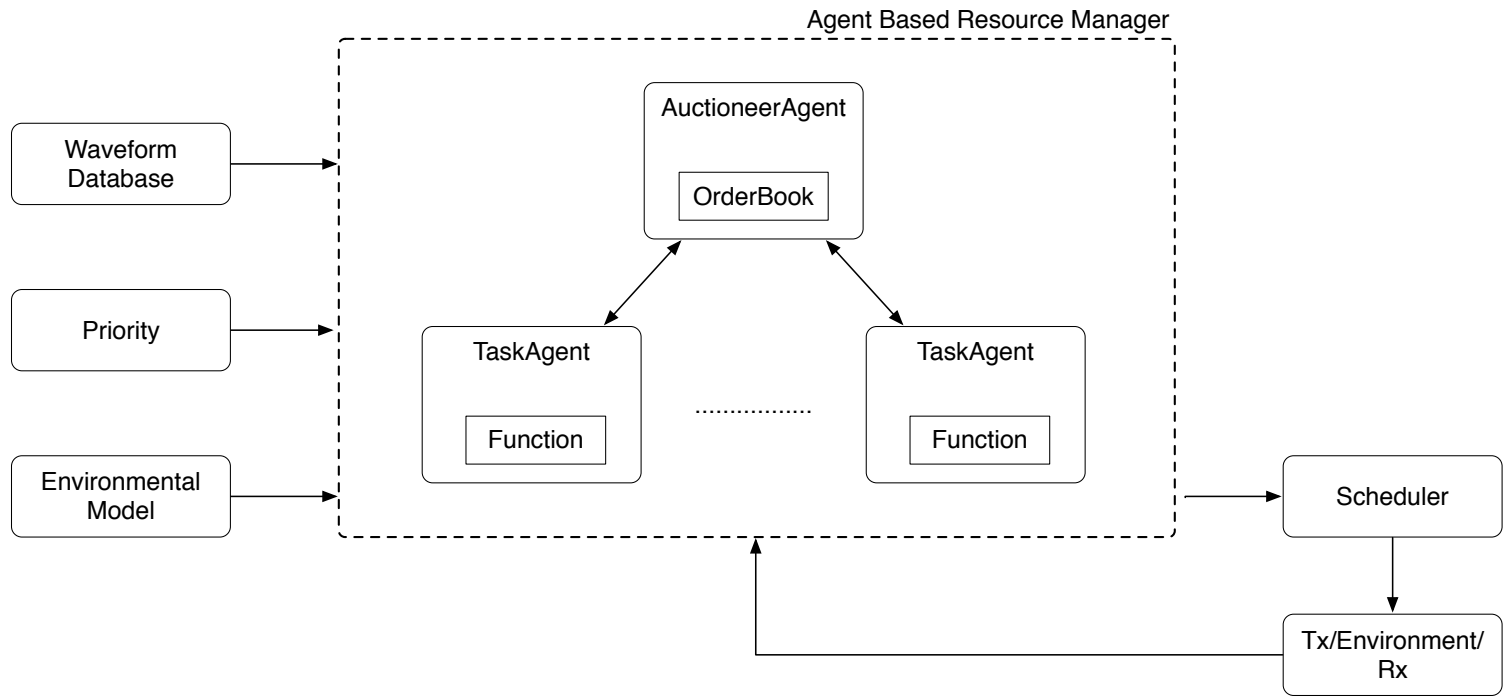
Figure 5.4: Agent based resource management architecture

Agent

|

TaskAgent

CDAAgent   RBAgent   FPAgent
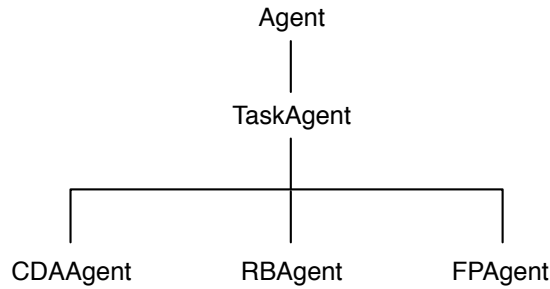
Figure 5.5: Inheritance structure for task agents

This design of the task agents was chosen to maximise code re-use whilst allowing agents representing new auction protocols to be rapidly included. Code re-use was ensured by allowing the task agent to represent any task, without any code specific to a certain radar function. Ability to rapidly extend was ensured by including common functionality in the *TaskAgent* class.

5.4.2.2   Auctioneer Agents

A variety of auctioneer agents can be selected to implement different auction social choice functions. The selected auctioneer agent organises the market place in which the numerous task agents, representing radar tasks, engage. Typical activities for the auctioneer agent is to facilitate trades between agents, as with a continuous double auction, or to declare the winner from a set of bids, as with the first price sealed bid auction. The inheritance structure for the auctioneer agents is shown in Fig. 5.6. *CDAAgent*, *RBAgent* and *FPAgent* correspond to continuous double auction, rule based and first price agent respectively; each of these agents inherit from the agent class and have differing behaviours which depend on the auction mechanism in which they engage, which are discussed further in Chapter 6 and Chapter 7. Auctioneer agents make use of the orderbook object which is described below.

Agent

|

AuctioneerCDA   AuctioneerRB   AuctioneerFP
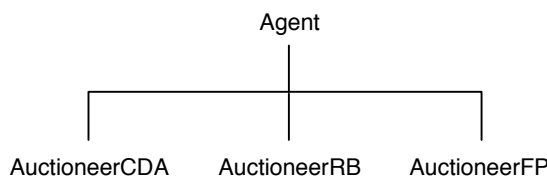
Figure 5.6: Inheritance structure for auctioneer agents

There is not much common functionality between auctioneer agents and so no common class is defined in the inheritance structure. Different auctioneer agents can be rapidly added to the software using the functionality provided by the JADE Framework.

5.4.2.3 Scheduler Agents

A variety of schedulers have been implemented to form radar timelines from sets of tasks requests which are sent from the task agents. The schedulers have been implemented as agents, primarily to exploit the ability to receive messages from the task agents. Each scheduler agent is a variant of the earliest deadline first scheduler which is selected for the desirable characteristics outlined in Sec. 3.4 and by Miranda et al. [2007a]. The inheritance structure for the functions is shown in Fig. 5.7.

Agent

|

EDF

|

HpEdf   EdfTws   HpEdfTws
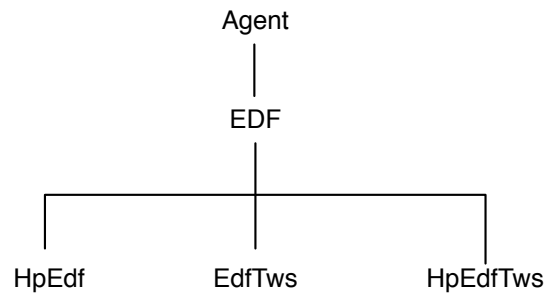
Figure 5.7: Inheritance structure for schedulers

Edf is the an implementation of a standard earliest deadline first scheduler. *HpEdf* includes respect to priority which ensures that no task is delayed by a lower priority task, which is similar to the Butler scheduler Butler [1998]. Finally *EdfTws* and *HpEdfTws* incorporates an additional track-while-scan scheduling mechanism which is used in the tracking control simulations. Again, this inheritance structure is chosen to maximise code reuse.

5.4.2.4 Auxiliary Agents

Additional agents have been implemented which aid the operation of the testbed either by creating and organising all the agents depending on the required simulation, or collecting the simulation data from agents in the system. They are not shown in Fig. 5.4 as they are not directly relevant to the operation of the agent based radar resource manager. They are:

- *DataMan* - The data manager collects the data from the individual agents when the simulation is complete and writes the data into Matlab data files for analysis.

- *MarketMan* - The market manager takes input from the simulation user, initialises and starts the appropriate simulation, which involves the creation of all the appropriate agents and ensures that they are synchronised.

These agents implement various behaviours but do not contribute to the novelty of this work and so are not detailed further.

5.4.3 Objects

There are a number of objects which can be possessed or used by the agents, which adhere to traditional object-orientated programming and so have some encapsulated state and invokable methods. Most importantly, individual radar tasks are implemented as objects, which inherit from the Function class, but there are also auxiliary functions which are used throughout the resource manager design.

5.4.3.1 Functions

Each task agent possesses an individual task which is instantiated from a function object. The agent has no knowledge of the specifics of the task, which enables a single agent design to represent all functions requiring support by the MFR. The inheritance structure for the function objects is shown in Fig. 5.8. Examples of functions are shown in the tree, *LRSurv*, *AccTrack* and *ExTrack* being long range surveillance, accuracy track and existence track respectively. Adding additional MFR functions is performed by adding a subclass which inherits from the function class which allows rapid development.
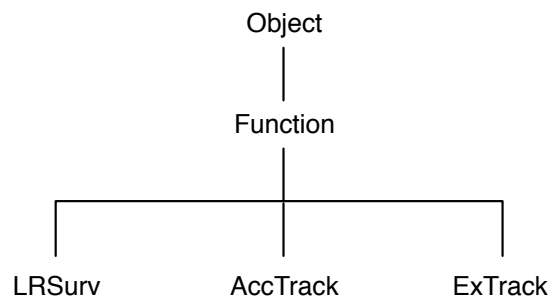
Figure 5.8: Inheritance structure for functions

The abstract function class requires subclasses to override a number of methods which varying depending on each function:

- *evalOpPoint* - Evaluates the utility associated with the passed operating parameters.

- *updateStatus* - Update the status of the task, used to check if a task is violating constraints.

- *receiveData* - Process data which results from an executed radar look.

In addition the *logResults* and *getResults* methods handle the collection of simulation data.

The function object also contains two methods which are invoked by the task agents to form bids:

- *hillClimb* - Perform local hill climbing search up or down.

- *findBestOffer* - Use hill climbing search to produce best offer.

which both rely on subclasses overriden version of the *evalOpPoint* method.

5.4.3.2 Auxiliary Objects

There are also a number of auxiliary objects which are used by the various agents to achieve their goals:

- *Look* - A *Look* object is a request which is passed to the scheduler. The look object contains requested task parameters as well as earliest, latest task execution times and the task priority.

- *Offer* - An *Offer* object is a request to trade resource which is submitted to the auctioneer. The offer contains a unit price and quantity and identifier.

- *OpDims* - Agents are capable of generating *OpDims* objects which describe a parameter selection.

- *OrderBook* - The *OrderBook* is an object which the auctioneer agents possess to organise the offers which are submitted by task agents.

These objects do not contribute to the novelty of this work and so are not discussed further.

5.5 Summary

Agent systems are collections of multiple agents who possess autonomy, a social ability, an independent thread of control and exhibit goal directed behaviour that generates desirable emergent behaviour through relatively simple interactions. Agent systems typically mimic human interaction mechanisms, of which auction mechanisms are especially relevant for resource management and allocation problems.

This chapter has detailed the architecture of an agent based radar resource manager which exploits the use of a mixture of objects and agents with functionality provided by the JADE Framework. Simulated measurement data is also generated using the theory and models from Chapter 2 and Chapter 4 respectively. The architecture was designed for the rapid addition of extra functions and for maximum code re-use. As an agent based radar resource manager it is the first of its kind and is used to produce the novel results for the auction mechanisms studied in the following chapters.

Chapter 6

Sequential First Price Auction

A first price auction is a mechanism which allows a central auctioneer to distribute a resource or commodity to the highest bidder drawn from a set of sealed bids. A sequential first price auction is a series of auctions which allows the auctioneer to distribute multiple units in succession. As auctioneers are only able to sell resource and participants are only able to purchase resource, the trades have one direction and so the auction is known as one-sided. The first-price auction has been implemented in various real world situations as well as seeing recent interest for optimisation problems [Rogers et al., 2007; Payne et al., 2006].

This chapter introduces the sequential first price auction mechanism and describes an application of the mechanism to multifunction radar resource management which results in the sequential first price auction resource management (SFPARM) algorithm. The resulting algorithm is analysed through simulated tracking control problems and similarities are drawn with existing best first schedulers and POMDP approaches.

6.1   Sequential First Price Auction Mechanism

In a first price auction each participant submits a single bid which is private and so not publicly available to other participants. This differs from the traditional English Auction where participants compete by revealing bids, however, it has been shown through the Revenue Equivalence principle [Vickrey, 1961] that both auctions generate equal profit. The single shot nature of the auction and the lack of price revelation reduces the communication and computation associated with the mechanism, which is a desirable characteristic for a radar resource manager which has a stringent real time requirement.

A first price auction is comprised of two periods, the trading and clearing periods. The trading period is the time over which bids can be submitted to the auction and the clearing period is the time required for the auctioneer to announce the winner or winners. A sequential first price auction mechanism is a series of first price auctions and so consists of a number of trading and clearing periods. A first price auction can distribute single or multiple units during one cycle of the auction. If a single unit is auctioned

then it is awarded to the highest bidder, if multiple units are auctioned then they are awarded to the set of highest bidders. All winning bidders are required to honour and pay the value of their submitted bid.

Implementations of sequential first price auctions incorporate the following concepts:

- *Bid* - A request to purchase resource. In a single unit auction this is simply a unit price, however, for a multi-unit auction this could be a preference profile over a range of resource quantities.

- *Trading Period* - The time over which bids are accepted to the auction.

- *Orderbook* - A collection of the best bid prices which is maintained by the auctioneer and wiped when the auction clears.

- *Bid Price Limit* - The minimum bid price allowed for admission to the auction.

The primary variations within a sequential first price auction are the trading period and the quantity of resource auctioned in each auction cycle. The choice of trading period depends on the application, it must be long enough for bids to be evaluated and collected, but short enough so that the multiple resource units are sold at an appropriate rate. The auctioned quantity depends on the nature of the resource; sequential auctions of single units places a greater computational demand on the participants, however, auctioning multiple units per trading period can reduce mechanism efficiency. The social choice function which translates the bids received during the trading period into an outcome is the same for all first price auctions and declares the maximum of the received bids as the winner. However, the social choice function can vary in the meaning of the application specific auction currency.

6.2 Sequential First Price Resource Management Algorithm

This section describes the Sequential First Price Auction Resource Management (SFPARM) algorithm which is the result of the application of the sequential first price auction mechanism from Sec. 6.1 to the radar resource management problem. The SFPARM algorithm allows a central auctioneer to distribute radar time between numerous competing tasks and so creates a schedulable radar time-line.

6.2.1 Mechanism

The SFPARM algorithm hosts a market where an auctioneer sequentially distributes radar time between agents representing competing radar tasks, such as tasks generated by the surveillance and tracking functions. The resource auctioned in the algorithm represents radar time, which is the next access to the radar. This access, or dwell, can be any desired length of time to accommodate the requirement of differing task dwell lengths. After collecting the bids the auctioneer declares the highest bidder the winner who is granted the next access. Once the winning radar task is finished, the next round of the auction occurs, to allocate the next access. This mechanism can be thought of as encompassing both the task request generation and also the scheduler, shown in the radar resource manager architecture in Fig.

3.3, as task parameter selection emerges from the agents and the auctioneer forms a radar timeline. The algorithm has been implemented for analysis through simulation using the radar resource management testbed from Sec. 5.4.

The SFPARM algorithm consists of the following auction cycle:

1. At the start of the auction cycle there exists a set a m task agents $T_A = \{t_1, ..., t_m\}$ which represent numerous tasks performing radar functions.

2. The auctioneer announces the start of the next trading period, which is equal in duration to the previously allocated dwell.

3. Each agent submits a private bid which comprises of a quantity q, unit price p and identifier. So a bid from agent n has the form $b_n(q_n, p_n, n)$.

4. By the end of the trading period the auctioneer has collected a set of active bids $B_A = \{b_1, ..., b_n\}$.

5. The auctioneer declares the highest bidder the winner which is scheduled next.

6. The auctioneer announces the start of the next trading period and the cycle repeats.

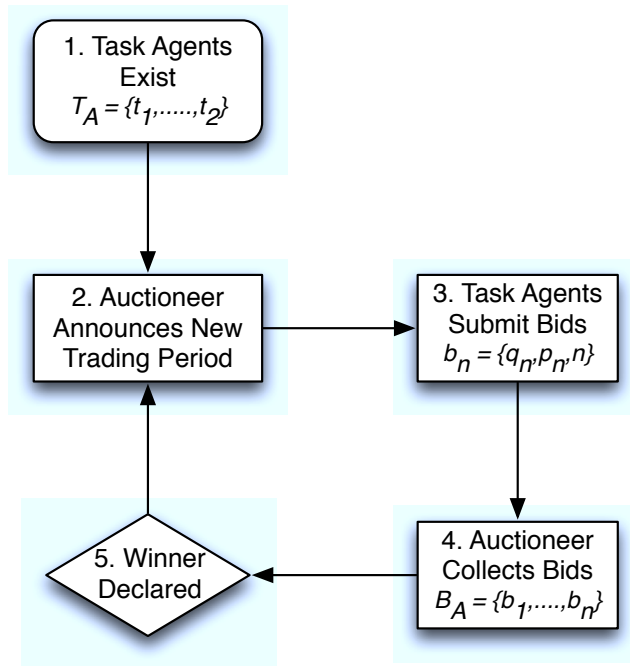This auction cycle is shown in Fig. 6.1.

Figure 6.1: Sequential first price auction cycle

In the following analysis one auction is held per access to the radar, which is a radar dwell of any length required by the task. Auctioning a single dwell in one auction cycle is more computationally demanding than several dwells, however, it is studied first as it will perform the best in terms of mechanism

efficiency in a dynamic scenario. The computational demand could be reduced by lengthening the radar time allocated in one auction cycle; a sensible value being approximately $0.1s$ which is a typical minimum revisit interval. A longer auction cycle duration would significantly reduce the radar performance against pop-up targets.

A conventional RM method is to use predefined rules to select task parameters and an earliest deadline (or highest priority) first scheduler. In this conventional approach the actual task parameters or at least the revisit interval is emergent depending on the radar loading, most notably in uncertain environments or under severe resource constraints. SFPARM is similar to this conventional approach, however, it allows for different criteria as the *'best'* task to schedule first which depends on the meaning of the auction currency. So this mechanism is not completely removed from existing best-first schedulers, but provides a framework to explore the criteria for determining the *'best'*.

This mechanism is also similar in form to a Partially Observable Markov Decision Process (POMDP), as sequential decisions are made under uncertainty subject to the reward function defined by the meaning of the auction currency. When representing the multifunction radar resource management problem as a POMDP, the possible actions which transition the discrete time decision process between states can be scheduling any of the tasks currently supported. The representation of SFPARM as a discrete time decision process is shown in Fig. 6.2. The reward associated with each action is equivalent to each agents' bid and is determined by the meaning of the auction currency. This interpretation of the SFPARM algorithm as a discrete time decision process lacks the non-myopic 'lookahead' present in a POMDP and so will undoubtedly perform worse, however, it is adequate for considering the initial feasibility of the application of POMDPs to existing radar control software.
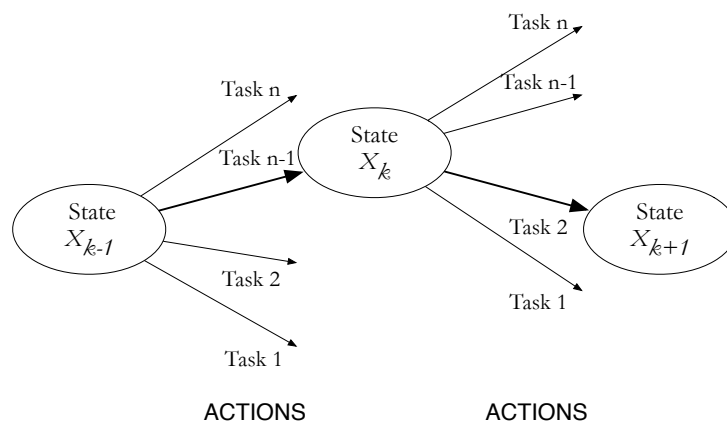

Figure 6.2: Multifunction radar discrete time decision process

Analysis of this mechanism can provide insight on agent systems but also best first schedulers and POMDP approaches.

