Mission-Driven Sensor Management

Fok Bolderheij

Introduction

Managing the sensor systems onboard modern naval vessels demands an increasing amount of operator knowledge due to the fact that these vessels are equipped with state-of-the-art sensor systems that provide more functionality and more accurate information at the cost of more complex control mechanisms. Furthermore, the shift of operational areas to littoral waters with often dense civil traffic and rapidly changing geographical and meteorological conditions calls for a much more dynamic adaptation of these sensor controls in comparison with the more stable environment of traditional operational areas in the Atlantic Ocean. The lowering of defence budgets on the other hand creates a need for crew reduction and shorter education/training times, thus reducing the synergy created within teams of operators and the knowledge and experience of individual operators.

From the above it can be concluded that sensor management requires an increasing amount of operator knowledge, while in effect, the available amount of knowledge is decreasing. The consequences of incorrect sensor management may however be severe: if the sensor systems of a ship fail to detect threatening objects, the vessel may be incapacitated or even neutralised and consequently mission objectives will not be met. This observation justifies research into the ways in which operators deploy and optimise the available sensor systems to observe the environment and how these observations contribute to mission success. This will result in the support of, or ultimately, the automation of the deployment of complex sensor systems in a versatile maritime environment.

This paper describes research into generic sensor management principles that enable the development of a support system that is capable of bridging the growing gap between the available knowledge and the required sensor management related knowledge.

Sensor management issues

As already mentioned in the introduction, sensor management is currently executed by operators, who have to translate the goals of a mission into technical sensor settings while taking operational and political constraints like Emission Control (EMCON) plans and Rules Of Engagement (ROEs) into consideration. Because these technical controls are sensor specific, the operator must be familiar with the meaning of each setting and how changing this setting affects the performance of the sensor. Furthermore, the operator has to account for and, if possible, compensate for the prevailing environmental influences on the quality of the information (QoI) that is delivered by the sensor. Furthermore, the operators must be aware of the complementary properties of the different sensors and actually have to consider the deployment of the complete sensor suite as opposed to setting each individual sensor.

The observation that system-specific sensor management is a complex task that requires extensive operational and technical knowledge is also recognised in literature and various papers can be found that propose methods and algorithms to support this task. Strömberg et al. [Strömberg et al., 2002] have conducted a literature survey that presents an overview of relevant principles and methods concerning sensor management. Most of the methods reviewed by them provide a technical, sensor-oriented approach that strives for obtaining optimal sensor settings but leaves the translation of the operational requirements into technical sensor settings to the operator and therefore do not provide a solution to the identified problem. McIntyre and Hintz have compiled *a Comprehensive Approach to Sensor Management*, consisting of three papers [McIntyre and Hintz 1999-I, 1999-II and 1999-III], that describe a survey of modern sensor management systems, a new hierarchical model and goal lattices. In their first paper [McIntyre and Hintz 1999-I], they present the concept of the sensor management process and recognise sensor management as a process that contributes to the realisation of the mission goals; how this may be achieved is however not made clear.

Interviews with operational experts [Bolderheij and Absil, 2006] showed that a good picture of the operational environment, also referred to as the Operational Picture (OP), the Recognised Picture (RP), the Recognised Maritime Picture (RMP) or the Common Operational Picture (COP) can be regarded as a critical success factor for mission success. According to these experts, the OP consists of all objects in the neighbourhood of the platform and that two important conditions have to be met in order to consider it a good picture:

1. the OP must be complete;
2. the OP must be accurate.

Sensor management should therefore support the compilation of a good OP because the sensors are the resources that provide the required information about the environment. This means that the deployment of the sensor systems must be aimed at satisfying the identified conditions. Bolderheij et al. [Bolderheij et al., 2005] argue that these requirements can be met by constructing the OP from objects that represent the mission-relevant elements in the environment. They state that the OP can be considered *complete* if each relevant element in the environment is represented by at least one (preferably by only one) object in the OP and that the *accuracy* of the OP can be pursued by reducing the uncertainty in the information about the object. To maintain the completeness of the OP, sensor systems have to be deployed to *search* the environment for the presence of these elements while the accuracy can be increased by *tracking* and *classifying* them.

Integrating sensor management in the Command and Control process

The discrepancy between the available and the required amount of sensor management related knowledge described in the introduction gave rise to the question whether and, if so, how the sensor management process could be embedded in the Command and Control (C2) process because this process is currently executed by an operator who utilises mission related data to deploy the sensor systems.

The Allied Joint Doctrine [AJP – 01(B)] defines C2 as the process that plans, directs, coordinates, controls and supports an operation and therefore inherently has to direct, coordinate, control and support the deployment of the available sensor systems. In the previous section we saw that the complete and accurate picture of the environment is a mission critical success factor and since this picture is compiled from sensor observations, it is clear that the sensor systems are important resources that are essential for mission success. The sensor management process can therefore be regarded as a key sub-process of the C2 process.