6.2.2 Task Agents

Each agent in the auction represents a radar task, such as surveillance or tracking, and aims to win as much radar time as possible during the auction, without paying above its valuation for the radar time. The resulting allocation is highly dependent on the meaning of the currency to evaluate and purchase radar time, which must be accurately related to the represented task. This critical meaning of the currency necessitated the research conducted in Chapter 4.

Through the research conducted in Chapter 4, three measures can be taken forward for consideration as the currency used in SFPARM:

- Time difference between current time and desired execution time as specified by rules.

- Mutual information production of next measurement as defined in Sec. 4.2.2.

- Any relevant task quality translated into utility using a utility function as described in Sec. 4.4.

Mutual information and task quality assessed through a utility function are chosen to further analyse the usefulness of these new measures. The time difference, using rules, produces behaviour which is equivalent to an earliest deadline first scheduler and so is similar to the Butler best-first scheduler [Butler, 1998]. It is included as it provides a basis of comparison with which to assess the alternative methods.

Each agent submits a bid when requested by the auctioneer, which is a valuation of the auction currency production the radar task would produce given the next access, which is a variable length dwell. It is assumed that the bidding strategy of each agent is to bid its true valuation for the dwell. This is not realistic in a real sequential first price auction mechanism as the agents would bid lower than their true valuation, depending on their assessment of the competition, in order to maximise their profit. However, truth revelation is used as a starting point and more complex bidding strategies can be developed once the merits of the mechanism have been assessed.

As implemented in the radar resource management testbed, the task agent possesses the behaviours described in Sec. 5.4.2.1. In addition the task agent also possesses the following behaviours:

- *ReceiveBidRequest* - Await announcement of a new trading period from the auctioneer, upon which evaluate and submit a new bid.

- *ReceiveWinNotif* - Await notification from the auctioneer of winning an auction, upon which inform the scheduler to update represented task.

these behaviours allow the task agents to participate in the auction.

6.2.3 Auctioneer Agent

The auctioneer synchronises the auction, collects bids for radar time and applies the auction social choice function. The auctioneer triggers a new auction cycle by requesting a bid from each task agent which

is required in terms of the selected auction currency. The auctioneer agent collects the bids, declares a winner which is scheduled while the auction process repeats.

This auction mechanism can be thought of as a best-first scheduler with the three proposed currencies in Sec. 6.2.2 giving different criterium for 'best':

- *Rule Based Earliest Deadline First (RB-EDF)* - The winning bid is the earliest deadline, which is equivalent to the greatest time between the deadline and current time. This is included as a conventional approach to provide a basis for comparison. RB-EDF implements the following social choice function:

$$f(T_A) = \arg\max_{o \in O} \sum_{t_k \in T_A} d_{t_k}(o) \tag{6.1}$$

  where t_k is an agent from set of agents T_A and $d_{t_k}(o)$ is a function giving the delay from the desired execution time encountered from the outcome o, which is scheduling task agent t_k next.

- *Greatest Mutual Information First (GIF)* - The winning agent will produce the greatest mutual information gain from the next radar access. This myopically maximises information production. GIF implements the following social choice function:

$$f(T_A) = \arg\max_{o \in O} \sum_{t_k \in T_A} i_{t_k}(o) \tag{6.2}$$

  where t_k is an agent from set of agents T_A and $i_{t_k}(o)$ is a function giving the mutual information production of the next measurement from the outcome o, which is scheduling task agent t_k next.

- *Lowest Quality First (LQF)* - The winning agent has the lowest quality in terms of utility. LQF implements the following social choice function:

$$f(T_A) = \arg\max_{o \in O} \sum_{t_k \in T_A} u_{t_k}(o) \tag{6.3}$$

  where t_k is an agent from set of agents T_A and $u_{t_k}(o)$ is a utility function.

In the following analysis the three variants of SFPARM are referred to as their best first scheduler equivalents.

As implemented in the radar resource management testbed, the task agent possesses the behaviours described in Sec. 5.4.2.2. In addition the task agent also possesses the following behaviours:

- *CollectBids* - Collect incoming bids from task agents and add bids to the orderbook. When trading period expires or all expected bids are received, notify the winner of the auction.

- *RestartAuction* - Once a winner is notified, clear the orderbook and announce the start of a new trading period.

these behaviours allow the auctioneer to interact with the task agents and execute the auction cycle.

6.3 Simulation Analysis

This section analyses the SFPARM algorithm presented in the preceding section in the context of single and multi-target tracking control problems.

6.3.1 Single Target

The SFPRAM algorithm allocates a radar dwell to the highest bidder, which has been taken in terms of time delay (RB-EDF), mutual information production (GIF) or utility relating to task quality (LQF). The allocation of resource based on these measures was discussed in Chapter 4. This section analyses the effect of using these measures on a target track given that competition for resource is resolved using SFPRAM.

It is first useful to visualise how the valuations submitted to the auction by a task agent representing a track, which determines which task is executed and hence the revisit interval, differ over time for the three variants considered. This is shown in Fig. 6.3(a), Fig. 6.3(b) and Fig. 6.3(c) for RB-EDF, GIF and LQF respectively. In this simulation a single target is tracked using a continuous white noise jerk model as the limiting form of the Singer model with a process noise intensity $\tilde{q} = 3.3$ unless otherwise stated. The target is on the radar boresight at $50km$, a unity probability of detection is assumed with a received SNR of $22dB$ unless otherwise stated and $1°$ beamwidth. The track is initiated using five dwells separated by $1s$ and the track is assumed to be in thermal noise with false alarm probability 10^{-4} and no clutter. The target dynamic is assumed matched to the tracker model, under this assumption the estimation error covariance correctly describes the uncertainty in the track and the RMSE is the trace of the estimation error covariance matrix. Therefore, only covariance matrices are propagated without generation and filtering of measurements, as under these assumptions the covariance matrices adequately describe performance. The rule used for RB-EDF is to maintain an angular accuracy below 0.1 of the beamwidth. The utility function used for LQF is a linear mapping from angular estimation error σ_p:

$$u_k(\sigma_p) = p_i \begin{cases} 1 & \text{if } \sigma_p > 0.15\theta_B \\ 1 - \frac{0.15\theta_B - \sigma_p}{0.075\theta_B} & \text{if } 0.075\theta_B \leq \sigma_p \leq 0.15\theta_B \\ 0 & \text{if } \sigma_p < 0.075\theta_B \end{cases} \tag{6.4}$$

note this is reversed from Eq. 4.53, as the measure used in the LQF is *lowest quality first*.

Trivially, the bid value submitted to the auction by a track agent in RB-EDF is equal to the time delay based on the execution time specified by the rules. For GIF, which is shown in Fig. 6.3(b), the bid value, which is the mutual information production of the potential measurement, does not increase linearly and is affected by the received SNR which suggests that when using this measure preference

is given to 'bright' targets which produce greater information. Fig. 6.3(c) shows LQF where it can be seen that the utility production of the next measurement increase is also non-linear and is affected by the target process noise intensity $\tilde{q}$ in the tracking filter model. This suggests that LQF gives a resource preference to targets performing extreme manoeuvres as they degrade in task quality, or angular accuracy the quickest. Therefore, RB-EDF would be preferable for maintaining many targets as few miss their deadlines, GIF would be preferable when all targets are required to be of greatest possible accuracy and LQF would be preferable when the tracks have differing quality requirements.

The preference towards specific environmental parameters in these results are an undesirable characteristic of each mechanism, the ideal preference under resource constraints is towards achieving best quality with minimum resource. So these results give a warning that none of these methods, including the conventional RB-EDF, are directly tackling the radar resource management problem, in the sense of globally maximising task performance given a finite resource.

(a) Rule based earliest deadline first

(b) Greatest information first

(c) Lowest quality first
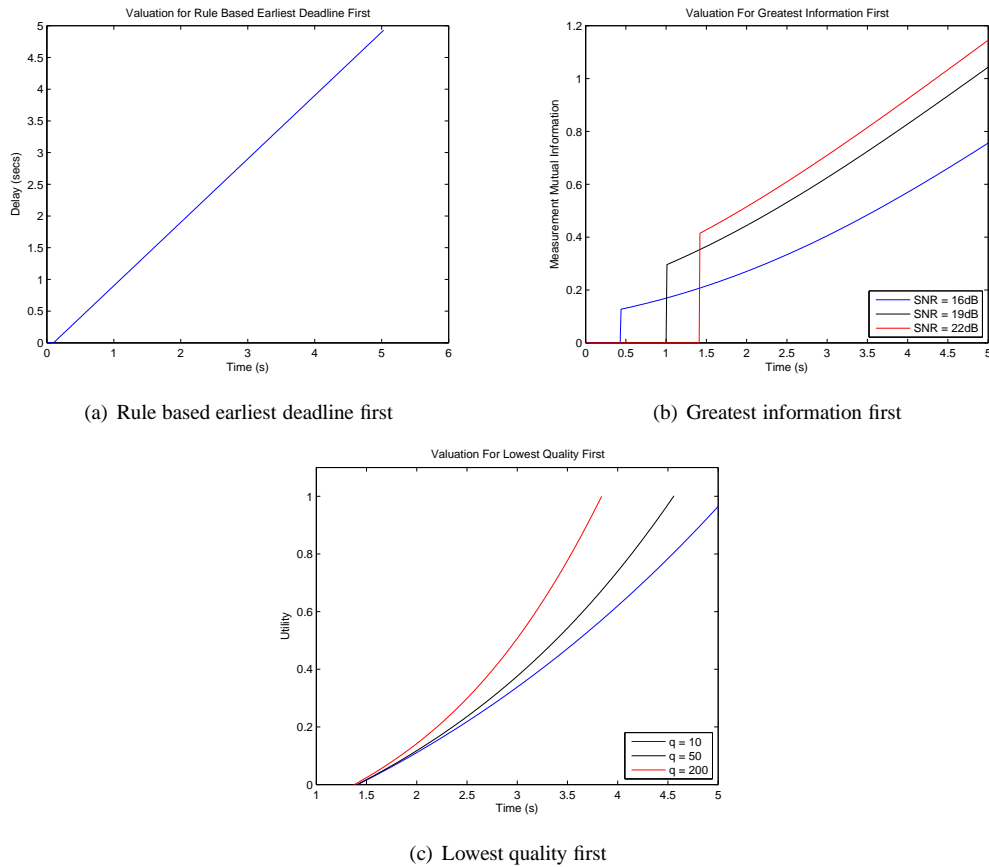
Figure 6.3: Valuations against time for three SFPRAM variants

These results indicate that the allocation produced using GIF and LQF will differ significantly from the conventional RB-EDF, by scheduling based on information regarding mutual information and task quality through utility. However, it has been found that the extent to which this improves or degrades performance can not be analysed from a single track perspective as it is not clear how the competition for

resource in SFPARM will manifest itself into the valuations shown in Fig. 6.3. Therefore to sufficiently analyse and compare the allocation mechanism it is necessary to fully model the competition for resource by simulating numerous competing tracking tasks.

6.3.2 Multiple Targets

This section analyses SFPARM by simulating the competition for finite resource which arises from numerous competing tracking tasks. In the simulation, targets can be tracked using a share of resource dedicated for active updates and it is desired to select parameters to optimise tracking performance subject to the finite resource available. The target environmental parameters considered outside of control are range, azimuth, radar cross-section and manoeuvre standard deviation. Operational parameters under control relate to the waveform selection, which is simplified to the choice of revisit interval and dwell length assuming that lengthening the dwell increases the SNR according to ideal coherent integration. The fixed radar parameters for the simulations in the section are listed in Table 6.1 and used in conjunction with the theory and models from Chapter 2 and Chapter 4 respectively.

Table 6.1: Fixed radar parameters for SFPARM multiple target simulation

Parameter	Value
Frequency	$3GHz$
Peak Power	$2kW$
Receiver Noise Figure	$6dB$
Transmitter Duty Factor	0.06
Losses	$6dB$
Boresight Gain	$36dB$

In the simulations in this section the RB-EDF methods uses the following specific rules:

- Earliest track update time is when the angular uncertainty is 0.1 beamwidths.

- Latest track update time is when the angular uncertainty is 0.15 beamwidths.

For all methods the following rules were applied:

- Coherent dwell length selected to maintain the received SNR above 19dB, given an estimate of the target radar cross section. $19dB$ is chosen as a compromise between the $26dB$ used in MESAR [Butler, 1998] and the minimum loading SNR suggested in Chapter 4.

- Minimum beamwidth selected such that the earliest track update time is causal.

The utility function used for LQF is a linear mapping from predicted angular estimation error σ_p:

$$u_k(\sigma_p) = p_i \begin{cases} 1 & \text{if } \sigma_p > 0.15\theta_B \\ 1 - \frac{0.15\theta_B - \sigma_p}{0.075\theta_B} & \text{if } 0.075\theta_B \leq \sigma_p \leq 0.15\theta_B \\ 0 & \text{if } \sigma_p < 0.075\theta_B \end{cases} \qquad (6.5)$$

note this is reversed from Eq. 4.53, as the measure used in the LQF is *lowest quality first*.

To analyse the effect of the resource allocations resulting from RB-EDF, LQF and GIF a simulation has been produced for the target scenario given by the target environmental parameters listed in Table 6.2. Where a parameters range has been given values are randomly generated within that range, and the same generated scenario is used for all three methods. Track starts were staggered and initiated using five dwells separated by $1s$ and deleted if the angular estimation error exceeded 0.3 beamwidths. The tracks were in a background of thermal noise with a false alarm probability of 10^{-4} and no clutter. The tracking filter uses a continuous white noise jerk noise model as the limiting form of the Singer noise model and it is assumed that the target dynamic is matched to the tracking filter. Under this assumption the estimation error covariance correctly describes the error in the track and as such individual measurements were not generated and filtered, and only the covariance matrices propagated. Hence, covariance analysis was used to determine the estimation errors in the simulations. The finite resource constraints were synthesised by extending the length of each radar dwell in proportion to resource availability, for example, 10% resource availability results in each dwell occupying ten times its required dwell length. The simulation is performed in the radar testbed described in Sec. 5.4, using the models of the measurement process in Sec. 4.1.2. In this simulation it is desired to globally optimise the predicted angular estimation error of the target tracks and maximise the number of targets able to be tracked.

Table 6.2: Target environmental parameters for SFPARM simulations

Parameter	Value
Number of Targets	300
Azimuth Region (°)	± 45
Range (km)	$10 - 80$
Priority	1
Radar Cross Section (m^2)	$0 - 20$
Process Noise Intensity	$1 - 50$

The parameter ranges for the environmental target parameters were chosen with a requirement to be representative of a realistic scenario and to be in the range where tracks could be successfully initiated given the fixed radar parameters. For example the azimuth $\pm 45°$ azimuth range is a typical field of view for an ESA, closer than $10km$ targets escape the beam too rapidly to be initiated and above $80km$ targets suffer an SNR which is too low for initiation. By randomly generating targets over a wide range of environmental parameter values, the general behaviour of the mechanism is revealed as the result does not depend on the specifics of a small section of one of the parameter dimensions. Hence the results were relatively insensitive between randomly generated target scenarios and a general assessment of performance is achieved without having to generate simulations for an exhaustive number of environmental parameter changes. The work can be extended to consider different parameter regions if the general behaviour is found to perform well.

Presenting the results from complex scenarios with varying parameters in a meaningful way can be problematic. Each target, of which there are 300, has its own parameter selection, resource loading, quality and so utility which all vary over time. Hence in the following presentations of results average values are taken over the 300 tracks in the simulation. Also, each result includes the result of 10 simulations which run over $120s$ for each resource availability.

As it is required to optimise tracking accuracy, the mean predicted angular estimation error against the resource available for tracking all the targets, which is shown in Fig. 6.4 for the three SFPARM types, allows comparison between the methods. It can be seen in Fig. 6.4 that the allocation resulting from the mutual information utility measure reduces the average angular estimation error in the active tracks. However, this is done by allocating more resource to high SNR targets than they require for maintenance, and less resource to low SNR targets than they require for maintenance. The result is that the tracks have an angular accuracy which is most likely better than what is required, which is a waste of resource. This highlights the important point first asserted in Chapter 4 that purely maximising information production is not the requirement of a multifunction radar.
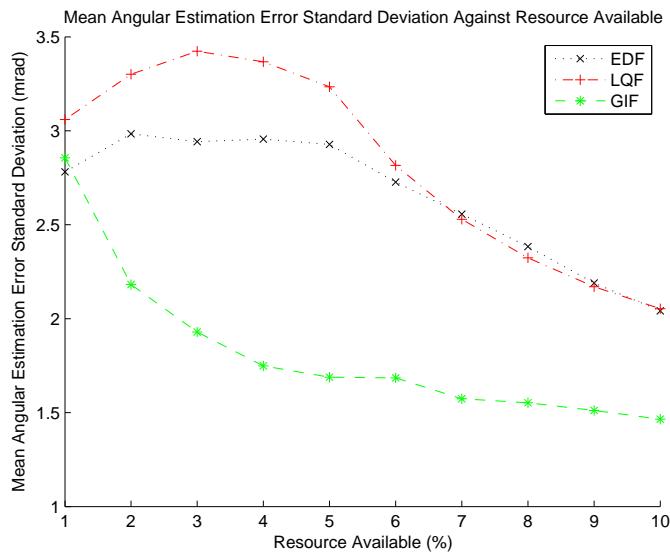
Figure 6.4: Mean track angular estimation error standard deviation for SFPARM types.

Inevitably the resource availability determines how many tracks the system is able to maintain. As it is desired to maximise the number of targets tracked this is also a useful measure of the allocation quality. The number of active tracks maintained against the resource available for tracking all targets is shown in Fig. 6.5 for the same simulation described previously. It can be seen that the number of targets tracked using GIF is low, which is undesirable. This is because allocating resource based on information production causes the radar to focus on a smaller number of 'bright' targets at the expense of many weaker targets, which are probably of more interest. For example, the maintenance of a large number of

low RCS tracks may be of more interest than a small number of high RCS tracks. This again demonstrates that optimising the information production, although theoretically appealing, does not match the requirement of the radar. Although it has been shown to be of value for sensor management in waveform selection, it is less suitable as an objective function for multifunction radar resource management.
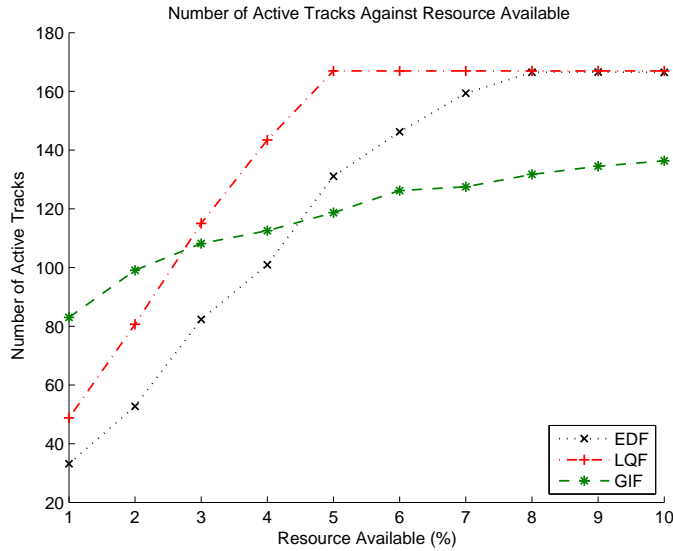
Figure 6.5: Number of active targets for SFPARM variants

The mean track utility is a useful measure of performance because it quantitatively describes the extent to which the allocation meets the requirement of the radar. In these simulations the mean track utility can describe both the predicted angular estimation error and the number of targets maintained in a single measure. The average utility production for the same simulation as described previously is shown in Fig. 6.6. It can be seen that the RB-EDF and LQF perform relatively similarly. This behaviour is contrary to the author's expectation which was that directly controlling task quality would produce superior allocations. This suggests that although task quality is an ideal measure, little performance is lost by translating the track quality requirement, i.e. angular estimation error, into the time domain for the EDF scheduler. However, an EDF scheduler is less computationally demanding as it can be easily implemented in a queue.