Command and Control

The definition of the C2 process that was presented above in itself explains why the RNLN considers the C2 process of vital importance and why a substantial amount of research has been and is still funded to analyse the nature and layout of this process. With this research the RNLN wants to increase its efficiency and support the automation of the process, thus enabling further crew reduction. A review of related literature yielded a *cognitive* C2 process model [van Delft and Schuffel, 1995] which forms the basis for subsequent research into the C2 process founded by the RNLN. This C2 process model distinguishes four main sub-processes:

1. The provision of Situational Awareness (SA): this process gathers data about events in the environment of the platform and compiles a picture of the environment.
2. The execution of Threat Assessment (TA): this process enhances the information compiled in the picture of the environment by reasoning about the imposed threat.
3. The support of Decision Making (DM): this process makes decisions about the deployment of the available systems based upon the threat in the environment.
4. The execution of Direction and Control (DC): this process executes the decisions with respect to the deployment of the ship's systems or resources, thus striving for mission completion.

The layout of the process model is shown in Fig. 1.

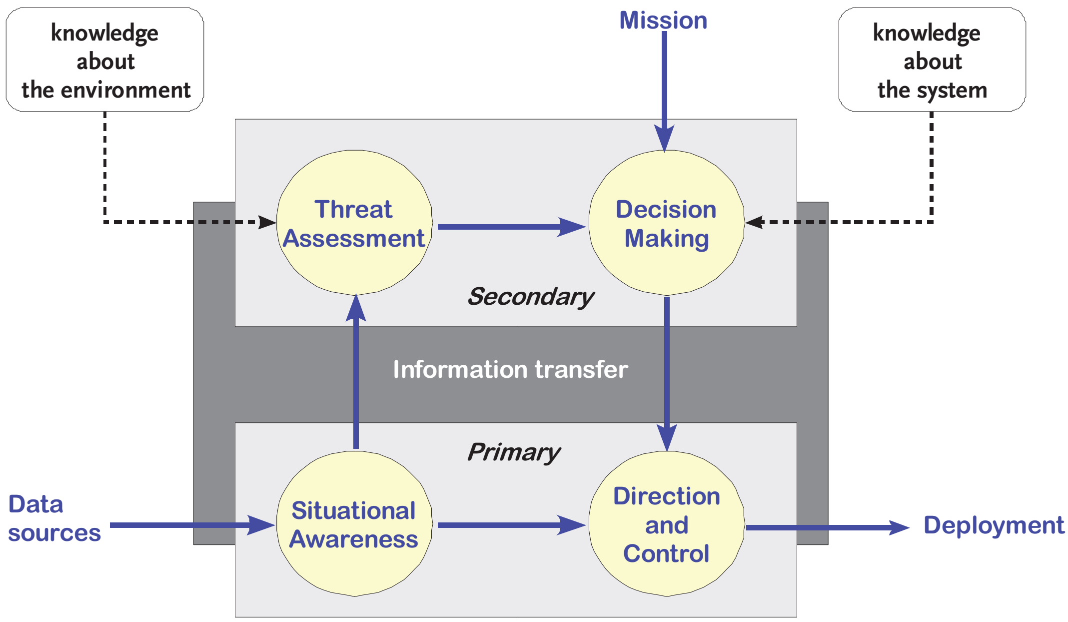

Figure 1. The cognitive C2 process model

Fig. 1 shows the three different types of input to the DM process:

1. Input from the TA process (information about threats in the environment).
2. Input from the mission (e.g. goals, requirements and constraints).
3. Knowledge about the system (e.g. available sensors and weapons).

From the description of the SA and the TA processes it can be concluded that these processes in effect execute the picture compilation process: the SA process gathers the information from the available sensors, associates, correlates and fuses this information and provides an *objective* view on the environment; the TA process now uses this information to infer the consequences of the elements in the environment in the (near) future. To accomplish this task, assumptions have to be made and therefore the picture becomes more *subjective* as a consequence of the reasoning process. At this point, the function of the direct connection between the SA process and the DC process became unclear. The report of Van Delft and Schuffel [Delft and Schuffel, 1995] states that predefined set points in the second level of information transfer enable the usage of this 'fast track'. Because the report takes a *Human Resources* point of view to the C2 process, this means that operators are required to control the combat systems by a nearly instantaneous appreciation of the situation based on intensive training. This *recognition-primed decision making process* as it is called by Klein and Grandall [Klein and Grandall, 1996] only functions if the situation at hand bears a resemblance to a training situation. This observation gave rise to the question whether this link can be maintained in the future, because due to the financial reasons mentioned in the introduction, training and education time will be limited. Because this research focuses primarily on the modelling and implementation of the sensor management related knowledge, research into the implementation of training would divert the attention too much from this objective and therefore this link was (temporarily) removed from the model.