Differences between RB-EDF and LQF can be identified in this tracking control example. In overload earliest deadline first inserts a time delay, whose magnitude is emergent, which has an uncontrolled effect on the quality of each task. In contrast, LQF inserts a delay in quality, whose magnitude is emergent. The delay inserted by EDF causes tasks with tight deadlines to miss their deadlines, whereas LQF delays tighter deadline tasks less than fluid deadline tasks. However, the non-myopic nature of the valuation in LQF means that tasks which have a low quality and a large resource demand drain resource from tasks, which only require a small amount of resource for maintenance. This demonstrates the need

126

to optimise over a time horizon, which is lacking these *'best first'* methods. The solution produced is the summation of numerous local optimisations and does not solve the global resource allocation problem. A 'lookahead' can be added to improve the allocation but this would further increase computation.
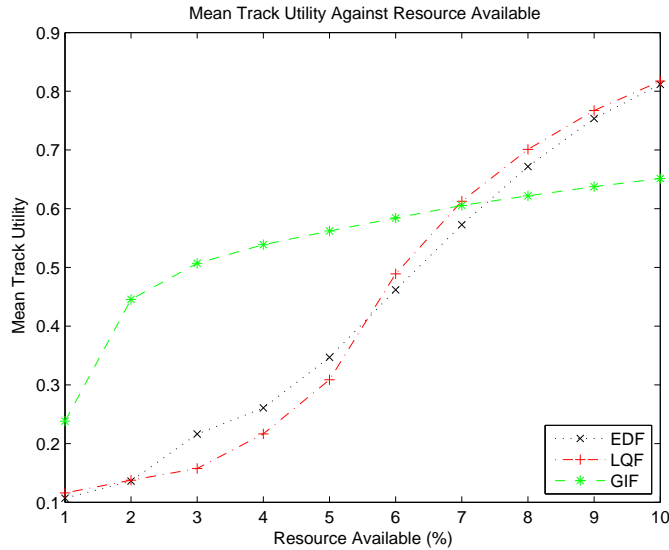
Figure 6.6: Utility production for SFPARM variants

SFPARM requires each agent to evaluate a bid and the auctioneer to compare all the bids for each dwell. As the dwell, and hence the auction process, occur on a fast time scale a significant computation requirement is produced. The similarity between the SFPARM algorithm and a best first scheduling approach allows insight to be drawn from these simulations on best-first scheduling. The conventional EDF can be implemented in a queue with a low computational requirement as the passing of time is the same for all task types. This makes the EDF computationally manageable which explains its application on existing MFR systems. However, if the criterion for the best task is a measure which does not pass equally for all tasks then recalculations of the measure are required on a fast time scale. This has been found in this application of SFPARM as mutual information and utility do not pass equally for tracking tasks with different parameters. Although the computation can be reduced for LQF and GIF by auctioning more than one dwell per cycle, it is still excessive. A better approach would be to utilise an auction mechanism for the resource management decisions, which evolve over a scale of seconds, and keep the scheduler, which operates on a fast time scale, as deterministic as possible and separated from the resource management. This would be more computationally manageable and so more suitable for application to existing radar control software but would require the development of a different auction mechanism.

It has been commented that SFPARM, as a discrete time decision process, is similar to a POMDP which lacks the 'lookahead' which is the effect of the action over an extended time horizon. It was found

in the simulations that LQF and GIF, which required re-computation of measures on a fast time scale, exerted significantly more computation than EDF. This highlights the key problem with stochastic control methods as the computation is known to render these methods intractable as the size of the problem is expanded. Future work could add a 'lookahead' to the calculation of the reward, which is each agents bid in SFPARM, however, this would only increase the computation further. This is not pursued in this thesis, as it is desired to develop methods which can be realistically developed for application to existing radar control software. Also, the POMDP formulation in Fig. 6.2 creates huge branches of actions as many tasks are able to be scheduled. In reality, measurements do not necessarily need to be processed before a different task can be scheduled. For example, the received measurement from a target track has no effect on a track which is well separated spatially.

6.4   Summary

A sequential first price auction is a mechanism which sequentially allocates resource between numerous participants. The auction mechanism has been applied to the radar resource management problem to develop the SFPARM algorithm. The three proposed variants of this algorithm can be thought of as best first schedulers, where the criterion for best relates to time delay, quality and information.

The research conducted in this section successfully produced some key conclusions and outcomes on designing radar resource management mechanisms as follows:

- EDF gives preference to tasks with fluid constraints, LQF gives preference to tasks with least degradation in quality and GIF gives preference to 'bright' targets with high SNR. These preferences are undesirable, to be in keeping with the radar resource management problem preference should be given to best quality using least resource.

- The computation for EDF is manageable, as the passing of time is the same for all tasks it can be handled in a queue. However, when the criterion is switched to a measure which does not pass equally, the computation is significantly increased and values need to be calculated frequently.

- It has been shown that maximising mutual information production leads to a smaller number of high accuracy tracks. Although optimising information is theoretically appealing, it is not what is required from the radar and so is not an appropriate choice of objective function for multifunction radar resource management optimisation. Instead, optimising a variety of task specific quality measures, represented through a utility function is preferred.

- It is necessary to optimise the radar resource allocation over a time horizon. However, mechanisms like SFPARM which are purely myopic can be improved upon by including a time-horizon over which to optimise. This improves the quality of the decision of where the finite resource is best placed.

- The inclusion of a time horizon would make the SFPARM mechanism similar to POMDPs. The computation of these methods is known to rapidly become intractable and so this is not pursued in this thesis which aims to produce a mechanism for realistic application to existing multifunction radar control software.

Despite producing these key research outcomes, the SFPARM mechanism has drawbacks, primarily in the computation required and the quality of the produced allocation. Extensions and improvements are not pursued in this thesis as it is believed alternative auction mechanisms will perform better. The next section will focus on producing a mechanism which is computationally efficient by separating the resource management decisions, which evolve over a slower time scale, from the scheduler. The mechanism should optimise task quality through utility functions and consider where the finite resource is best placed to achieve the global optimum solution.

Chapter 7

Continuous Double Auction

The continuous double auction (CDA) is an auction mechanism which allows numerous participants to trade quantities of a finite resource or commodity, and so closely resembles a free market. The CDA is continuous as trades can occur at any time and is double sided as participants can assume the role of both buyer and seller. The continuous double auction has been applied in numerous stock exchanges and financial institutions, for real-world applications such as the New York Stock Exchange. As such, it has evolved as a scalable, trusted mechanism for rapidly allocating large resource volumes.

This chapter introduces the CDA and describes the application of the CDA to multifunction radar resource management, which has delivered the Continuous Double Auction Parameter Selection (CDAPS) algorithm. The algorithm is analysed in terms of theoretical concepts such as mechanism efficiency and optimality and the scale of improvement verified through simulation. To demonstrate the multifunction capability, simulations are produced for the control of tracking and surveillance functions in example MFR scenarios.

7.1 Continuous Double Auction Mechanism

Variants of the continuous double auction mechanism have been applied in a variety of financial institutions and exchanges, however, the protocols defining interaction vary depending on the application domain. This section details common characteristics of continuous double auctions highlighting potential variations before introducing the New York Stock Exchange as an example of a continuous double auction mechanism.

7.1.1 Market Theory

When a commodity is traded in a free market the higher the price the lower the demand and conversely the lower the price the greater the demand. This can be represented in a supply and demand curve, which is a plot of the quantity and price of the supply and demand in the market. An example supply and demand curve is shown in Fig. 7.1.

The point at which the supply and demand curves meet is known as the competitive market equilib-
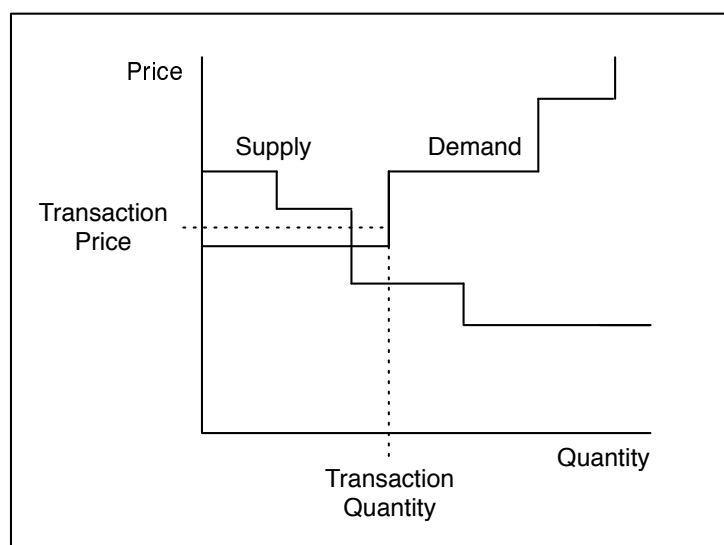
Figure 7.1: Example supply and demand curves

rium. If the market trades at this price and quantity, social welfare is maximised through maximisation of participant profit. The market efficiency is defined as the ratio of the participant profit achieved with a mechanism in relation to the profit associated with an optimal market trading at the competitive equilibrium. A centralised market mechanism is able to find this equilibrium by compiling the complete preferences of all participants and hence achieve optimal market efficiency. Due to the decentralised nature of the continuous double auction, no central auctioneer is in possession of the complete preferences of all the participants. However, transaction prices in the continuous double auction do converge to the competitive market equilibrium and so adequate market efficiencies can be generated.

7.1.2 Mechanism

A market mechanism is defined by the protocol for participant interaction. The protocol determines the format of admissible resource offers, the information which is publicised to the auction participants and the conditions under which transactions can occur. Although variants of the continuous double auction exist, most research is based around the structure first developed in Smith [1962]. In this form of CDA, offers are announced for single unit quantities with a spread improvement and no order queue both of which are explained below.

Despite wide variations between implemented continuous double auction, the following concepts are always present:

- *Bid* - A request to purchase resource, which typically contains a unit price, quantity and identifier.

- *Ask* - A request to sell resource, which typically contains a unit price, quantity and identifier.

- *Offer* - A request to purchase or sell resource.

- *Trading Round* - The duration between cleared transactions over which bids and asks are announced.

- *Trading Day* - The allowed trading period, which contains numerous trading rounds.

- *Orderbook* - A collection of the best active bid and ask prices. For single unit mechanisms it is only required to hold the best single bid and ask meaning there is no order queue.

- *Bid/Ask Price Limit* - The minimum bid price or maximum ask price allowed for admission to the market.

Whilst the above concepts are open to variation, the primary variations in a continuous double auction are the pricing rule, clearing rule and offer queue. The clearing rule determines when the auction clears to generate a transaction. For single unit auctions this is trivially when the bid price exceeds the ask price, however, for multiple unit auctions this can become more complicated. The pricing rule determines at what price the subsequent transaction occurs. The transaction price $\tilde{p}$ is commonly found using the k-pricing rule:

$$\tilde{p} = \hat{k}p_b + (1 - \hat{k})p_a \tag{7.1}$$

where p_b is the bid price, p_a is the ask price and $\hat{k}$ is commonly taken as 0.5. The queuing rules determine the nature of the orderbook which records the active offers. For single unit auctions a spread improvement rule can be enforced which reduces the number of announced offers by requiring each new offer to be an improvement upon the last announced offer. For multiple unit auctions a similar offer limit can be applied to reduce the communication overhead of uncompetitive offers.

7.1.3 New York Stock Exchange

To cement the concepts described in the preceding subsections it is useful to examine the New York Stock Exchange (NYSE) as an example of an operational continuous double auction. As the worlds largest stock exchange with average daily trades of hundreds of billions of U.S dollars this is a relevant example of an efficient and trusted mechanism.

The NYSE is split between the upstairs and the downstairs markets, where the upstairs market specialises in large stock volumes. In the downstairs market there are seventeen trading posts which are split across numerous rooms at which traders are able to trade stocks. Each stock listed in the exchange has a specialist located at a trading post who is responsible for facilitating the trading between brokers. Although the specialist helps to match the bids and asks of traders, they are involved in a small number of trades and the majority of trades occur by traders self-matching without the specialist. The specialist is also responsible for stabilising the market by limiting successive transaction price changes, buying when the price drops and selling when the price rises. During the trading day, which runs between 9.30 and

16.00 EST, the NYSE functions as a continuous double auction. However, when the NYSE is closed, offers are received and stored in Opening Automated Report Service (OARS). When the NYSE opens the specialist decides on a transaction price to clear the trades in the OARS.

A general architecture for the system involving the NYSE is shown in Fig. 7.2. Within the NYSE specialists are able to interact with the traders, who are also able to interact with each other. The subsequent trades which emerge from this interaction determines the valuation of the companies's stock. The responsibility of the board of directors is to create value, which is directly related to the stock valuation and so the NYSE directly affects the actions taken by the company. The action taken by the company determines how it is perceived by potential investors and fund managers which determines how the stock is valued by the NYSE.
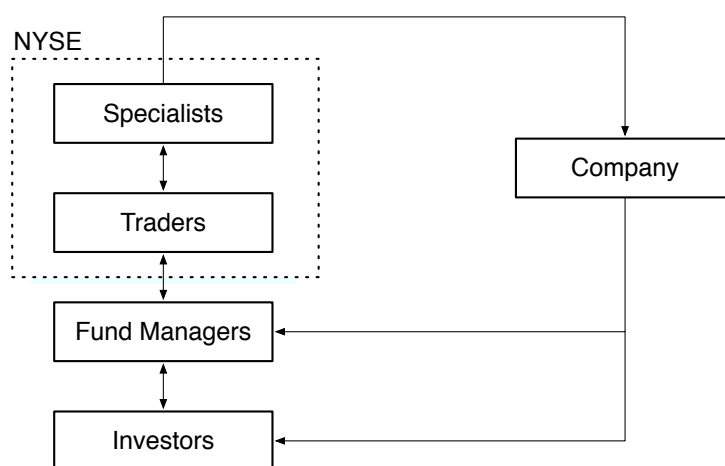
Figure 7.2: New York Stock Exchange system architecture

Simple analysis shows that this system contains numerous elements identified as required in a cognitive system. Namely:

- *Memory* - There are several memory elements throughout the system. The specialists provide short term memory, whereas investors and fund managers provide longer duration memory and the inclusion of fixed knowledge.

- *Hierarchical Memory Structure* - There are a small number of specialists and so a small memory capacity, however, these specialists can act with a very low latency. Conversely there are a large number of investors providing a vast memory capacity, however, this memory has a high latency.

- *Micro Feedback* - There are numerous micro feedback loops between elements on each hierarchical memory level as well as between levels.

- *Macro Feedback* - There is a global macro feedback loop which incorporates the company, the

investors and the NYSE.

Clearly this implementation of a continuous double auction creates a sophisticated data processing system which is capable of processing a large amount of uncertain information.

The CDA as an economic paradigm can be adapted as a methodology for radar resource management and signal processing optimisation. In a CDA, such as the NYSE, available information is efficiently processed leading to reconfigurations of resource allocations which match the market to the dynamic financial environment. This precisely echoes the desire for next generation sensor systems to efficiently process received radar measurements to match the signal processing applied by the radar to the dynamic and uncertain sensing environment.

The desire for cognitive signal processing is a response to the functionality of next generation radar systems being fundamentally limited, not by hardware, but by the radars ability to utilise sensor information to generate autonomous and adaptive behaviour. By incorporating desirable information processing characteristics, this form of economic paradigm can be used as a tangible step towards developing cognitive sensor signal processing techniques.

7.2 Continuous Double Auction Parameter Selection Algorithm

This section describes the Continuous Double Auction Parameter Selection (CDAPS), which is the result of the application of the continuous double auction from the previous section to the radar resource management problem. The CDAPS algorithm implements a market mechanism which manages the allocation of the resources of the radar system through the selection of parameters for the individual radar tasks. This section describes the CDAPS mechanism and the participating agents.

7.2.1 Mechanism

The CDAPS algorithm hosts a market mechanism where agents representing the numerous radar tasks, such as tracking or surveillance, can trade resource. The distributed, decentralised nature of the mechanism provides scalability and allows each task agent to be designed independently. The CDAPS mechanism is implemented in the radar resource management testbed as shown in Fig. 5.4. This architecture allows easy integration into a typical radar resource management architecture, shown in Fig. 3.3, by replacing the task request modules. The CDAPS algorithm selects task parameters from a usable waveform database given the model of the current scenario which includes priority assignment, the usable waveform database and the current state of the radar task function. The global feedback enables the update of the model of the current scenario from the received measurements and local feedback within the CDAPS mechanism ensures the parameters are selected subject to the finite resource available. The selected task parameters are used to issue tasks requests which are formed into a timeline by the scheduler.

The resource traded by radar task agents in the mechanism represents radar loading as described in

Sec. 4.1.1.1 and Eq. 4.1. The resource held by radar task agent k, which is denoted r_k, represents the allowed sensor loading of its represented task. The total resource held by all radar task agents, denoted r_T, cannot exceed the radar resource loading available for all tasks, i.e. $\sum_k r_k \leq r_T$. Trading is driven by each radar task agent having differing and potentially dynamic valuations of the resource in terms of the system currency, known as utility, as described in Sec. 4.4. Each task agent can simultaneously assume the role of both buyer and seller to facilitate the continuous resource trades. A single auctioneer agent is present to implement the protocol of the mechanism which can be formalised as follows:

- There exists a set of m task agents $T_A = \{t_1, ..., t_m\}$ which represent numerous tasks performing radar functions.

- Each agent may publicly announce an offer to trade as a bid to buy, an ask to sell or both. The offer comprises of a quantity q, unit price p and identifier. An ask from agent m has the form $a_m(q_m, p_m, m)$ and a bid from agent n has the form $b_n(q_n, p_n, n)$.

- At a given time there is a set of active asks $A_A = \{a_1, ..., a_m\}$ and a set of active bids $B_A = \{b_1, ..., b_n\}$.

- After each new offer is announced the auctioneer attempts to find a valid transaction set of asks $I \subseteq A$ and bids $J \subseteq B$. The value and quantity of the transaction ask set I is $V_I = \sum_i p_i q_i$ and $Q_I = \sum_i q_i$ respectively. The value and quantity of the transaction bid set J is $V_J = \sum_j p_j q_j$ and $Q_J = \sum_j q_j$ respectively.

- The transaction clearing rule declares a bid and asks set valid if $V_I < V_J$ and $Q_I > Q_J$.