The sensor control loop

After the removal of the link between the SA and the DC process, the C2 process showed a striking resemblance to the Observe, Orient, Decide and Act loop (OODA-loop) as proposed by Boyd [Boyd, 1987-1992]. This loop was initially intended to explain victory in air-to-air combat, but is nowadays also used within a wide variety of applications like in the design of business strategies. It describes how data is upgraded into information that in its turn leads to knowledge that can be used for actions that contribute to the realisation of mission goals. The consequences of these actions can now be observed as changes in the environment, observations that, after analysis may trigger more decisions and subsequent actions. From this description it is clear that the C2 processes from Fig. 1 can be directly mapped on the OODA processes. This provided an opportunity to redesign the C2 process as a loop, see Table 1.

Table 1. OODA and C2 processes

OODA Process	C2 Process
Observe	Provide Situational Awareness
Orient	Perform Threat Evaluation
Decide	Perform Decision Making
Act	Execute Direction and Control

To close this C2 loop, a feedback connection from the DC process to the SA process needs to be implemented. At first sight it may not be apparent what form this connection should take, but if one recalls the statement from the previous section that sensor management consists of at least two sub processes of which one sub process controls/tunes the sensor system, it can be seen then, that the loop can be closed by means of the sensor systems: the sensor settings that are produced within the DC process generate new sensor observations that are again inserted into the SA process. This results in a revised, OODA-loop based C2 concept. Fig. 2 depicts this new C2 process model in combination with the knowledge required for the sensor management process including some other resources that are controlled by the DC process.

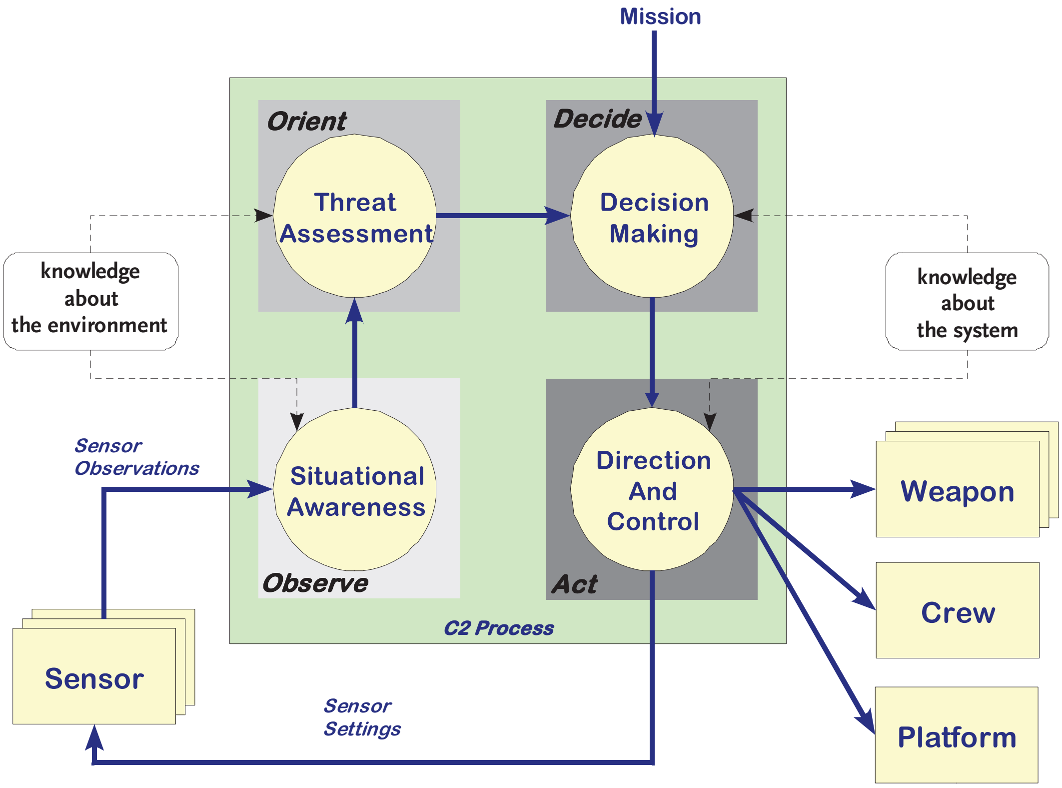

Figure 2. The revised C2 process model including the controlled resources

The adapted C2 process model now shows the outline of a sensor control cycle that induced further research.

The object-centred C2 model

The C2 process that was developed in the previous section is described in a *functional* way. Fig. 3 shows the processes (functions) that provide situational awareness and assess threats by processing and analysing the data received from the sensor systems to make decisions about the direction and control of the available resources. Data/ information flows from one process to another following the arrows in the diagram. In the development process this data flow has to be described because it is not evident what data is exchanged between the processes. Interviews and literature reviews made it clear however, that all data exchange has some relation to the OP which acts as a central element within the C2 processes: all processes either add information to the OP, update the information that is contained by its objects, or use this information to make decisions about deploying the available resources. These results enabled the first step towards a new, OP-centred, or *object-centred* C2 concept.