- The transaction price $\hat{p}$ is a weighted average of the lowest ask price in the ask transaction set i_{min} and the lowest bid price in the bid transaction set j_{min} so $\hat{p} = 0.5 i_{min} + 0.5 j_{min}$.

As it is required to trade in multi-unit quantities the transaction clearing rule is more complicated than implemented by Smith [1962]. The transaction clearing rule allows transactions with unequal quantities, in which case the excess $e = Q_J - Q_I$ is held by the auctioneer and included in the next transaction. Each offer remains active until it is cleared or updated by the agent, and only the best fifty bids and asks are kept active. Trading rounds are generated as the auction clears continuously for the duration of the radar's operation, and so there is no set trading day.

This mechanism is suited for dynamic resource allocation as the resource for arriving task agents is met by taking resource away from the tasks who lose the least amount of utility per unit resource. When a task agent becomes inactive the resource is purchased by task agents who gain the most utility per unit resource.

7.2.2 Task Agents

Each agent represents a radar task and aims to maximise its utility production by acquiring as much resource, which is radar loading, as possible given the competition from the other task agents. To ensure the validity of the allocation, each agent's valuation of potential increases or decreases in resource must be accurately related to its radar task. This can be achieved by assessing the quality and resource loading of potential task parameters as described in Sec. 4.4.

Following a similar model to Hansen et al. [2006], each agent has a current task parameter selection γ_k from the potential parameter space Γ. Each task parameter selection has a different resource loading which is derived from the task's resource function, $g_k : \Gamma_k \to \Re$, which provides a mapping from parameter to resource space. Each task parameter selection also produces a different task quality which is derived from the task's quality function $q_k : \Gamma_k \to \hat{Q}$, which provides a mapping from parameter to quality space. Finally, a utility function is required, $u_k : \hat{Q}_k \to \Re$, which maps the task quality into auction utility, giving the satisfaction associated with each point in the tasks quality space. As the primary quality metric of interest varies between radar task types, the utility function provides a comparable measure, used as currency, between tasks assessed by different quality metrics.

As each combination of task parameters has a different radar resource loading, produces a different task quality and so also a different utility, they occupy different points in resource-utility space, as shown in Fig. 7.3. It is desired to select parameters to maximise utility per resource. Potential changes from the current parameter selection can be evaluated as the difference in utility, Δu, given the change in resource, Δr which is the gradient between resource-utility points or the difference in utility per unit resource. This gradient is the agents true price valuation, $p*$, of the potential change in parameters:

$$p* = \frac{\Delta u}{\Delta r} \tag{7.2}$$

A hill climbing search is used to find the potential change in parameters which gives the best ask and bid offers. The criterion for the best bid is the largest bid price, or largest increase in utility per unit resource. Conversely the criterion for the best ask is the lowest bid price or smallest decrease in utility per unit resource. Large gradients and offer prices, observed in the lower resource region of Fig. 7.3, can produce larger utility increases per unit resource than the high resource region, where the task becomes saturated. An example of the gradient used to calculate the offer price is shown in Fig. 7.3.

Crucially, the best bid and ask prices are local due to the monotonic nature of each parameter dimension, which reduces the search space. Additionally, new bids and asks are generated over a time scale of seconds, as new data is received or the environment changes, which spreads the search over time. This synchronises changes in the allocation to the changes in the environment and so greatly reduces the computational demand.
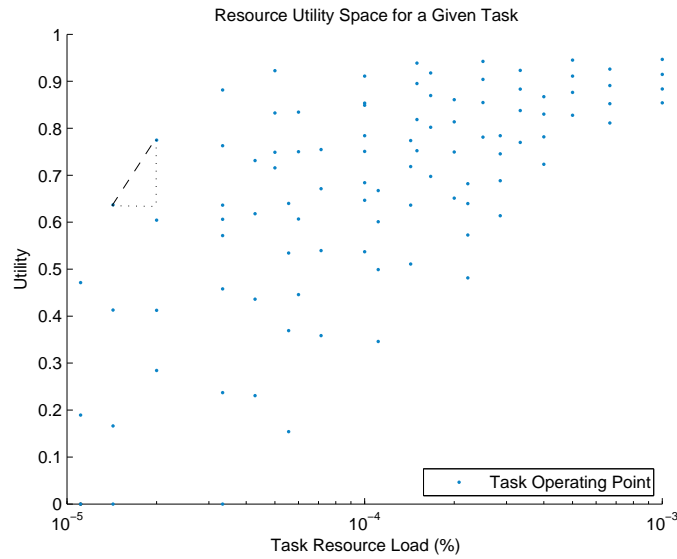
Figure 7.3: Resource utility space with example gradient between parameter selections marked

Given the agent's true valuation it must decide on what price and resource quantity to announce to the market. Various bidding strategies [Gode and Sunder, 1993; Gjerstad and Dickhaut, 1998] have been studied, all of which depend on the agent's true valuations of the resource, but may additionally use information on the current market state or market history. The CDAPS algorithm assumes truth telling agents, who offer their true valuations, with no knowledge of the market state. The quantity of the offer is the change in radar resource loading resulting from the potential change in task parameters. Subsequent trades occur with an increase in total utility, which improves overall system performance. The resulting competitive market equilibrium price can be visualised as selecting parameters on constant gradients across all the tasks in resource-utility space.

As implemented in the radar resource management testbed, the task agent possesses the behaviours described in Sec. 5.4.2.1. In addition the task agent also possesses the following behaviours:

- *Process Transactions*: Update the task parameter selection according to a previous transaction.

- *Submit Offers*: Evaluate the utility of the current parameters and search for the best offer to announce to the market.

these behaviour allow the agent to participate in the continuous double auction mechanism.

7.2.3   Auctioneer Agent

The auctioneer agent organises the market mechanism, which involves the maintenance of a public list of the best active bids and asks, called the orderbook. After each new offer is announced the auctioneer attempts to generate a transaction by clearing the orderbook. If the orderbook can successfully clear the auctioneer facilitates the transaction by communicating between the buyers and sellers. The auctioneer

agent is implemented through the following behaviours:

- *Collect Offers*: Listen to announced offers, update the orderbook if necessary and check if the auction can clear.

- *Process Clearing Details*: Communicate a potential transaction to the agents involved.

The size of the workbook is limited to fifty bids and asks, to limit the computation burden placed on the auctioneer. The auctioneer publicises the lowest bid and the highest ask price required for addition to the orderbook. Task agents do not announce offers which do not meet this criterion which reduces the number of offers announced.

The auction clears and a transaction is generated when a set of bid and asks has higher bid value than ask value, with bid prices greater than the ask prices, i.e. $V_I < V_J$ and $Q_I > Q_J$. The requirement on the value in addition to the price is not common in continuous double auctions but is implemented in the CDAPS algorithm as a consequence of offers having multiunit quantities and each offer representing a parameter selection. This ensures that if the bid and ask quantities are unequal, the subsequent switch in parameters resulting from the trade has a positive effect on the total utility of the system.

The decision process used to clear the orderbook attempts to produce a transaction set of asks and bids and is shown in Fig. 7.4. The transaction set starts with the best bid, and bids and asks are added depending on the set quantity until no further offers can be added. The best valid transaction set, if it exists, generates a transaction. An example of this decision process is shown in Table 7.1 and Table 7.2. In Table 7.1, which is based on the supply and demand curve in Fig. 7.1, a valid transaction set of asks $I = \{a_3\}$ and bids $J = \{b_7, b_2\}$ can be found with values $V_I = 25$, $V_J = 26$ and quantities $Q_I = 5$ and $Q_J = 4$. However, in Table 7.2, the previous transaction set is not valid as $V_I = 25$ and $V_J = 24$ and no other valid transaction set exists. The orderbook can be thought of as an incomplete estimate of the current market supply demand, which differs from a centralised market, or optimisation RRM, which compiles the complete supply and demand preferences. The efficiency of the continuous double auction mechanism can in part be attributed to the compilation of the incomplete preference estimate instead of the full preferences.

Table 7.1: Example of an orderbook which is able to clear

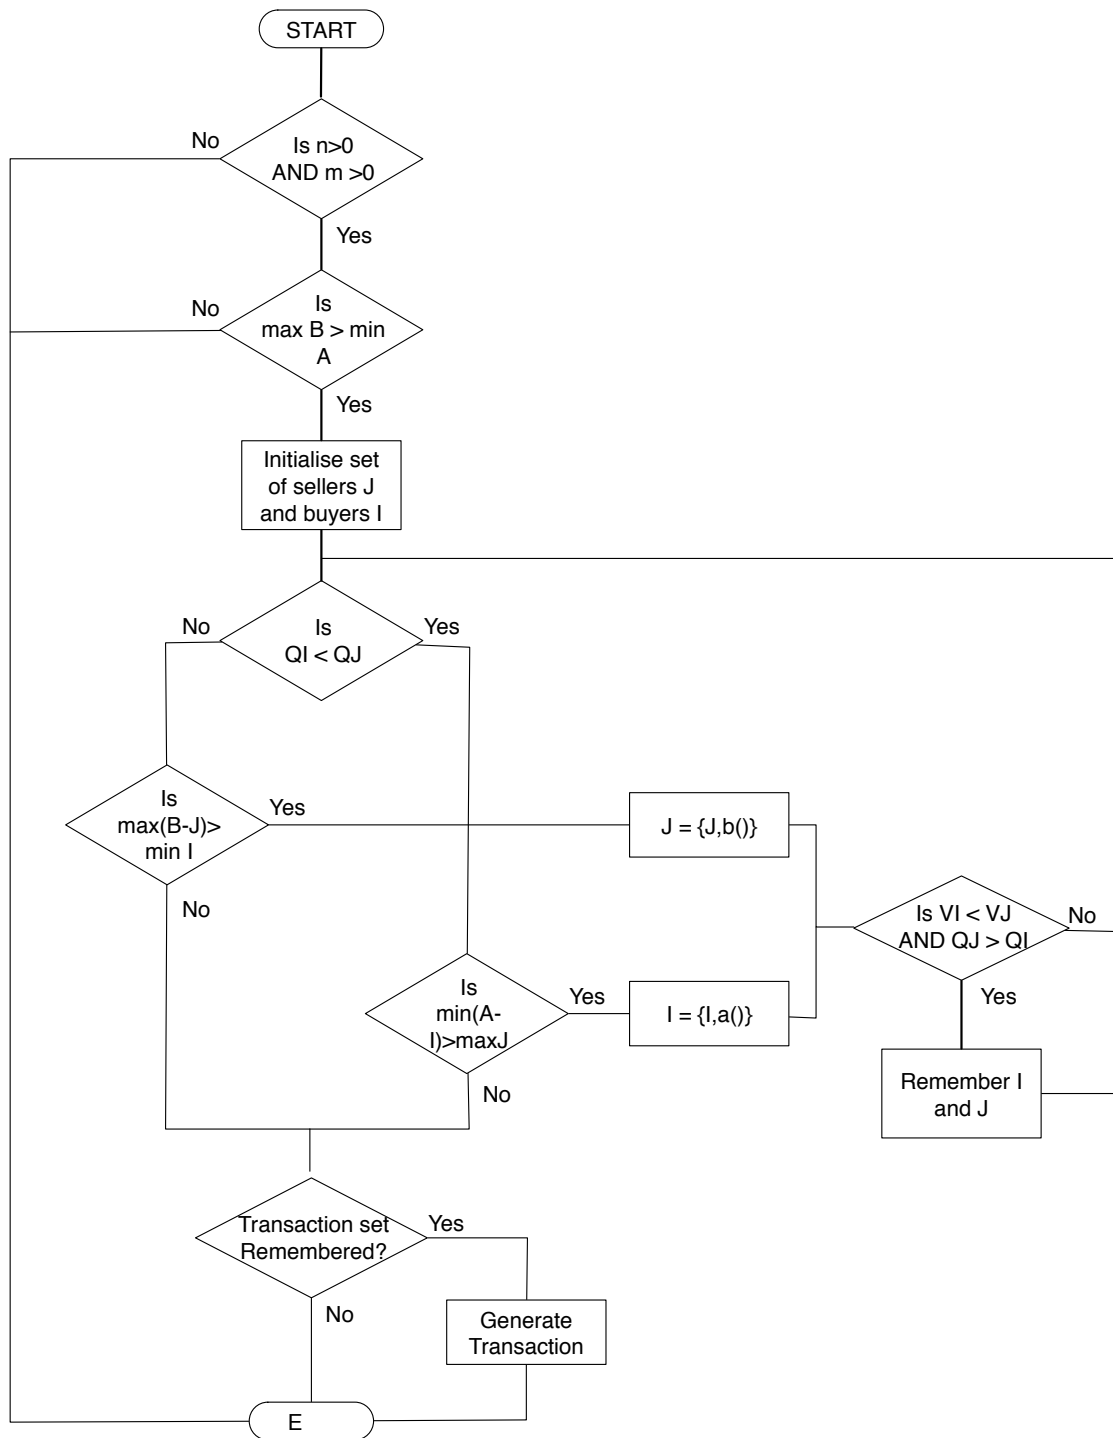| Bids (Buyer, price, quantity)	Asks (Buyer, price, quantity)
$b_7(7, 7, 2)$	$a_3(3, 5, 5)$
$b_2(2, 6, 2)$	$a_6(6, 7, 4)$
$b_9(9, 4, 3)$	$a_1(1, 9, 2)$

Figure 7.4: Auction clearing decision process in CDAPS

Table 7.2: Example of an orderbook which is unable to clear

Bids (Buyer, price, quantity)	Asks (Buyer, price, quantity)
$b_7(7, 7, 2)$	$a_3(3, 5, 5)$
$b_2(2, 5, 2)$	$a_6(6, 7, 4)$
$b_9(9, 4, 3)$	$a_1(1, 9, 2)$

7.3 Theoretical Analysis

Analysis of resource management algorithms is difficult for a number of reasons. Firstly, no single metric exists to assess performance and there are a wide variety of scenarios in which the resource management algorithm is required to operate. Secondly, the quality of the behaviour generated by the resource management algorithm can only be assessed in the context of what is required from the system, which may also be dynamic. Finally real data is of limited use, unless it is sufficiently oversampled, as the data capture must have already applied some resource management, even if this is a fixed as for a mechanically scanned system. Hence, theoretical considerations such as the mechanism efficiency and optimality are key tools for initial assessment of performance, which are described in this section.

7.3.1 Mechanism Efficiency

The efficiency of a market mechanism is a ratio describing the ability of a specific non-ideal market to maximise participant profit in comparison to the ideal market. Centralised mechanisms are capable of obtaining optimum market efficiency by compiling the complete participant preferences, and hence complete market information, to find the optimal transaction price which is found at the point where supply equals demand. This is conceptually equivalent to, upon entering a local high street market, being required to submit bids and asks for all combinations and quantities of available commodities. This is not only an excessive computational burden for each participant but also for the central auctioneer due to the number of bids. Although the optimum solution would be found, it is impractical to implement such a mechanism due to the computation burden; it is straightforward to see why this is not applied for market mechanisms in human societies. However, it has been shown by Smith [1962] that decentralised mechanisms, specifically the continuous double auction, can achieve close to the optimum efficiency using just a fraction of the participant preferences, which is only part of the market supply and demand and hence a fraction of the computation.

In the context of radar resource management, the mechanism efficiency represents how close to the optimum solution the resource allocation mechanism is capable of achieving. Existing optimisation approaches for radar resource management, such as dynamic programming and Q-RAM, are similar to centralised mechanisms where a large quantity of preferences are collected leading to a heavy computational demand. This is undesirable as the greater the computational demand, the less scalable the

allocation mechanism, the less tasks it will be able to maintain and the poorer the system performance against pop-up tasks which require rapid reaction. The result in Smith [1962] indicates that the CDA converges on similar quality allocations as centralised, or existing optimisation based RRM methods, with a fraction of the information processing. Hence this result suggest the CDAPS algorithm can achieve a performance equivalent to existing optimisation methods with a fraction of the computational burden. This is highly desirable as it means a significant improvement in the performance of the numerous radar tasks executed can be achieved with computation that is realistic and suitable for application to existing multifunction radar control software.

7.3.2 Optimality

Optimality is a very important theoretical aspect as if it can be shown that the solution is optimal then it is ensured that the optimisation, i.e. maximisation of the objective function, cannot be improved. A necessary condition for a non-linear programming solution to be optimal is that the Karush-Kuhn-Tucker (KKT) conditions are satisfied whereby the marginal utilities, or gradients in resource-utility space, are equal. This concept was used to develop the Q-RAM [Hansen et al., 2006] algorithm which produces solutions which maximise resource utilisation whilst satisfying KKT conditions. The CDAPS algorithm also relies on the KKT conditions by producing an optimal solution with equal marginal utilities as the emergent competitive equilibrium from the market mechanism. However, as the possible parameter selections are discrete, the solution is optimal for the given discrete parameter set, but only near optimal in contrast to a continuous parameter set [Irci et al., 2010].

This principal is demonstrated in Fig. 7.5-7.6. Fig. 7.5 shows the possible parameter selections in resource-utility space for three different example long range surveillance tasks. The possible parameter selections are the same for all three tasks, however, as the environmental parameters outside of control differ, each task produces a different utility for identical parameter selections. It is desired to select parameters along the concave majorant where utility per resource is maximised, and the concave majorant can be followed using a hill climbing search. The concave majorants for the three surveillance tasks are shown in Fig. 7.6. If parameters are selected such that the gradients in resource-utility space are equal then the KKT conditions are satisfied. One such selection is shown in Fig. 7.7. These selected parameters use a total of 0.5% of the available resource. By satisfying the KKT condition, this is the optimal parameter selection for 0.5% resource loading, and so produces the maximum utility.

CDAPS naturally satisfies the KKT conditions whilst using the maximum possible resource available. This ensures that the parameters selected for the varied radar tasks collectively produce the global maximum utility. As utility is a mapping from a variety of relevant radar task quality measures, this selected parameter set is the best combined quality for all radar tasks that can be achieved given the finite resource. So, optimal optimisation of the radar resource management problem is achieved.
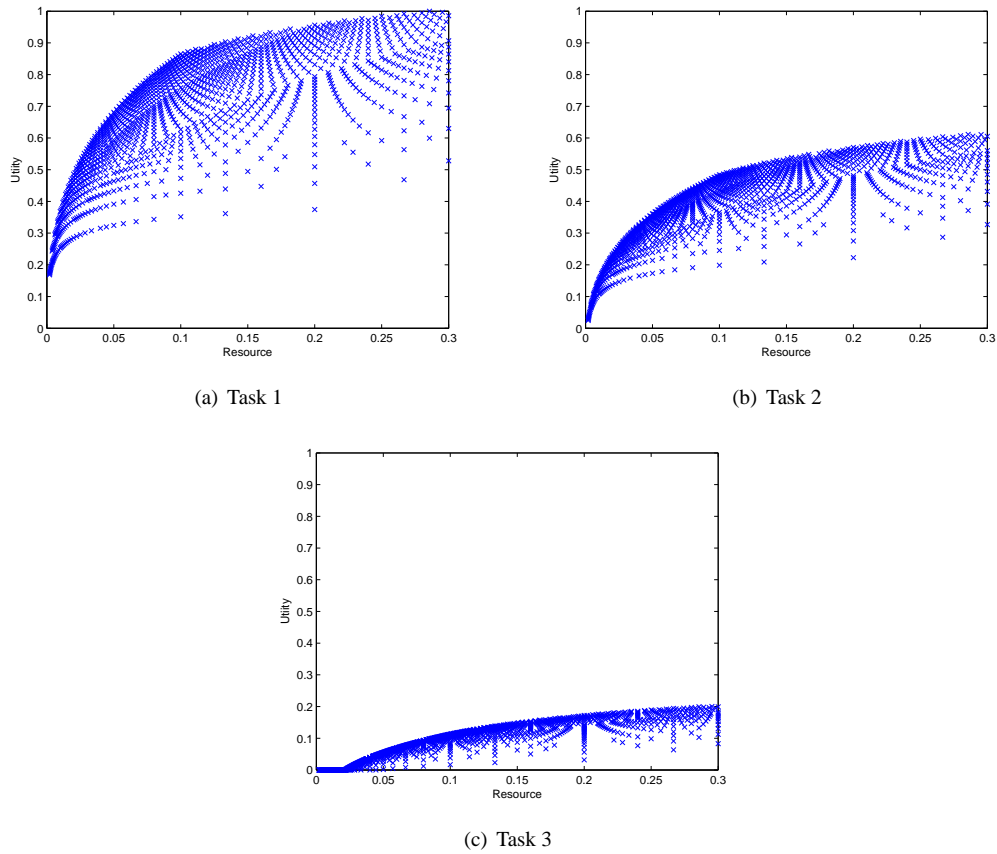
(a) Task 1

(b) Task 2

(c) Task 3

Figure 7.5: Possible parameter selections for three example surveillance tasks

From a resource allocation mechanism perspective this is very satisfying result as it guarantees the optimal performance of the allocation mechanism. Subsequently, the ultimate performance is solely dependent on how well the different radar tasks can be modelled, measured and translated into utility, which may be problematic in some situations. For example, if the target dynamic noise is modelled incorrectly for a target track then the quality of the task and hence utility will be incorrectly calculated and resource incorrectly allocated, despite the optimal operation of the allocation mechanism. This highlights the importance of the measures and models described throughout Chapter 4.