The introduction of the OP as the central element in the C2 process is visualised in Fig. 3. This rearrangement of the C2 processes provides the option to utilise the OP as a virtual *blackboard*: all processes write their object related information onto this blackboard or use the information that is written on it. The OP may also be regarded as a virtual *marketplace* where *agents* representing objects in the environment negotiate their need for resources.

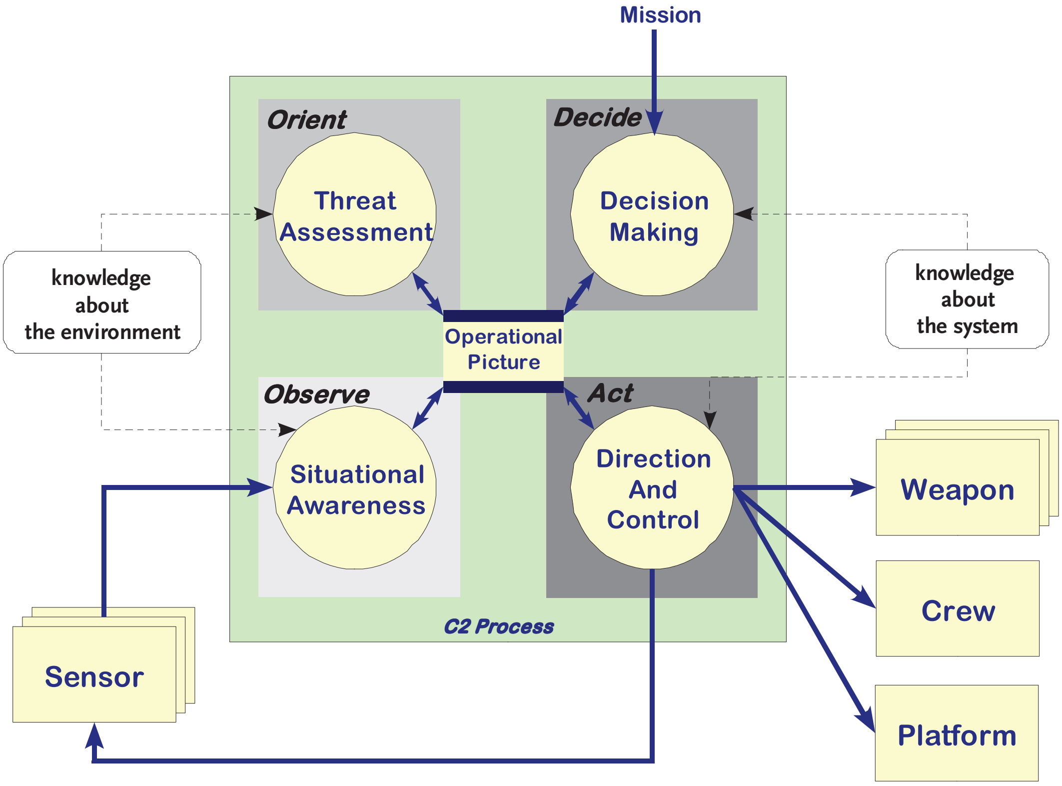

Figure 3. The object-centred C2 process

The reorganisation of the C2 processes that is introduced here also enables the implementation of fast, pre-programmed reactions triggered by specific observations and results in the reinstatement of the *recognition-primed decision-making* process [Klein and Grandall, 1996] that was briefly discussed previously.

A closer look at Fig. 3 shows that the initial OODA (SA-TA-DM-DC) loop based C2 concept of Fig. 2 is in fact abandoned, because all processes interact directly with the OP without a specific sequence. This is even more evident if the SA, TA, DM and DC processes are broken up in their constituting sub-processes: there is no apparent reason to maintain a specific execution sequence among those sub-processes. These sub-processes may even be executed concurrently if this is facilitated by the infrastructure. It can be seen however, that the concept of a sensor control cycle introduced previously still holds, as observations are provided by the sensor, are stored in the OP and are used to select a sensor and determine the sensor settings. This sensor control cycle is depicted in Fig. 4.

The upper-half of the cycle is formed by the processes that use the sensor information to compile the OP and the lower-half is composed of the processes that use the information stored in the OP to select the most appropriate sensor(s) and to control each sensor.

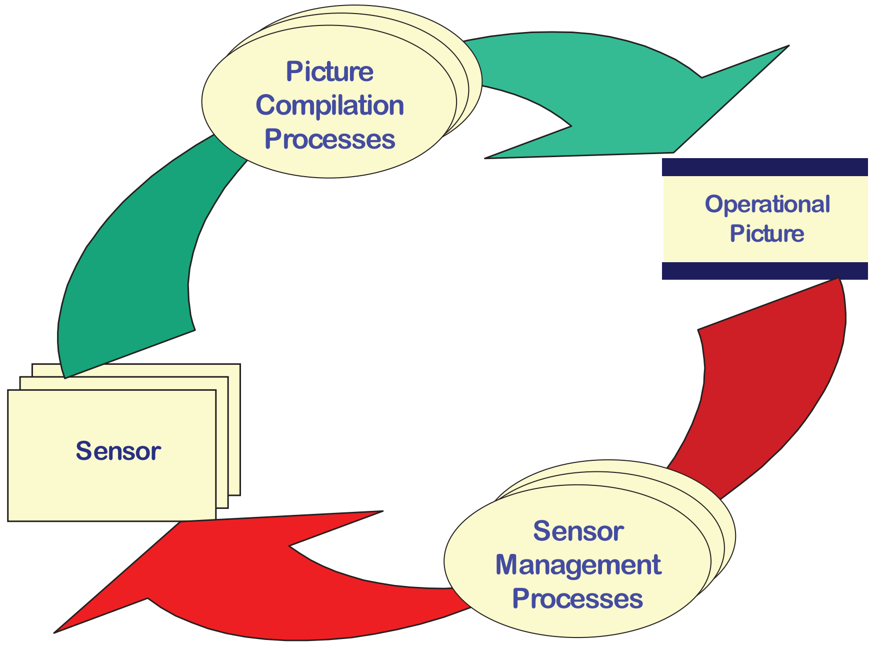

Figure 4. The sensor control cycle

The sensor management process

In their overview of sensor management methods and principles, Strömberg et al. [Strömberg et al., 2002] pose two relevant questions with respect to sensor management: "tasks want to know 'what sensor to select?' and sensors want to know 'what action to take?'" Based on these questions, two sensor management sub-processes can be identified:

1. A sensor tasking process that decides which sensor is the most appropriate for a specific task.
2. A sensor scheduling process that controls the sensor that is tasked by the sensor tasking process.

These sub-processes however are not identical to the sensor management functionalities proposed by Blackman and Popoli [Blackman and Popoli, 1999] who describe a template for a sensor management design consisting of two loops:

1. A loop controlled by a 'Macro Sensor Manager' assigning the tasks that need to be accomplished to satisfy the overall goals (here: the generation/maintenance of the Operational Picture);
2. A loop controlled by a 'Micro Sensor Manager' optimising the assigned tasks.

Both sets of functionalities can be reconciled by considering that the micro-level functionality, being the process that determines how a specific task is performed, is in fact a combination of the sensor tasking process and the sensor scheduling process. The combination of the sensor management tasking functionalities and the sensor management sub-processes yields a new sensor management concept that takes the shape of a three-stage sensor manager:

1. A task-composing stage that produces sensor tasks.
2. A sensor-selecting stage allocates (a) sensor(s) to these tasks.
3. A controlling stage that optimises the sensor settings with respect to the allocated task.

The first stage of this sensor manager analyses the information stored in the attributes of each object in the OP and uses information about the mission to determine if more *accurate* information is required and composes a sensor task based on these requirements. The Quality of Information (QoI) is determined by the accuracy of the sensor that provided the observation and the quality of the process that filters the observations (if present). Therefore this accuracy, or rather uncertainty, should be stored in combination with the information itself in the attributes of the object. This stage of the sensor manger was implemented as a rule-base. The second stage uses the available knowledge about the system from Fig. 1 to select the most appropriate sensor or combination of sensors. The third stage uses the information about the object, the knowledge of the system and the knowledge about the environment to determine the sensor settings. If too many tasks are assigned to this sensor, the task is returned to the second stage of the sensor manager and this stage now assigns the task to the next best sensor. These stages can therefore be combined in a single algorithm as described in [Van Norden et al., 2005].

From the description of the sensor management stages it can be seen that the first and second stage of the sensor manager can be situated within the DM process as they decide about tasks that need to be executed and the sensor that will be assigned to execute them. The third stage however has to be positioned within the DC process because the sensor is *controlled* to optimally execute these tasks.

Initiation of the object store

An analysis of the object-centred C2 process in combination with the three-stage sensor manager from the previous section reveals an initiation problem: at start-up time of the C2 process, no sensor observations are available and therefore no objects are present in the OP. Because the three-stage sensor manager uses the information stored in the object attributes to compose a sensor task, to select a sensor and to determine the corresponding sensor settings, no new measurements are generated. In order to enable the detection of objects in the environment of the platform, surveillance functions have to be initiated. These initial surveillance functions could be operator-controlled, but then the original problem that prompted this particular research reappears and therefore the sensor control loop needs to be initiated in a different way.

A solution to this problem is provided by a further exploitation of the mission information showed in Fig. 1. Within the planning stage of a mission, among other activities, the resources of the opposing forces are reviewed. This review process yields a set of potential weapon carriers and related weapon systems likely to be deployed by the opposing forces. These objects can also be stored in the OP as *virtual objects* because they represent real world objects that might be present in the environment of the platform, but are not (yet) observed. The information stored in the attributes of the instantiated objects, can be utilised by the sensor manager to select the most appropriate surveillance sensor and to determine the optimal sensor settings for these surveillance tasks. There may also exist more or less *uncertainty* about some aspects of the opponent's resources: depending on *quality* of the available intelligence information, it may not be known for sure what type of weapon systems the carriers are fitted with and their location may not be precisely known. This uncertainty resembles the uncertainty resulting from the sensor accuracy and this uncertainty can also be exploited by the sensor manager to determine the most appropriate sensor for the task and the corresponding task parameters.