Ensuring optimality through the KKT conditions is a very important concept as it ensures that the numerous localised agents are able to collectively solve the global resource utilisation objective of the radar resource management problem. This proves that CDAPS applied in an agent system is highly suitable for multifunction radar resource management.

7.4 Simulation Analysis

The theoretical concepts in Sec. 7.3 show that the CDAPS algorithm should be computationally efficient in comparison to existing radar resource management optimisation methods whilst also producing a globally optimal allocation which improves upon existing rule based methods for multifunction radar
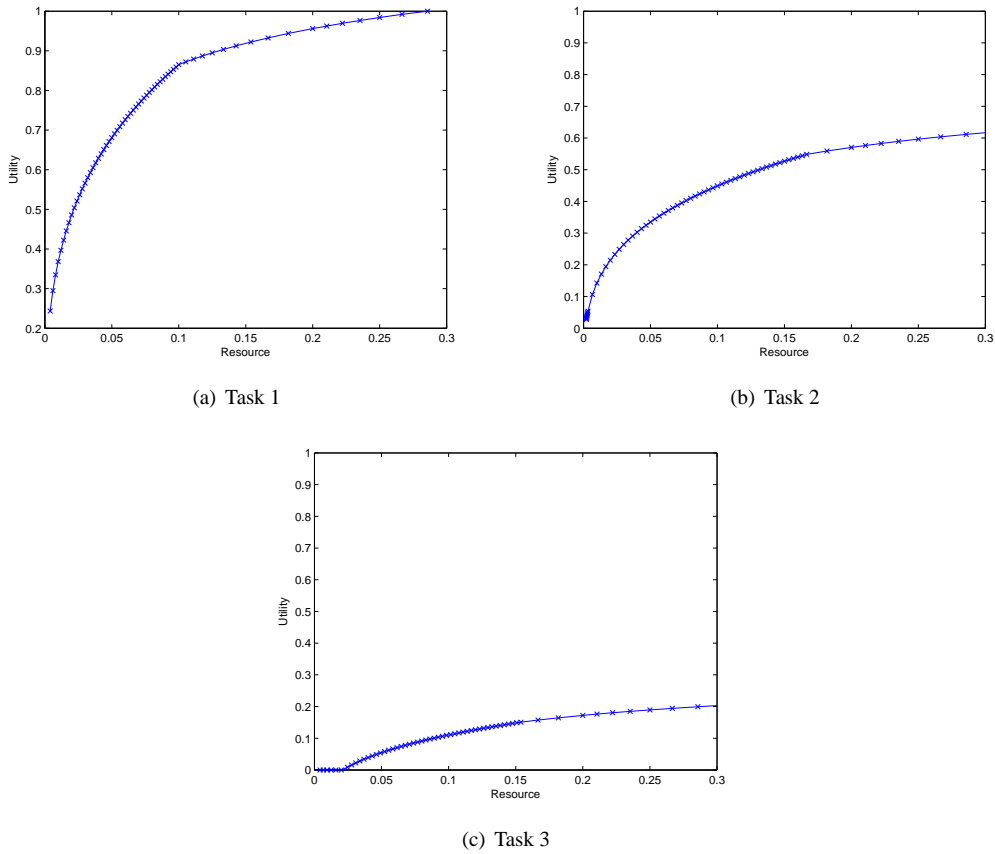
(a) Task 1

(b) Task 2

(c) Task 3

Figure 7.6: Concave majorant for three example surveillance tasks

resource management. In this section, simulation is used to assess the performance of CDAPS for example radar resource management problems. The tracking and surveillance functions for a MFR are used to demonstrate the potential improvement in resource allocation using CDAPS.

7.4.1 Tracking Control

A number of tracking control simulations have been generated, using the radar resource management testbed from Sec. 5.4, to analyse the performance of CDAPS in contrast to existing techniques. In the simulations, targets can be tracked using measurements from fixed surveillance, known as track-while-scan, or by using a share of additional resource dedicated for active updates. It is desired to select parameters to optimise tracking performance, which is the angular estimation error and number of targets tracked, subject to the finite resource available. The target environmental parameters considered outside of control are range, azimuth, radar cross-section and manoeuvre model standard deviation. Operational parameters under control relate to the waveform selection, which is simplified to the choice of revisit interval, dwell length and beamwidth, assuming that lengthening the dwell increases the SNR according to ideal coherent integration. The target dynamic was assumed matched to the tracker which used a continuous white noise jerk model. As under this assumption the estimation error is correctly described
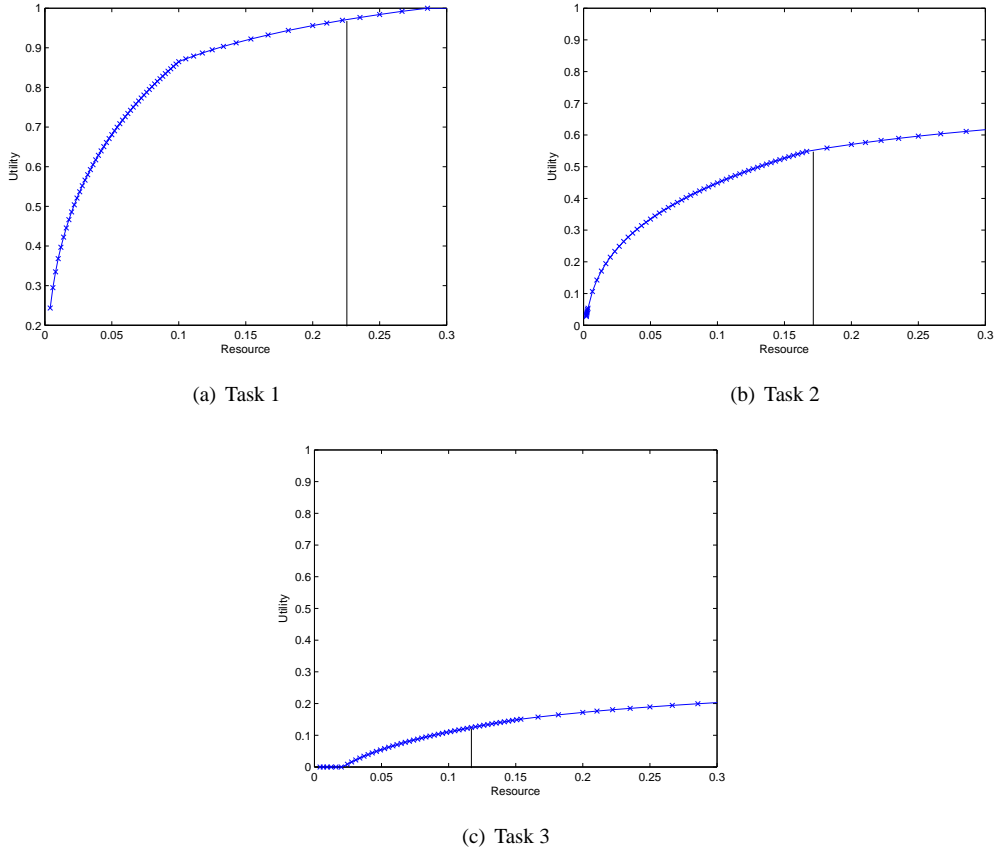
(a) Task 1

(b) Task 2

(c) Task 3

Figure 7.7: Equal gradient points for three example surveillance tasks

by the filter covariance, covariance analysis was used to assess performance and only the covariance matrices propagated without generating or filtering measurements. In all cases the tracks were in thermal noise with false alarm probability of 10^{-4} and no clutter. Track were initiated using five dwells separated by $1s$ and deleted when the angular estimation error exceed 0.3 beamwidths.

A variable finite resource was synthesised by prohibiting any other task to be scheduled for a duration after each task was scheduled. This duration length depends on the desired resource availability. For example, given a $1ms$ dwell and a 5% resource availability it would not be possible to schedule another task for a further $19ms$. Likewise, for a $1ms$ with a 10% resource availability it would not be possible to schedule another task for an additional $9ms$. Although this simulates resource being used for surveillance, the surveillance does not generate new tracks.

As a basis of comparison a rule based parameter selection (RBPS) strategy has been used with rules chosen to be similar to those used in existing MFR systems [Noyes, 1998; Butler, 1998]. This strategy specifies desirable regions in quality space, such as maintaining the angular uncertainty beneath a fraction of the half beamwidth as described in Sec. 3.1.2.1. These rules do not specify how or which parameters should change given a resource constraint leaving the scheduler to mediate access to the resource. An earliest deadline first scheduler is used to form the radar timeline for both RBPS and

CDAPS [Butler, 1998].

The rules used to select parameters for rule based parameter selection are:

- Earliest track update time is when the angular uncertainty is 0.1 beamwidths.

- Latest track update time is when the angular uncertainty is 0.15 beamwidths.

- Coherent dwell length selected to maintain the received SNR above $19dB$, given an estimate of the target radar cross section. $19dB$ is chosen as a compromise between the $26dB$ used in MESAR [Butler, 1998] and the minimum loading SNR suggested in Chapter 4.

- Minimum beamwidth selected such that current track accuracy is less than 0.1 beamwidths.

These rules are selected to be aligned with the studies on tracking control parameter selection described in Sec. 3.1.2.1.

The utility function used for CDAPS is a linear mapping from angular estimation error standard deviation σ_p:

$$u_k(\sigma_p) = p_i \begin{cases} 0 & \text{if } \sigma_p > 0.15\theta_B \\ \frac{0.15\theta_B - \sigma_p}{0.075\theta_B} & \text{if } 0.075\theta_B \leq \sigma_p \leq 0.15\theta_B \\ 1 & \text{if } \sigma_p < 0.075\theta_B \end{cases} \qquad (7.3)$$

The fixed radar parameters used to produce the following simulations are detailed in Table 7.3. these parameters are used in conjunction with the theory and models from Chapter 2 and Chapter 4

Table 7.3: Fixed radar parameters for CDAPS simulations

Parameter	Value
Frequency	$3GHz$
Peak Power	$2kW$
Receiver Noise Figure	$6dB$
Transmitter Duty Factor	0.06
Losses	$6dB$
Boresight Gain	$36dB$

respectively.

7.4.1.1 Static Scenario

The first scenario consisted of targets having uncontrollable environmental parameters as listed in Table 7.4, which all require tracking subject to the finite resource available. The target environmental parameters remained static over the duration of the simulation. Where a parameter range is given values are random generated uniformly across the range. Given this target scenario it was desired to optimise tracking performance for all targets, in terms of the predicted angular estimation error. The agents were generated in the radar resource management testbed described in Sec. 5.4.

Table 7.4: Target environmental parameters for static scenario

Parameter	Value
Number of Targets	300
Azimuth Region ($°$)	± 45
Range (km)	$10 - 80$
Priority	1
Radar Cross Section (m^2)	$0 - 20$
Process Noise Intensity	$0 - 50$

Examining the delay inserted by the earliest deadline first scheduler can provide insight into the resource allocation mechanism's ability to degrade gracefully. Fig. 7.8 shows the delay from the desired task execution time inserted by the scheduler for RBPS and CDAPS over the duration of an overloaded and underloaded simulation. It shows that RBPS inserts a delay depending on the system loading, whereas CDAPS balances the time budget regardless of resource availability. In CDAPS, the total resource held by all agents is matched to the sensor resource available, and so the selected parameters maximise utility whilst balancing the time budget. In contrast, the collective operating points specified by RBPS may or may not exceed the time budget, requiring adjustment by the scheduler. The inserted delay for RBPS has a differing, uncontrolled effect on the quality of each task, which leads to non-graceful degradation. The selection of parameters by CDAPS has a controlled effect on task quality and produces graceful degradation.
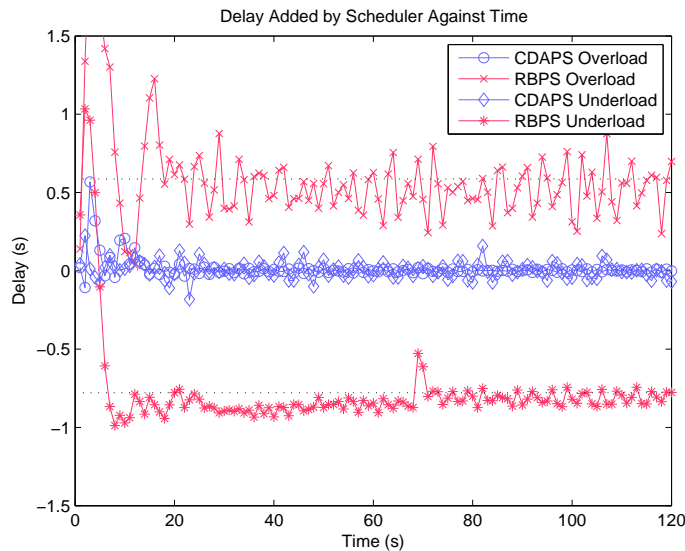

Figure 7.8: Time delay inserted by the scheduler for CDAPS and RBPS

As the requirement of the tasks is to globally optimise the track angular estimation error standard deviation, the average of the track estimation errors is also a useful indication of the quality of the resource allocation. Fig. 7.9 shows the average track angular estimation error standard deviation across all tracks

for RBPS and CDAPS with varying active track resource. Simulations of duration $120s$ were produced for each resource availability with identical target environmental parameters according to Table 7.4, with the resource availability synthesised by lengthening each radar dwell. It shows that CDAPS significantly improves the angular estimation error in the tracks. This is a result of CDAPS trading in differential utility and so resource is allocated where the greatest improvements in quality can be achieved. Although the scale of this improvement depends on the scenario and the rules, globally optimised CDAPS continually outperforms the locally optimised RBPS due to satisfying the KKT conditions. Equal performance is achieved when there is 0% resource available for active tracking, as all tracks are supported using TWS.
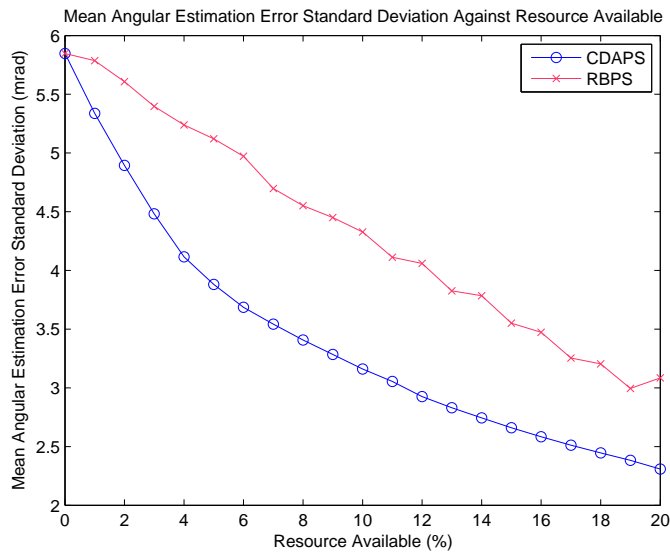
Figure 7.9: Mean track angular estimation error standard deviation for CDAPS and RBPS

These simulated static scenario tracking control problems have demonstrated how CDAPS can degrade gracefully and provide an improvement in global task quality, which in this case is tracking accuracy, in contrast to a conventional rule based approach.

7.4.1.2 Dynamic Scenario

In a multi-target tracking application, a dynamic environment is typically manifested by dynamic target number, target kinematics and measurement origin uncertainty. An effective resource management mechanism is required to adapt in a timely fashion to the evolving dynamic environment.

To analyse adaptation to an evolving target number, a simulation was produced where the number of targets increased for the first minute from 150 to 300 and decreased for the second minute back to 150. Although this scenario is highly contrived, it is adequate to demonstrate the behaviour of the mechanism. The uncontrollable target environmental parameters used for the simulation are listed in Table 7.5, the fixed radar parameters are listed in Table 7.3. Given this target scenario it was desired to optimise tracking performance across all targets, in terms of the angular estimation error. Covariance

analysis was used to determine the estimation error and the target dynamic was assumed matched to the tracker which uses a continuous white noise jerk model. The agents were generated in the radar resource management testbed described in Sec. 5.4.

Table 7.5: Target environmental parameters for dynamic scenario

Parameter	Value
Number of Targets	$150 - 300$
Azimuth Region $^\circ$	± 45
Range (km)	$10 - 80$
Priority	1
Radar Cross Section (m^2)	$0 - 20$
Process Noise Intensity	$0 - 50$

As with the static scenario, the delay inserted by the earliest deadline first scheduler can provide insight into the ability of the mechanism to allocate resource. Fig. 7.10 shows the delay inserted by the scheduler for CDAPS and RBPS. This has been produced by generating simulations with identical target parameters with resource availability of 10%. It can be seen that for RBPS the delay changes as the target scenario evolves, however, for CDAPS the delay is always balanced. This caused the quality of each task for RBPS to change unpredictably as the scenario evolves, whereas CDAPS produces predictable behaviour by selecting parameters given the dynamic competition for resource. In this example the delay for RBPS has an uncontrolled effect on the angular estimation error, whereas CDAPS adjusts the parameter selection to ensure best tracking accuracy is achieved given the dynamic target number. This further demonstrates the ability of CDAPS to degrade gracefully.
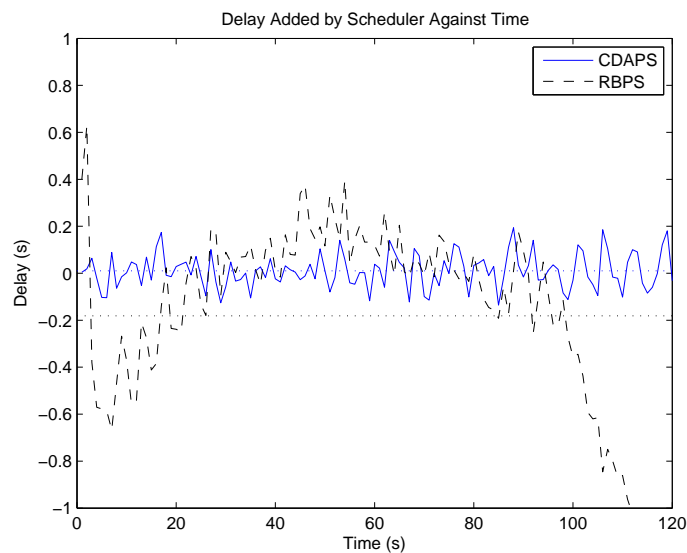

Figure 7.10: Time delay inserted by the scheduler for CDAPS and RBPS with a dynamic target scenario.