Budget allocation and prioritisation needs

The object-oriented approach that was described in the previous sections supports the scheduling of the available sensor systems in the case of sufficient resources. In traditional sensor suites, different types of tasks are assigned to different types of sensors: search tasks are assigned to (long range) surveillance radars; target acquisition is accomplished by radars that provide a higher range and cross range resolution than the standard surveillance radars and tracking and illumination is done by track radars. Because dedicated radars are available to perform different tasks, there is not much need for dividing the sensor budgets. The only experience available in task scheduling and budget allocation is related to the deployment of mechanical Single Target Trackers (STTs) for Weapon Assignment (WA) purposes: once the decision has been made to deploy a guided weapon system, a scheduler selects the missile in combination with a fire control radar. The characteristics of this type of scheduling mechanism fit the needs of a sensor manager for tracking and weapon direction; it does not however reserve sensor capacity for not yet detected but potentially more dangerous objects (the so-called virtual objects) and is therefore not suited for scheduling Multi Function Radars (MFRs, sensors that provide surveillance, tracking and sometimes weapon guidance and/or classification capability) or complete sensor suites. In modern active MFRs, allocated search budget is not available for tracking purposes, and the illumination of an object for weapon guiding purposes will seriously drain the available time/energy budget (TEB). This means that in a scenario with a lot of 'neutral' traffic, these tracking tasks would consume the entire TEB while surveillance tasks are omitted and any missile in the vicinity of the platform would remain undetected. Therefore priorities have to be assigned to these different tasks.

Various scheduling mechanisms dealing with this problem are proposed in literature. For instance Huizing and Bloemen [Huizing and Bloemen, 1996] suggest a scheduling algorithm for an MFR, based upon an operator assigned priority of the sensor function type. The question that remains to be answered here is on what basis these priorities have to be assigned and furthermore, two similar sensor function types may require different priorities. Komorniczak et al. [Komorniczak et al., 2002], describe a prioritising mechanism based upon the kinematical properties, its Identification Friend or Foe (IFF) identity and an operator assigned rank of a threat object once this object is detected; this mechanism could be used to assign the priorities required for the tracking functions [Huizing and Bloemen, 1996] but it needs to be expanded for assigning priorities to surveillance functions.

According to Huizing and Bloemen, the prioritising mechanism has to placed in an operational perspective; this requirement closely fits the demand for the maximisation of the probability of mission success mentioned earlier, because depending on the intentions of the operators of the objects in the environment of the platform, these objects may damage or even destroy our platform. In a novel approach to solve his problem, the probability of mission success is maximised by ranking these objects with respect to their capability to cause mission failure and assigning the available sensor budget in accordance with this ranking. This *threat ranking* process can be executed by estimating the risk composed of the lethality of the object and the probability of occurrence of the damage that can be inflicted by this object. The risk estimation process is described in detail in [Bolderheij and van Genderen, 2004].

The new C2 concept with embedded sensor management
From the descriptions in the previous sections, the C2 model shown in Fig. 3 and inherently the sensor control cycle from Fig. 4 was developed in more detail, by identifying the processes that contribute to the picture compilation process and combining them with the three-stage sensor manager. The results are shown in Fig. 5. This figure outlines how sensor observations are merged into the OP and how this information can be used to track and classify objects and to infer their threat. The information is then utilised in combination with the related uncertainty to update the deployment of the sensors in order to keep the OP complete and accurate.

Simulation and results

To demonstrate the validity of the newly developed C2 model and the sensor management principles and the sensor manager that was designed along the lines of these principles, a prototype was developed and subsequently tested in a simulation environment. Operational experts were asked to assist with the composition of a sufficient realistic maritime scenario. In this scenario the deployment of a MFR consisting of four active antenna arrays was simulated because it has been shown in practise that this type of sensor is hard to control as it incorporates different sensor functions that need to be deployed simultaneously. The results of the deployment are logged and analysed after the mission has been completed.