As the number of targets change, the competition for resource also changes. This causes the CDAPS

market equilibrium price, which represents the global valuation of the resource, to vary accordingly. Fig. 7.11 shows the transaction prices in the simulation over time for the same simulations used to generate the previous scheduler time delays. It can be seen that the market equilibrium price increases as targets arrive, and reaches its maximum when the competition for resource is greatest. This relationship between transaction price and competition for resource is in keeping with the market paradigm described in Sec. 7.1.1 and demonstrates that the mechanism is functioning correctly.

The market equilibrium price also has an interpretation through the Karush-Kuhn-Tucker conditions. The market equilibrium price represents the constant gradient in resource-utility space which is maintained for all tasks and so determines the individual parameter selection for each varying task. As the market equilibrium price increases, each task is forced to select parameters at a higher gradient, which is towards the lower resource region in resource-utility space. So, the requirement for resource for arriving targets is met by taking resource away from the tasks who lose the least amount of utility per unit resource. As targets become inactive the resource is purchased by targets who gain the most utility per unit resource. This shows how CDAPS dynamically adjusts the parameter selection to globally maximise utility production which for this example produces the optimal parameter selection for accurate tracking.

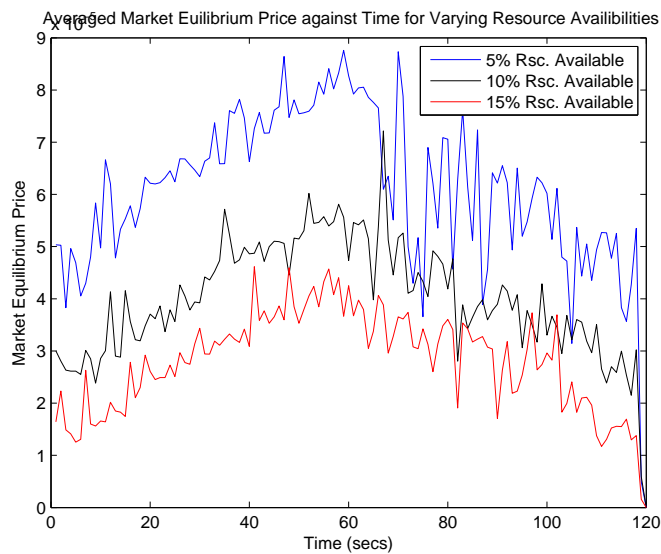Figure 7.11: Market equilibrium prices for CDAPS with a dynamic target scenario.

These simulations of dynamic tracking control scenarios have further demonstrated graceful degradation of CDAPS, as well its ability to adapt to a dynamic environment in a timely fashion to globally optimise utility production and hence improve resource allocation performance in contrast to a conventional rule based approach.

7.4.1.3   Task Priority

Task priority is an essential aspect of resource management which reflects the fact that different tasks have differing importance. The priority is typically a value which represents the tasks entitlement to resource relative to other tasks. Recent priority assignment methods [Miranda et al., 2007b] provide a continuous priority value in contrast to previous methods which simply rank based on task type. However, this is only useful if the allocation mechanism effectively translates the priority into behaviour.

The selection of parameters should fundamentally depend on the required quality of the task, however, when the radar becomes overloaded priority determines which tasks are degraded in quality. A typical implementation of task priority in a earliest deadline first scheduler only allows tasks to be delayed by tasks with a higher priority, which means high priority tasks fully meet the requirements of their rules. CDAPS incorporates priority by weighting the utility function, shown in Eq. 7.3, which scales the quality of the task depending on the priority. Hence, a task with twice the priority of another task, is able to produce half the increase in utility per unit resource and so a higher quality for an equal resource consumption.

Simulations have been produced to analyse the affect of priority assignment. The uncontrollable target environmental parameters used for the simulations are listed in Table 7.6, the fixed radar parameters are listed in Table 7.3. Given this target scenario it was desired to optimise tracking performance across all targets, in terms of predicted angular estimation error. Covariance analysis was used to determine the estimation error and the target dynamic was assumed matched to the tracker which uses a continuous white noise jerk model. The agents were generated in the radar resource management testbed described in Sec. 5.4.

Table 7.6: Target environmental parameters for priority simulation

Parameter	Value
Number of Targets	300
Azimuth Region (°)	± 45
Range (km)	$10 - 80$
Priority	1 or 2
Radar Cross Section (m^2)	$0 - 20$
Process Noise Intensity	$0 - 50$

Fig. 7.12 shows the angular estimation error for CDAPS and RBPS against resource availability for priorities values of 1 or 2 where the higher the value the greater the priority ranking. This has been produced by generating simulations with identical target parameters for a resource availability range between $0\% - 20\%$ which is synthesised by extending each radar dwell length. It can be seen that RBPS gives all resource to the higher priority tracks until the required quality defined by their rules are met which is marked by the dashed line, after this point resource is given to the lower priority tasks. In

contrast CDAPS controls the quality of each task according to the priority. This shows that CDAPS utilises the information in the priority assignment for enhanced control.
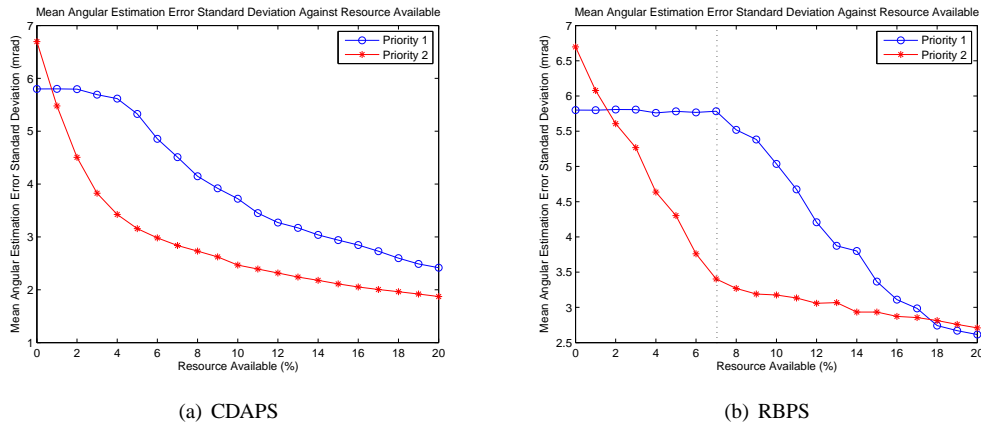

(a) CDAPS　　　　　　　　　　　　　　　(b) RBPS

Figure 7.12: Mean track angular estimation error standard deviation for CDAPS and RBPS with differing task priorities

Mission critical tasks may require fixed parameters which should not be degraded under any circumstance. In this instance these mission critical tasks can be excluded from the CDAPS algorithm. The rest of the architecture can remain unchanged and providing the total resource held by all agents in CDAPS represent the load available, accounting for the consumption of the excluded tasks, the time balance will remain balanced.

These simulations have shown how CDAPS is able to incorporate the value of the priority to directly control the quality of the radar tasks, whereas conventional methods do not.

7.4.1.4　Comparison to Sequential First Price Mechanism

A final tracking control simulation has been produced to compare the difference in performance between the three SFPARM types from Chapter 6 to CDAPS. This comparison is made to determine whether CDAPS adequately answers the issues related to the SFPARM mechanism, which were the computation and quality of the allocation.

The scenario consisted of 300 targets requiring tracking with uncontrollable environmental parameters as listed in Table 7.4. The target environmental parameters remained static over the duration of the simulation. Given this target scenario it was desired to optimise tracking performance across all targets, in terms of predicted angular estimation error. Covariance analysis was used to determine the estimation error and the target dynamic was assumed matched to the tracker using a continuous white noise jerk model. The agents were generated in the radar resource management testbed described in Sec. 5.4.

Fig. 7.13 shows the mean utility for the target tracks for CDAPS and the three SFPARM types from Chapter 6. This was produced by running simulations of $120s$ duration for each resource availability, with resource availability synthesised by lengthening each radar dwell. It can be seen that CDAPS

significantly outperforms the SFPARM variants. This is a result of CDAPS globally optimising utility production over a continuous time horizon, in contrast to the SFPARMs which produce local and myopic optimisation.
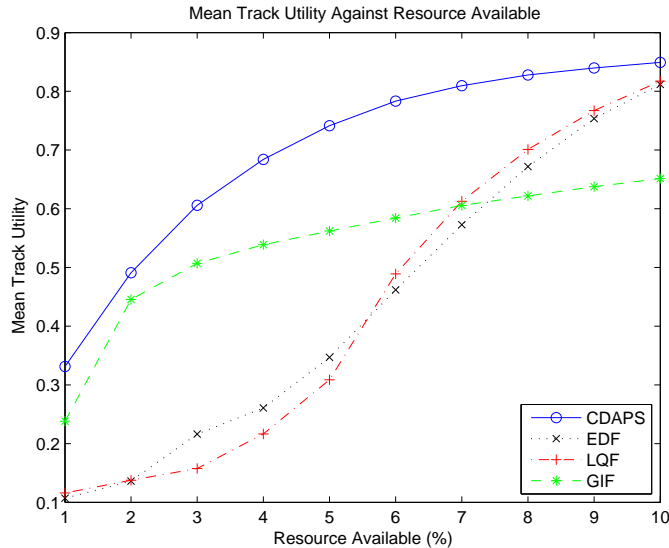

Figure 7.13: Comparison of mean utility for CDAPS and SFPARM types

By considering task quality over a continuous time horizon and optimising global utility production through satisfying the KKT conditions CDAPS resolves the issues associated with the quality of the SFPARM mechanism. Also, the computational is manageable as demonstrated by CDAPS executing in real time on a personal laptop.

7.4.2 Surveillance Control

A multifunction radar is also required to produce intelligent surveillance behaviour which can adapt to a variety of requirements in dynamic and uncertain environments. The suitability of CDAPS to achieve this has been assessed through simulation using the resource management testbed described in Sec. 5.4. In the simulation it was required for the resource manager to select parameters to perform the long range surveillance function, as described in Sec. 4.1.1.5 and Sec. 4.1.1 for a multifunction system. The requirement was to survey a $\pm45°$ azimuth and $0 - 5°$ elevation region centred on the antenna boresight. Using a $b_w = 1.5°$ beamwidth and a triangular lattice with $0.9b_w$ spacing produced 301 beam positions with each beam position represented by an agent. A further requirement was that the revisit interval must not exceed 16 seconds.

As a basis of comparison, a simple rule based parameter selection (RBPS) algorithm was generated. The rules produced a fixed surveillance pattern which requests an 8 second revisit interval and a constant energy dwell in each beam position. These parameters extracted from Butler [1998] are considered typical of a multifunction system such as MESAR, ARTIST or SAMPSON. An earliest deadline first

scheduler Blackman and Popoli [1999] was used to form the radar timeline for both CDAPS and RBPS.

7.4.2.1 Static Scenario

The CDAPS algorithm was first assessed for a static scenario where it was desired to perform long range surveillance against expected targets with a radar cross section of $10m^2$ and a radial velocity of $300ms^{-1}$. Figure 7.16(a) shows the resulting cumulative detection range performance across azimuth angle in the lowest elevation plane. As the RBPS method compensates for the loss of gain from scanning off-boresight by lengthening the coherent dwell times at greater angle, constant energy dwells across azimuth are produced. It can be seen in the figure that by selecting a fixed revisit interval and constant energy dwell, RBPS maintains a constant cumulative detection range performance across azimuth. In contrast, CDAPS extends the cumulative detection range on the boresight where maximum gain is available whilst degrading performance to an adequate level in the off-boresight angles which maximises the utilisation of the finite resource. The cumulative detection range averaged over azimuth angle is improved when using the CDAPS algorithm, which relates to an improvement in task quality.
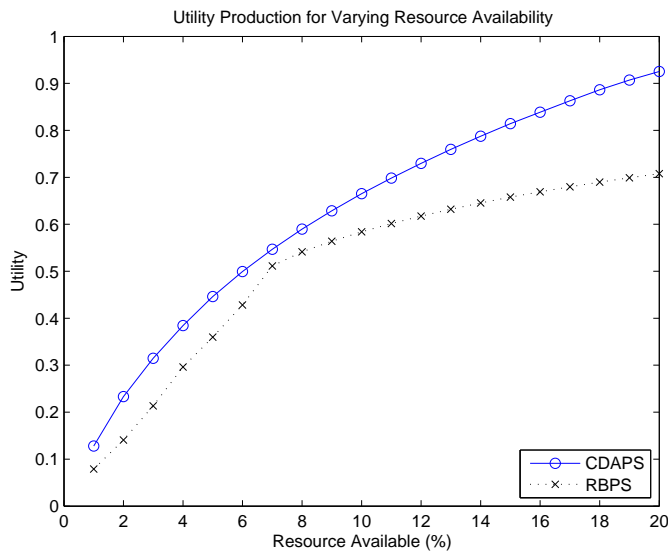
Figure 7.14: Average utility per task for CDAPS and RBPS allocation in a static environment

Allocation performance can be assessed in terms of utility which is a mapping from quality, i.e. cumulative detection range, as it describes what is required from the task. Fig. 7.14 shows the average utility per task against resource availability. It can be seen that, as expected by the KKT conditions, CDAPS outperforms the locally optimised RBPS. At 7% resource availability RBPS performs closest to CDAPS as the rules are matched to the available resource and all beam positions are maintained and requested dwells are not delayed by the scheduler. For RBPS below 7% tasks are delayed or dropped which has an adverse effect on the average utility, and so the quality of the tasks. For RBPS above 7% the excess resource causes tasks to be scheduled early, which reduces the revisit interval. In all

cases CDAPS chooses the parameters which optimises the quality and so average utility, given the finite resource. Although the RBPS rules used here are simple and can be improved, complex rules can be difficult to develop and are not guaranteed to achieve global optimisation. CDAPS naturally balances the time budget whilst solving the global resource utilisation objective.

7.4.2.2 Dynamic Scenario

CDAPS was assessed using a dynamic scenario where the expected radar cross section, radial velocity and priority in different regions changes over time. This represents a more realistic non-uniform and dynamic requirement against which it is desired to optimise the cumulative detection range. The situation assessment changes occur at 30 sec., 60 sec. and 90 sec. for the different regions depicted in Fig. 7.15 with parameters detailed in Tab. 7.7.
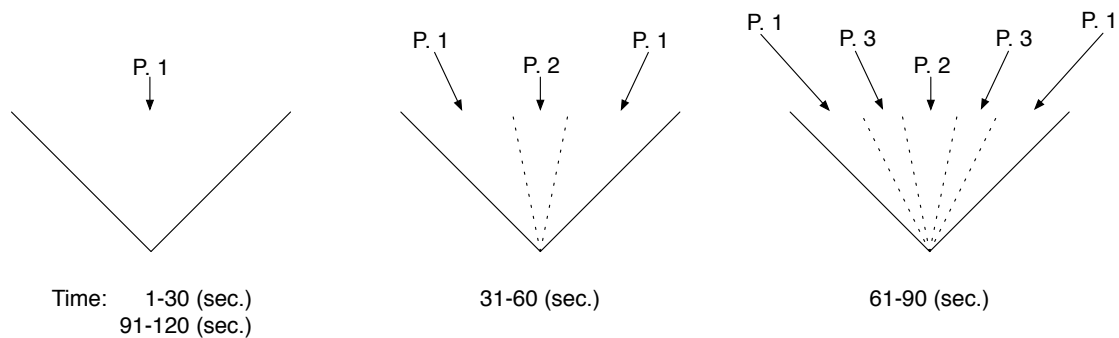
Figure 7.15: Dynamic scenario over dynamic simulation

Table 7.7: Task parameters in simulation for dynamic environment

Region	Radial Velocity (ms^{-1})	RCS (m^2)	Priority
P1	300	20	1
P2	800	1	3
P3	500	10	2

Fig. 7.16 shows the cumulative detection range performance across azimuth angle in the lowest elevation plane at varying stages in time. It can be seen in Fig. 7.16(b) that the increase in expected target velocity and reduction in expected radar cross section reduces the performance of the fixed RBPS allocation where the threatening targets are expected. This produces a degradation in task quality in the high priority region which is undesirable. CDAPS however adjusts the selection of task parameters to respond to the changing environment by improving performance in the threatening region. This can be further seen with the addition of another two medium threat regions at 60s shown in Fig. 7.16(c). Again, RBPS does not optimise task quality and so performance is degraded, whereas CDAPS reacts to the changing priority and expected target parameters to adjust parameters for improved performance in the

threatening regions. These figures show that the CDAPS algorithm is capable of effectively reallocating the finite resource in a dynamic environment to globally optimise task utility production. Previously [Mallett and Brennan, 1963; Mathews, 2005] cumulative detection range has been optimised for static uniform environments, but to the author's knowledge this application is the first optimisation for non-uniform dynamic environments.

(a) Cumulative detection range R_{90} (km) at 30 sec.

(b) Cumulative detection range R_{90} (km) at 60 sec.

(c) Cumulative detection range R_{90} (km) at 90 sec.

(d) Cumulative detection range R_{90} (km) at 120 sec.
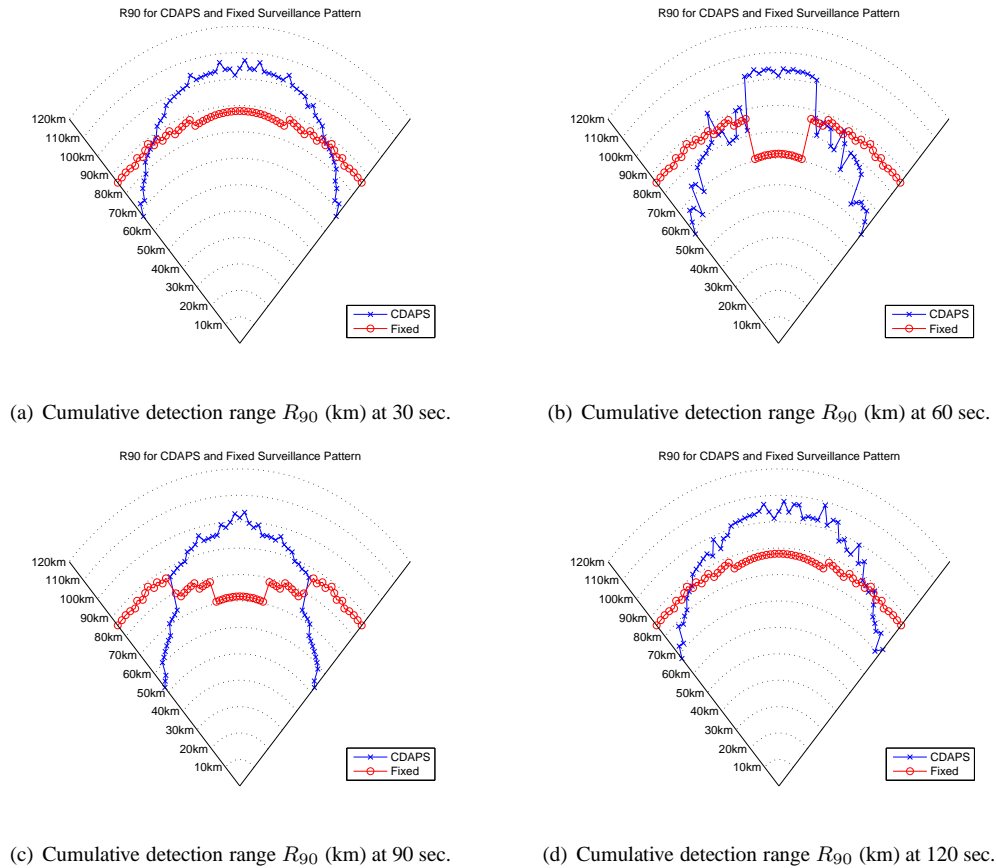
Figure 7.16: Cumulative detection range for CDAPS and RBPS allocation.

Fig. 7.17 shows the average utility per task against resource availability for the dynamic environment. A similar improvement is seen as with the static environment, across resource availability. However, the improvement of CDAPS over RBPS is slightly greater in the dynamic case than in the static case.

The market transaction price in CDAPS represents the current gradient in resource-utility space resulting from the current competitive equilibrium in the market. This gradient determines which parameters are selected from the concave majorant for each individual task and ensures the KKT conditions are satisfied. Fig. 7.18 shows the transaction prices in the auction over the simulation time, for varying resource availabilities. It can be seen that the transaction prices are higher when the resource availability is lower. A higher transaction price means parameters are selected from the lower resource region of each task's concave majorant in resource-utility space. It can also be seen that the transaction prices

Figure 7.17: Average utility per task for CDAPS and fixed (RBPS) allocation in a dynamic environment

change and settle on different equilibriums to respond to the changes in the scenario. These equilibrium changes represent the parameter selections being adjusted to respond to the changing environment.
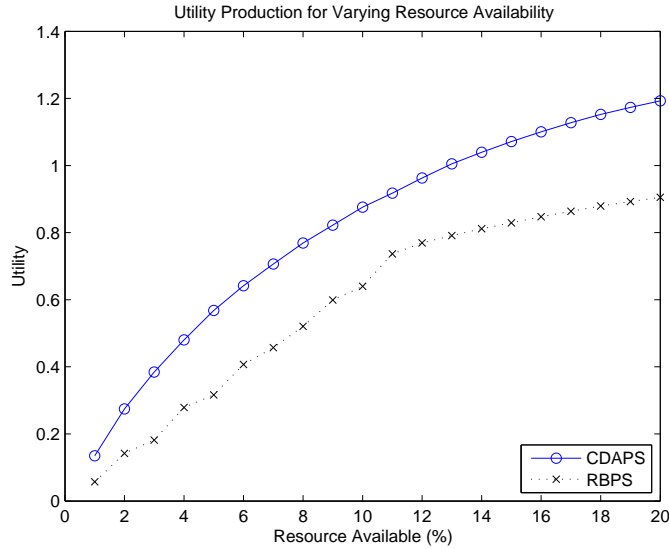
7.4.3 Discussion

It should be noted that the performance improvement shown in these simulations is not representative of an absolute improvement in performance. These simulations are heavily affected by the choice of utility function, the selected rules for RBPS, the models used and parameters which define the models. These results have been fairly arbitrarily chosen to be indicative of a reasonable performance gain, and are by no means selected as the best performance improvement achievable. However, although the performance improvement may vary it is important to note that through satisfying the KKT conditions CDAPS will continually outperform the locally optimised RBPS, which is evident in all the simulations in this chapter.

The CDAPS algorithm is implemented in real time and all simulations have been produced on a personal laptop. CDAPS is computationally efficient as the continuous nature of the mechanism allows adjustments to be made to the allocation as and when required, without having to recompute the entire allocation. Additionally, only points local to the current parameter selection are evaluated, instead of the entire concave majorant as in Q-RAM. This combined effect reduces the search load and spreads it over time. Whilst reducing the computation burden, CDAPS is able to maintain rapid reaction to pop-up tasks due to the continuous nature of the mechanism. Both these aspects suggest CDAPS is realistic and suitable for application to existing multifunction radar control software.

The performance analysis given in this section has been completely produced through simulation. This is valid in this case because the mechanism operates at a high level after significant data processing and so the mechanism is purely responsible for maximise global utility production. The CDAPS
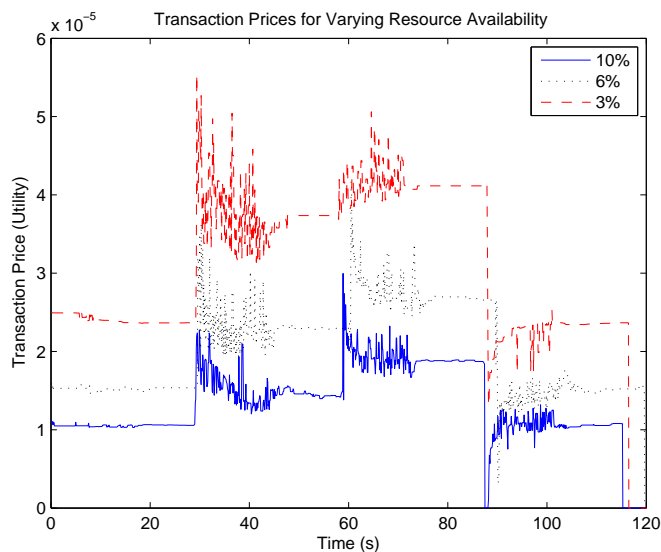
Figure 7.18: Transaction prices in CDAPS market for varying resource availability over simulation

algorithm has been shown to be superior in this respect than conventional rule based approaches. The performance improvement when applied to real data would be affected by how well the measures and models relate to the true situation, however, this is an aspect of performance for the measures and models and not for CDAPS. Hence it is not essential for CDAPS to be validated on real data.

7.4.3.1 Performance Issues and Drawbacks

To produce the simulations in this analysis, the CDAPS algorithm was implemented in real time for execution on a personal laptop. This demonstrates the efficiency of the algorithm but also provided insight into bottlenecks and performance issues which affect the algorithm. A key performance metric for the system is the reaction time, which is the time taken to generate, send and process an offer. It is desirable to keep the processing time involved in each of these steps to a minimum. To reduce the time to generate an offer it is desirable to minimise the computation required to evaluate each potential parameter selection. To minimise the time taken to send and process the offer messages it is desirable to minimise the number of message being passed in the system. This is achieved in the mechanism by restricting the price of the offer above a threshold for a bid and below a threshold for an ask. Given a finite computational resource, the response time will increase as the number of task agents increases. However, CDAPS was found to give a reaction time in the order of $1 - 10ms$ where as the Q-RAM [Hansen et al., 2006] approach is in the order of seconds.

When a large change occurs simultaneously for a number of tasks, it may be required for numerous tasks to sell or purchase a substantial amount of resource. This occurred in the long range surveillance function simulations when the scenario changed abruptly. In this case the time taken to the market to settle on the new competitive equilibrium, which is the optimal allocation, was longer than expected. Fu-

ture work can address how task agents can buy or sell larger resource volumes, which could be provided in a similar way to the upstairs market in the NYSE which specialises in larger resource volumes.

CDAPS is highly suitable for allocating resource to tasks such as surveillance and tracking as they can readily be described by numerous quality measures. However, there are tasks which need to be performed which do not lend themselves to be described by measures so easily. However, as these types of tasks are quite fixed, there is little flexibility for optimisation. It is suggested that a multi-layer resource manager would be required to accommodate such tasks.

A final limitation for exploration in further work is that the parameter selections are considered to be discrete. Extending the mechanism to allow continuous parameter selections could further improve the quality of the allocation, and find the optimal solution in the true sense [Irci et al., 2010].

7.5 Summary

The continuous double auction mechanism has been successfully applied in many real world resource allocation problems, such as the NYSE. The continuous double auction parameter selection (CDAPS) algorithm has been developed to select parameters for individual radar tasks in a multifunction radar, hence allocating the finite resource. The algorithm hosts a market mechanism which enables numerous localised agents representing radar tasks to solve the global resource utilisation optimisation problem.

Theoretical concepts indicate that CDAPS is able to produce close to the optimality of existing optimisation approaches to RRM with a fraction of the computational burden. By satisfying the KKT conditions the mechanism can be shown to tackle the global optimisation problem of maximising task quality subject to the resource constraint.

Realistic and complex simulations of surveillance and tracking scenarios have verified that the algorithm enables a worthwhile improvement in task performance and hence resource utilisation which continually outperforms a locally optimised rule based method. Results from the simulations have shown graceful degradation with adaptation to dynamic environments. To ensure feasibility for application to real systems, the prototype mechanism used for the simulations ran in real time.

As utility maps from any quality metric a wider variety of measures than considered here can be implemented. For example in tracking, resources can be allocated based on track existence. The algorithm has the potential for worthwhile extensions to sensor suites and networks, as well as relevance for many resource allocation applications.

Chapter 8

Conclusions

Multifunction radar resource management addresses how to effectively automate the allocation of the finite radar time-energy resource between numerous, potentially conflicting tasks that support multiple radar functions. This thesis has advanced the sensor management field through the application of autonomous agent systems to the multifunction radar resource management problem.

8.1 Summary

Multifunction radars are capable of supporting differing functions by utilising an agile beam and configuring the radar operation for each task, allowing the radar to tailor performance to various roles or applications. The multifunction radar system is controlled by the automated resource manager and hence overall performance is dictated by the resource managers ability to effectively adapt performance to a dynamic and uncertain environment. This thesis provided a thorough overview of the principles of multifunction radar operation. These principles determine the parameters which the automated resource manager must control. This is an important basis for the development of the algorithms in this thesis.

Through analysis of existing literature on radar resource management it was seen that existing methods of radar resource management do not adequately meet the requirements of emerging systems. Specifically it was identified that a resource allocation mechanism is required which is computationally light like rule based methods whilst producing near optimal solutions of the optimisation methods. To adhere to the notion of multi-functionality the mechanism is required to optimise a variety of requirements, with adequate consideration of the finite resource, to produce behaviour which adapts to dynamic and uncertain environments. This analysis also showed that information theoretic measures could be beneficially applied.

As a step towards the ultimate goal of the application of agent systems to multifunction radar resource management, this research derived and compared some measures and models which are suitable for allocating multifunction radar resource. The quality of the measures and models used are critical as they ultimately limit the quality of the resource allocation. This research concluded that information

theoretic measures as surrogate functions are useful for the optimisation of tasks in isolation, but are less useful for making higher level resource allocation decisions as information production is not the primary requirement of the radar system. It was subsequently concluded that as multifunction radar resource management inherently aims to optimise multiple functions, it is desirable to use as wide a variety of measures as possible, which must be accommodated by the mechanism. It was found that this can be achieved by defining utility functions which give the satisfaction associated with each point in quality space and allow a variety of quality metrics to be represented by a single utility measure. It was also found that when tracking with significant measurement origin uncertainty the resource allocation can be improved by using the modified Riccati equation.

As a result of this research a novel agent based radar resource manager was developed which exploits the use of a mixture of objects and agents with functionality provided by the Java Agent Development Framework. This development included the generation of simulated measurement data to stimulate the subsequent agent systems. This agent based radar resource manager is believed to be the first of its kind and knowledge was gained on how to design an agent based resource manager which allows rapid expansion and maximum code re-use. This resource manager was then used to produce results on two mechanisms, the sequential first price and continuous double auctions.

The sequential first price auction was applied to the radar resource management problem to develop the novel Sequential First Price Auction Resource Management (SFPARM) algorithm and gave significant insight into multifunction radar resource manager design. It was found that a best-first scheduler based on delay gives preference to fluid constraints, information gives preference to 'bright' targets with high SNR and quality gives preference to tasks which degrade in quality the least. These preferences are undesirable as they do not adequately address the global resource management objective. An EDF has manageable computation as the passing of time is the same for all tasks. However, when the criteria is switched to a measure which does not pass equally, the computation is significantly increased and values need to be calculated frequently. It was confirmed that information theoretic measures do not adequately describe the requirements of the radar and so are not a suitable choice of objective function. This research found that myopic allocation performs poorly and so optimising over an extended time horizon is preferred. Despite these numerous conclusions it was found that SFPARM had a high computational burden and inadequate performance.

The continuous double auction mechanism was applied to the radar resource management problem to produce the novel Continuous Double Auction Parameter Selection (CDAPS) algorithm. The CDAPS algorithm hosts a market mechanism where agents representing radar tasks can trade radar loading enabling parameters to be selected for individual radar tasks, hence allocating the finite resource. Crucially, satisfying the Karush-Kuhn-Tucker conditions ensures that the CDAPS algorithm converges on the global optimal solution, in that maximum utility is produced from the finite resource which equates to

the best quality for a set of tasks. This ensures that the numerous localised agents are solving the global resource utilisation optimisation problem. Realistic and complex simulations of surveillance and tracking scenarios have verified that the algorithm enables a worthwhile improvement in task performance and hence resource utilisation which continually outperforms a locally optimised rule based method. Results from the simulations have shown graceful degradation with adaptation to dynamic environments. To ensure feasibility for application to real systems, the prototype mechanism used for the simulations ran in real time. The successful development of this algorithm has opened up numerous potential extensions for future work.

8.2 Key Research Achievements and Contributions

The research conducted for this thesis has explored many aspects of the emerging field of radar resource management. The key research achievements which have advanced this field are believed to be:

- Information theoretic measures for multifunction radar resource management have been derived and developed for estimation and discrimination problems. Although previously applied to sensor management this is believed to be the first assessment of suitability for multifunction radar resource management. This study found that information theoretic measures, contrary to assertions in the literature, do not adequately describe the requirements of the radar and so are less suitable as objective functions for multifunction radar resource management.

- Tracking control in clutter has also been examined, it has been shown that measurement origin uncertainty strongly affects the parameter selection choice to achieve minimum track loading and also that the maximum measurement mutual information production peak is dependent on revisit interval and false target density. The Modified Riccati Equation is shown to improve tracking allocation performance under significant measurement origin uncertainty.

- An agent based multifunction radar resource manager using the JADE framework has been developed. This is the first application of agents to multifunction radar resource management. The research has delivered a suitable architecture which allows rapid extension and maximises code reuse which can be used as a basis for future research.

- The sequential first price sealed bid auction mechanism has been applied to the multifunction radar resource management problem for the first time. This has resulted in the development of two novel best first schedulers, the lowest quality first and greatest information first. Through comparison to existing radar resource management techniques the research suggested a number of novel conclusions and guidelines for radar resource manager design.

- The novel application of the continuous double auction mechanism to the multifunction radar re-

source management problem has led to the development and assessment of the Continuous Double Auction Parameter Selection (CDAPS) algorithm. This algorithm improves the current state of the art by producing a quality allocation with low computational burden. As an economic paradigm, CDAPS is a tangible step to developing cognitive sensor signal processing techniques.

In summary this work represents the first application of agent systems to multifunction radar resource management and has significantly advanced the knowledge in this field.

8.3 Future Work and Extensions

This thesis has delivered the key achievements listed in the previous section. However, it has also opened up many avenues for future work. Much of the future work could involve the development of the successful CDAPS algorithm.

The measures and models considered in Chapter 4 are by no means exhaustive. Different models could be used, of particular relevance would be a model similar to the Van Keuk model which incorporates measurement origin uncertainty. Additional tracking measures can be considered such as track continuity, however it may be difficult to accurately relate these quality measures to a definite resource loading.

The CDAPS algorithm has demonstrated improved performance over conventional techniques for the accurate tracking and long range surveillance functions. These functions, however pertinent, are only two examples of the functions that need to be performed. This work could be extended by modelling the resource loading, quality and utility of additional functions such as self protect or medium range search. The CDAPS algorithm also allows for new measures, such as track existence, to be allocated resource within the same mechanism. This capability, which can be realised through future work allows for new functionality to be added to the system. Also different target geometries and scenarios can be considered, such as a requirement to defend a point spatially separated from the platform.

The CDAPS algorithm has yet to be applied to an application involving real data. Simulated data has been used in this thesis due to the difficulties in performance assessment outlined in Sec. 4.5. Real data could be used from a mechanically scanning system and posing a finite resource constraint, such as only allowing a few beam positions to be used from each scan. Another imposed resource constraint could be using a reduced number of PRFs from a fixed set of bursts. Ultimately, these resource constraints are not ideal and the application of the algorithm on a real MFR would provide definitive insight into how well the task measures and models work in reality. This is relevant as if the task measures and models are the greater bottleneck then improving the performance of the allocation mechanism will have less effect.

The implementation of the CDAPS algorithm in this thesis can be improved. As it is crucial that the mechanism operates quickly, ways in which the mechanism protocol can be changed to reduce communication overhead are of benefit. This could include developing methods which allow agents to trade

directly with each other without having to advertise in the workbook. Alternatively this could mean reducing the size of the workbook and minimising the conditions under which agents announce offers. Also as the market equilibrium convergence time is desired to be as short as possible, the mechanism could be adapted to include a separate facility for trading in large resource volumes, similar to the provision in the New York Stock Exchange.

As the CDAPS algorithm is distributed and decentralised it is inherently scalable. This makes both the algorithm and economic paradigms in general particularly suitable to be adapted for application to a variety of applications within the sensor management field such as the control of multiple UAVs. The abstraction for the next development of the work is clear; where resource allocation for single multifunction radar can be represented as an auction, the allocation for a sensor suite can be represented as a market and the allocation for sensor network can be represented as an economy.

Finally, as CDAPS is general it has the potential for application outside of the sensor management field, for any application requiring a finite resource to be allocated between conflicting tasks, such as in grid computing, factory automation or communication networks.

Bibliography

Y. Bar-Shalom, X. Rong Li, and T. Kirubarajan. *Estimation with Applications to Tracking and Navigation: Theory Algorithms and Software*. Wiley, 2001.

F. Barbaresco, J.C. Deltour, G. Desodt, B. Durand, T. Guenais, and C. Labreuche. Intelligent M3R radar time resources management: Advanced cognition, agility; autonomy capabilities. In *International Radar 09 Conference*, pages 1–6, Bordeaux, Oct. 2009.

A. Barbato and P. Giustiniani. An improved scheduling algorithm for a naval phased array radar. In *IEE International Radar Conference*, pages 42–45, Oct. 1992.

D.K. Barton. *Radar Systems Analysis and Modelling*. Artech House, rev. 1st. edition, 2004.

M.R. Bell. Information theory and radar waveform design. *IEEE Transactions on Information Theory*, 39:1578–1597, 1993.

S.G. Bier, P.L. Rothman, and R.A. Manske. Intelligent sensor management for beyond visual range air-to-air combat. In *IEEE National Aerospace and Electronics Conference*, pages 264–269, 1988.

E.R. Billam. Parameter optimisation in phased array radar. In *International Radar 92 Conference*, pages 34–37, 1992.

E.R. Billam. The problem of time in phased array radar. In *International Radar 97 Conference*, pages 563–575, Oct. 1997.

D.R. Billetter. *Multifunction Array Radar Design*. Artech House, July 1989.

S. Blackman. *Multiple-Target Tracking with Radar Applications*. Artech House, Dec. 1986.

S. Blackman and R. Popoli. *Design and Analysis of Modern Tracking Systems*. Artech House, 1999.

W.D. Blair, G.A. Watson, and S.A. Hoffman. Benchmark problem for beam pointing control of phased array radar against maneuvering targets. In *American Control Conference*, volume 2, pages 2071–2075, 1994.