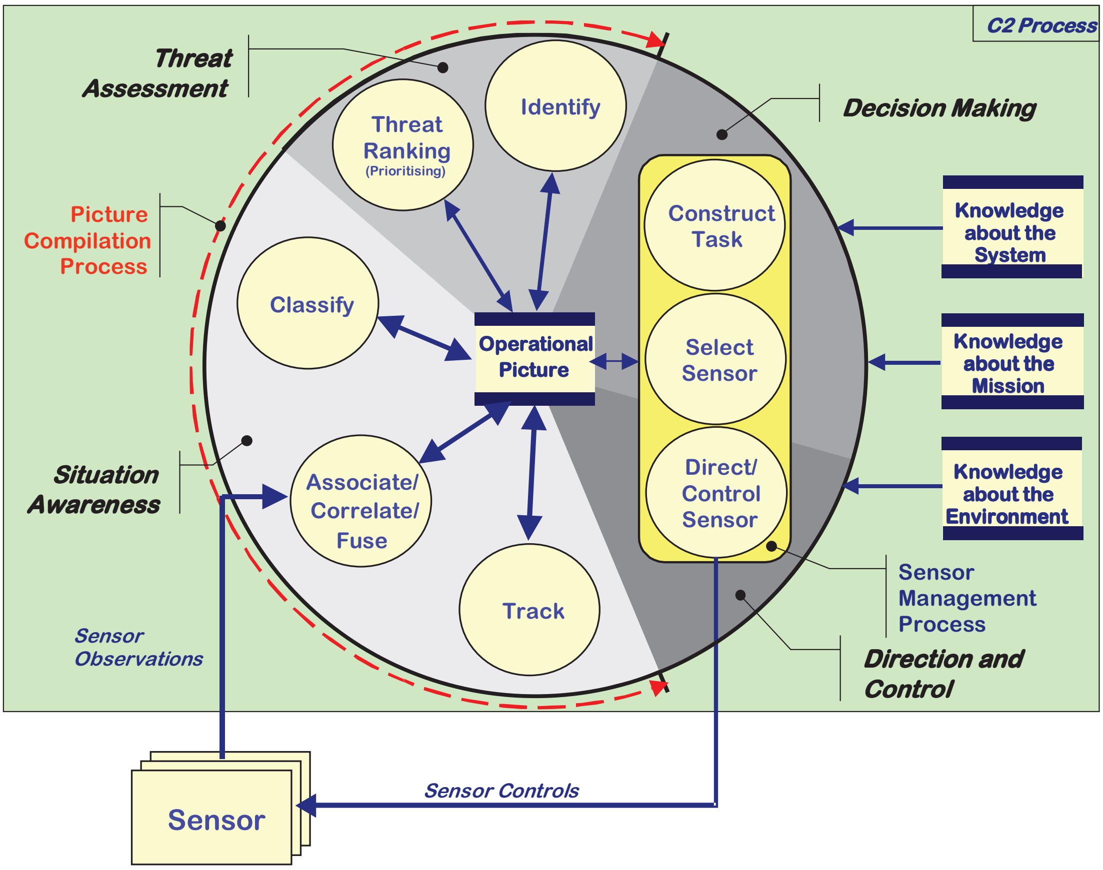

Figure 5. The more detailed C2 concept with embedded sensor manager

The scenario

In the constructed scenario that is depicted in Fig. 6, the important regional power Orange Country has just extended its territorial waters to include some important maritime oil fields, a claim that is heavily disputed by the international community.

To demonstrate the determination of the international community in this matter, an RNLN Air Defence and Command Frigate (ADCF) is tasked to sail along a navigation track (solid blue line) which is laid out just outside the original territorial waters (dashed green line) but well within the new territorial waters. The sensor suite of the ADCF consists of a Volume Search Radar (VSR), an MFR, several navigation radars, an Electronic Support Measures (ESM) system, an Infrared Search and Tracking system (IRST) and a Trainable Electro-Optical Observation System (TEOOS). Orange Country has deployed two land-based mobile missile launchers and a small aircraft carrier, which is positioned just within the border of the newly claimed territorial waters. Both launcher 1 (SSM site 1) and the aircraft carrier are fitted with subsonic Surface-to-Surface Missiles (SSMs) that are launched in the direction of a predicted hitting point and activate their internal radar after a pre-programmed time delay. Launcher 2 (SSM site 2) is loaded with an SSM type that first follows a set of predetermined waypoints before it activates its internal radar. Intelligence sources have made this information also available onboard the ADCF.