W.D. Blair, G.A. Watson, G.L. Gentry, and S.A. Hoffman. Benchmark problem for beam pointing control of phased array radar against maneuvering targets in the presence of ECM and false alarms. In *Proc. American Control Conference*, volume 4, pages 2601–2605, 1995.

W.D. Blair, G.A. Watson, T. Kirubarajan, and Y. Bar-Shalom. Benchmark for radar allocation and tracking in ECM. *IEEE Transactions on Aerospace and Electronic Systems*, 34(4):1097–1114, Oct. 1998.

D. Blatt and A. O Hero. Optimal sensor scheduling via classification reduction of policy search (crops). In *Conference on Automated Planning and Scheduling*, Cumbria, UK, June 2006.

Y. Boers, H. Driessen, and J. Zwaga. Adaptive MFR parameter control: fixed against variable probabilities of detection. *IET Radar, Sonar and Navigation*, 153(1):2–6, 2006.

J.M. Butler. *Tracking and Control in Multi-Function Radar*. Ph.D. Thesis, University College London, 1998.

J.M. Butler, A.R. Moore, and H.D. Griffiths. Resource management for a rotating multi-function radar. In *International 97 Radar Conference*, pages 568–572, Oct. 1997.

D.A. Castanon. Approximate dynamic programming for sensor management. In *Proc. 36th IEEE Conference on Decision and Control*, volume 2, pages 1202–1207, 1997.

D. Cochran, S. Suvorova, S. Howard, and B. Moran. Waveform libraries. *IEEE Signal Processing Magazine*, 26(1):12 –21, 2009.

T.M. Cover and J.A. Thomas. *Elements of Information Theory*. Wiley, 2nd edition edition, 2006.

E. Daeipour, Y. Bar-Shalom, and X. Li. Adaptive beam pointing control of a phased array radar using an IMM estimator. In *American Control Conference*, pages 2093–2097, Baltimore, USA, June 1994.

R.A. Dana and D. Moraitis. Probability of detecting a swerling I target on two correlated observations. *IEEE Transactions on Aerospace and Electronic Systems*, 17(5):727–730, 1981.

G. Davidson. Cooperation between tracking and radar resource management. In *IET International Conference on Radar Systems*, pages 1 –4, Edinburgh, UK., Oct. 2007.

S. Eguchi and J. Copas. Interpreting Kullback-Leibler divergence with the Neyman-Pearson lemma. *Journal Multivariate Analysis*, 97:2034–2040, Oct. 2006.

D.G. Finch. A detailed simulation of a 'closed loop' multi function radar. In *IEE Colloquium on Radar System Modelling*, pages 9/1–9/7, 1998.

T. Fortmann, Y. Bar-Shalom, M. Scheffe, and S. Gelfand. Detection thresholds for tracking in clutter–a connection between estimation and signal processing. *IEEE Transactions on Automatic Control*, 30 (3):221–229, 1985.

S. French. *Sequencing and Scheduling: An Introduction to the Mathematics of the Job-shop*. Ellis Horwood Ltd , Publisher, 1981.

S. Ghosh, R. Rajkumar, J. Hansen, and J. Lehoczky. Scalable resource allocation for multi-processor QoS optimization. In *Proc. 23rd International Conference on Distributed Computing Systems*, pages 174–183, Providence, USA, May 2003.

S. Ghosh, J. Hansen, R. Rajkumar, and J. Lehoczky. Integrated resource management and scheduling with multi-resource constraints. In *IEEE International Real-Time Systems Symposium*, pages 12–22, Lisbon, Portugal, Dec. 2004.

B. Gillespie, E. Hughes, and M. Lewis. Scan scheduling of multi-function phased array radars using heuristic techniques. In *IEEE International Radar Conference*, pages 513–518, Arlington, Virginia, May 2005.

W.H. Gilson. Minimum power requirements of tracking. In *IEEE International Radar Conference*, pages 417–421, May 1990.

S. Gjerstad and J. Dickhaut. Price formation in double auctions. *Games and Economic Behavior*, 22: 1–29, 1998.

D. K. Gode and S. Sunder. Allocative efficiency of markets with zero-intelligence traders: Market as a partial substitute for individual rationality. *Journal of Political Economy*, 101(119137):119–137, 1993.

J.R. Guerci and E.J. Baranoski. Knowledge-aided adaptive radar at DARPA: an overview. *IEEE Signal Processing Magazine*, 23(1):41 – 50, Jan. 2006.

J. Hansen, R. Rajkumar, J. Lehoczky, and S. Ghosh. Resource management for radar tracking. In *IEEE International Radar Conference*, pages 140–147, Verona, NY., May 2006.

S. Haykin. Cognitive radar: a way of the future. *IEEE Signal Processing Magazine*, 23(1):30–40, 2006.

Y. He and K.P. Chong. Sensor scheduling for target tracking in sensor networks. In *Proc. IEEE Conference on Decision and Control*, pages 743–748, Bahamas, Dec. 2004.

Y. He and K.P. Chong. Sensor scheduling for target tracking: A monte carlo sampling approach. *Digital Signal Processing*, 16(5):533–545, 2006.

A.O. Hero, D.A. Castanon, D.Cochran, and K. Kastella. *Foundations and Applications of Sensor Management*. Springer, 2007.

K.J. Hintz. A measure of the information gain attributable to cueing. *IEEE Transactions on Systems, Man and Cybernetics*, 21(2):434–442, 1991.

K.J. Hintz and E.S. McVey. Multi-process constrained estimation. *IEEE Transactions on Systems, Man and Cybernetics*, 21(1):237–244, 1991.

J. Holloway. Design considerations for adaptive active phased-array 'multifunction' radars. *Electronics & Communication Engineering Journal*, 13(6):277–288, 2001.

A.G. Huizing and A.A.F. Bloemen. An efficient scheduling algorithm for a multifunction radar. In *IEEE International Symposium on Phased Array Systems and Technology*, pages 359–364, Oct. 1996.

A. Irci, A. Saranli, and B. Baykal. On optimal resource allocation in multifunction radar systems. In *IEEE International Radar Conference*, pages 684–691, Verona, NY., May 2006.

A. Irci, A. Saranli, and B. Baykal. Study on Q-RAM and feasible directions based methods for resource management in phased array radar systems. *IEEE Transactions on Aerospace and Electronic Systems*, 46(4):1848 –1864, Oct. 2010.

A. Izquierdo-Fuente and J.R. Casar-Corredera. Approach to multifunction radar scheduling simulation. In *Proc. IEEE Telesystems Conference*, pages 67–70, May 1994a.

A. Izquierdo-Fuente and J.R. Casar-Corredera. Optimal radar pulse scheduling using a neural network. In *IEEE International Conference on Neural Networks*, pages 4588–4591, Orlando, USA, June 1994b.

N.R. Jennings and M.J. Wooldridge. *Agent Technology: Foundations, Applictions and Markets*. Springer, 1998.

R. Rajkumar J.P. Hansen, S. Ghosh and J. Lehoczky. Resource management of highly configurable tasks. In *Proc. 18th International Parallel and Distributed Processing Symposium*, page 116, New Mexico, April 2004.

K. Kastella. Discrimination gain to optimize detection and classification. *IEEE Transactions on Systems, Man and Cybernetics, Part A: Systems and Humans*, 27(1):112–116, 1997.

D.J. Kershaw and R.J. Evans. Optimal waveform selection for tracking systems. *IEEE Transaction on Information Theory*, 40(5):1536–1550, 1994.

D.J. Kershaw and R.J. Evans. A contribution to performance prediction for probabilistic data association tracking filters. *IEEE Transactions on Aerospace and Electronic Systems*, 32(3):1143–1148, 1996.

D.J. Kershaw and R.J. Evans. Waveform selective probabilistic data association. *IEEE Transactions on Aerospace and Electronic Systems*, 33(4):1180–1188, 1997.

T. Kirubarajan, Y. Bar-Shalom, W.D. Blair, and G.A. Watson. IMMPDAF for radar management and tracking benchmark with ECM. *IEEE Transactions on Aerospace and Electronic Systems*, 34(4): 1115–1134, Oct 1998.

W. Komorniczak and J. Pietrasinski. Selected problems of MFR resources management. In *Proc. 3rd International Conference on Information Fusion*, pages 3–8, Paris, France, July 2000.

W. Komorniczak, T. Kuczerski, and J.F. Pietrasinski. The priority assignment for detected targets in multi-function radar. In *13th International Conference on Microwaves, Radar and Wireless Communications*, pages 244–247, 2000.

C. Kreucher and A.O. Hero. Monte carlo methods for sensor management in target tracking. In *IEEE Nonlinear Statistical Signal Processing Workshop*, pages 232 –237, Sept. 2006.

C. Kreucher, A.O. Hero, K. Kastella, and D. Chang. Efficient methods of non-myopic sensor management for multitarget tracking. In *43rd IEEE Conference on Decision and Control*, pages 722–727, Bahamas, Dec. 2004.

C. Kreucher, III Hero, A.O., and K. Kastella. A comparison of task driven and information driven sensor management for target tracking. In *44th IEEE Conference on Decision and Control*, pages 4004–4009, Dec. 2005a.

C. Kreucher, K. Kastella, and A.O. Hero. Multitarget tracking using the joint multitarget probability density. *IEEE Transactions on Aerospace and Electronic Systems*, 41(4):1396–1414, 2005b.

C. Kreucher, K. Kastella, and A.O. Hero. Sensor management using an active sensing approach. *Signal Processing*, 85(3):607–624, 2005c.

J.D. Mallett and L.E. Brennan. Cumulative probability of detection for targets approaching a uniformly scanning search radar. *Proceedings of the IEEE*, 51(4):596 – 601, April 1963.

J. I. Marcum. A statistical theory of target detection by pulsed radar. Technical report, RAND Corporation, Santa Monica, CA, 1947.

J. I. Marcum. A statistical theory of target detection by pulsed radar: Mathematical appendix. Technical report, RAND Corporation, Santa Monica, CA, 1948.

B.D. Mathews. Optimal dwell time for approach-warning radar. *IEEE Transactions on Aerospace and Electronic Systems*, 41(2):723 – 728, April 2005. ISSN 0018-9251.

S. Miranda, C. Baker, K. Woodbridge, and H. Griffiths. Knowledge-based resource management for multifunction radar: a look at scheduling and task prioritization. *IEEE Signal Processing Magazine*, 23(1):66–76, 2006.

S.L.C. Miranda. *Resource Management in Multifunction Array Radar*. PhD thesis, University College London, 2004.

S.L.C. Miranda, C.J. Baker, K. Woodbridge, and H.D. Griffiths. Comparison of scheduling algorithms for multifunction radar. *IET Radar, Sonar & Navigation*, 1(6):414–424, 2007a.

S.L.C. Miranda, C.J. Baker, K. Woodbridge, and H.D. Griffiths. Fuzzy logic approach for prioritisation of radar tasks and sectors of surveillance in multifunction radar. *IET Radar, Sonar & Navigation*, 1 (2):131–141, 2007b.

G. Morris and L. Harkness. *Airborne Pulsed Doppler Radar*. Artech House, 2nd revised edition edition, Jan. 1996.

S. Musick and R. Malhotra. Chasing the elusive sensor manager. In *Proc. IEEE National Aerospace and Electronics Conference*, pages 606–613, 1994.

D. Musicki, R. Evans, and S. Stankovic. Integrated probabilistic data association. *IEEE Transactions on Automatic Control*, 39:1237–1241, 1994.

J.M. Nash. Optimal allocation of tracking resources. In *IEEE Conference on Decision and Control*, volume 16, pages 1177–1180, 1977.

A. Nelander and D. Stromberg. Energy management in multiple-beam phased array radars. In *International Radar 97 Conference*, pages 553–557, Edinburgh, UK, Oct. 1997.

G. W. Ng and K. H. Ng. Sensor management - what, why and how. *Information Fusion*, 1(2):67–75, 2000.

D.O. North. An analysis of the factors which determine signal noise discrimination in pulsed carrier systems. *Proceedings of the IEEE*, 51(7):1016–1027, July 1963.

S.P. Noyes. Calculation of next time for track update in the MESAR phased array radar. In *IEE Colloquium on Target Tracking and Data Fusion*, pages 2/1–2/7, 1998.

A. J. Orman, C. N. Potts, A. K. Shahani, and A. R. Moore. Scheduling for a multifunction phased array radar system. *European Journal of Operational Research*, 90(1):13–25, 1996.

T. R. Payne, E. David, N. R. Jennings, and M. Sharifi. Auction mechanisms for efficient advertisement selection on public display. In *European Conference on Artificial Intelligence*, pages 285–289, Trentino, 2006.

R.J.E. Pollard. Adaptive radar system control. In *IEE Colloquium on Real-Time Management of Adaptive Radar Systems*, pages 4/1–4/4, 1990.

R. Popoli and S. Blackman. Expert system allocation for the electronically scanned antenna radar. In *American Control Conference*, pages 1821–1826, Minneapolis, USA, June 1987.

L. Powis, D.A. Lambert, P.W. Sarunic, A. Jakobsson, and M. Stone. Adaptive radar control of multi-function radars using artificial intelligence. In *International Radar 92 Conference*, pages 426–429, 1992.

B.T. Rhatigan, P.R. Kalata, and T.A. Chmielewski. An alpha-beta target tracking approach to the bench-mark tracking problem. In *American Control Conference*, pages 2076 – 2080, Baltimore, USA, June 1994.

B. Ristic and M.L Hernandez. *Tracking Systems*. IEEE Radar Conference Tutorial, 2008.

B. Ristic, S. Arulampalam, and N. Gordon. *Beyond the Kalman Filter*. Artec House, 2004.

A. Rogers, E. David, T.R. Payne, and N.R. Jennings. An advanced bidding agent for advertisement selection on public displays. In *Sixth International Joint Conference on Autonomous Agents and Multiagent Systems*, pages 251–258, Honolulu, Hawaii, USA, May 2007.

S. Russell and P. Norvig. *Artificial Intelligence: A Modern Approach*. Prentice Hall, 3rd edition, Dec. 2009.

S. Sabatini and M. Tarantino. *Multifunction Array Radar: System Design and Analysis*. Artech House, 1994.

W.W. Schmaedeke and K.D. Kastella. Information-based sensor management and IMMKF. In *Signal and Data Processing of Small Targets*, pages 390–401, Orlando, FL, USA, 1998.

R.A. Singer. Estimating optimal tracking filter performance for manned maneuvering targets. *IEEE Transactions on Aerospace and Electronic Systems*, AES-6(4):473 –483, July 1970.

S.P. Sira, Ying Li, A. Papandreou-Suppappola, D. Morrell, D. Cochran, and M. Rangaswamy. Waveform-agile sensing for tracking. *IEEE Signal Processing Magazine*, 26(1):53–64, 2009.

M. I. Skolnik, editor. *Radar Handbook*. McGraw-Hill, New York, 3rd ed edition, 2008.

V. L. Smith. An experimental study of competitive market behaviour. *Journal of Political Economy*, 70: 111–137, 1962.

T.L Song and D. Musicki. Target existence based resource allocation. *Signal Processing, IEEE Transactions on*, 58(9):4496–4506, Sept. 2010.

W.K. Stafford. Real time control of a multifunction electronically scanned adaptive radar (MESAR). In *IEE Colloquium on Real-Time Management of Adaptive Radar Systems*, pages 7/1–7/5, 1990.

W.K. Stafford. MESAR, sampson & radar technology for BMD. In *IEEE International Radar Conference*, pages 437–442, Boston, May 2007.

G.W. Stimson. *Introduction to Airborne Radar*. Raleigh, NC :SciTech Publishing, Inc, 2nd edition, 1998.

A.P. Stoffel. Heuristic energy management for active array multifunction radars. In *Proc. IEEE National Telesystems Conference*, pages 71–74, San Diego, May 1994.

D. Stromberg. Scheduling of track updates in phased array radars. In *Proc. IEEE International Radar Conference*, pages 214–219, Oct. 1996.

D. Stromberg and P. Grahn. Scheduling of tasks in phased array radar. In *IEEE International Symposium on Phased Array Systems and Technology*, pages 318–321, Oct. 1996.

S. Suvorova, S.D. Howard, W. Moran, and R.J. Evans. Waveform libraries for radar tracking applications: Maneuvering targets. In *40th Annual Conference on Information Sciences and Systems*, pages 1424–1428, Princeton, New Jersey, March 2006.

H. Van Trees. *Detection, Estimation, and Modulation Theory, Part III*. Wiley, 2001.

G.V. Trunk, J.D. Wilson, and P.K. Hughes. Phased array parameter optimization for low-altitude targets. In *IEEE International Radar Conference*, pages 196–200, May 1995.

G. van Keuk and S.S. Blackman. On phased-array radar tracking and parameter control. *IEEE Transactions on Aerospace and Electronic Systems*, 29(1):186–194, Jan 1993.

N.A. Vaughan. MFR surveillance optimisation-a pragmatic systems engineers' approach. In *IEE Multifunction Radar and Sonar Sensor Management Techniques*, pages 6/1–6/3, 2001.

W. Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *The Journal of Finance*, 16 (1):8–37, 1961.

J. Vidal. *Fundamentals of Multiagent Systems*. 1st edition, 2007.

M.T. Vine. Fuzzy logic in radar resource management. In *IEE Multifunction Radar and Sonar Sensor Management Techniques*, pages 5/1–5/4, 2001.

K.D. Ward, R.J.A Tough, and S. Watts. *Sea Clutter: Scattering, the K Distribution and Radar Performance*. IET Radar, Sonar and Navigation Series 20, 2006.

R.B. Washburn, M.K. Schneider, and J.J. Fox. Stochastic dynamic programming based approaches to sensor resource management. In *Proc. 5th International Conference on Information Fusion*, pages 608–615, Annapolis, USA, July 2002.

G. Weiss. *Multiagent Systems: A Modern Approach to Distributed Artifiial Intelligence*. Cambridge, Mass.;London: MIT Press, 1st edition, 1999.

A. Whitewood, G. Davidson, and A. Stove. Resource management for airborne surveillance radars. In *3rd IET Seminar on Intelligent Sensor Management*, pages 1–8, London, 2007.

J. Wintenby and V. Krishnamurthy. Hierarchical resource management in adaptive airborne surveillance radars. *IEEE Transactions on Aerospace and Electronic Systems*, 42(2):401–420, 2006.

E. Winter and P. Baptiste. On scheduling a multifunction radar. *Aerospace Science and Technology*, 11 (4):289 – 294, 2007. COGIS '06.

E. Winter and L. Lupinski. On scheduling the dwells of a multifunction radar. In *International Radar Conference*, China, Oct 2006.

W.D. Wirth. *Radar Techniques Using Array Antennas*. Institution of Engineering and Technology, April 2001.

P.M. Woodward. *Information and Probability Theory, with Applications to Radar*. Artech House, new edition edition, Dec. 1980.

M. Wooldridge. *An Introduction to Multi-agent Systems*. John Wiley & Sons, 2002.

J.H. Zwaga and H. Driessen. Tracking performance constrained MFR parameter control: applying constraints on prediction accuracy. In *8th International Conference on Information Fusion*, pages 546–551, PA, USA., July 2005.

J.H. Zwaga, Y. Boers, and H. Driessen. On tracking performance constrained mfr parameter control. In *Proc. 6th International Conference of Information Fusion*, pages 712–718, Cairns, Australia, July 2003.