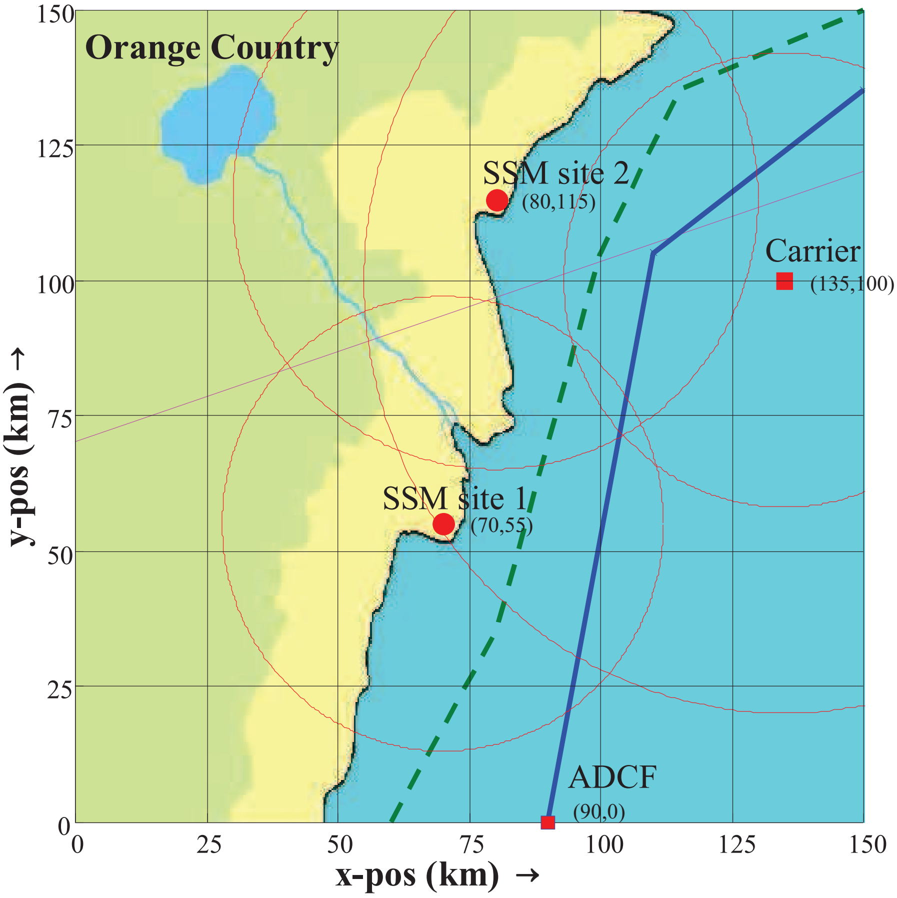

Figure 6. The maritime test scenario

At the starting time of the mission, the conflict between Orange Country and the international community has just escalated to a full-scale war. After the ADCF has entered the range of SSM site 1, it is discovered by a patrol aircraft originating from the carrier and subsequently four missiles are launched. Furthermore, two missiles are fired from SSM site 2 at ADCF when it comes within range. While the scenario is rolling, several civilian aircraft move through the area following the airway represented in Fig. 6 by the solid cyan-coloured line. Because intelligence sources are not sure whether all platforms with weapon carrying capabilities operated by Orange Country have been identified, a constant threat of Subsonic Sea-skimming missiles (SBSs) is assumed during the mission.

Results

The sensor control cycle was initiated by inserting the SSM launchers and the aircraft carrier in the system as virtual objects and the related prior information was stored in their attributes: this information initiated Limited Volume Search (LVS) functions when the ADCF entered the known (estimated) range of the particular missile range. Similarly, the generic SBS triggered a Horizon Search (HS) function because no specific direction could be assigned to this threat. Observations resulting from these surveillance functions caused the real objects within sensor range to be detected, tracked and eventually illuminated when this was operationally required. The allocation of the sensor budget during the execution of the mission is shown in Fig. 7. In this figure, the allocation of the time budget is shown as bar graphs in time steps of 10 seconds for each of the four antenna faces of the MFR. It can be seen from this figure that during the first stage ($t = 0$ s until $t \approx 4{,}500$ s) of the mission, only HSs are executed, resulting from the presumed presence of a generic SBS. The scheduler only allocates time to this task because no other tasks need to be performed. The amount of time is based on the object properties and (in future versions) on environmental factors. As the ADCF enters the range of the SSMs from site 1, an LVS is assigned in the direction of its presumed location (Face 1 and 4). The beam of this function is determined by the uncertainty in the position of this site. The height of the search pattern is derived from the maximum flight level of the missile type that may be launched from this type of launcher. While the ADCF continues its navigation track, several airliners are transiting following the air lane situated at an altitude of approximately 10 km. Because of the curvature of the earth these airliners may be detected coincidentally by an LVS. The reconnaissance aircraft takes off from the aircraft carrier at $t \approx 6{,}000$ s, and is finally detected by the HS function; the aircraft then turns and when it is outbound, the priority is reduced. The four missiles that are fired from launcher 1 at $t \approx 6{,}700$ s after the detection of the ADCF, are detected by the LVS, are subsequently tracked and a weapon guidance function is initiated. This is the only moment within the simulation that the sensor manager had to drop a sensor function due to insufficient budget. The HS is temporary dropped in favour of the illumination and the LVS because this function has the lowest priority.

The scenario had to be reconstructed several times to create this overload situation. If in future versions of the sensor manager more sensors are incorporated, the HS function has to be assigned to the next best sensor.

After the ADCF enters the range of SSMs from site 2 ($t \approx 14{,}500$ s), two missiles are fired, but since these missiles first follow a predetermined track leading them away from ADCF and because the launcher is still behind the radar horizon, these missiles are only picked-up at a very late stage by the HS function. At that time the risk posed by these missiles has already risen substantially and therefore a weapon guidance function is scheduled immediately at $t \approx 15{,}500$ s. In the remainder of the scenario, the MFR performs LVS functions to observe the aircraft carrier.

Because presently no comparable sensor management systems exist, operational experts were consulted about this prototype. When presented with the loggings of the MFR deployment, they remarked that this way of controlling a sensor would likely put too much strain on human operators, as it requires a constant update of the parameters that *drive* the deployment of the different sensor functions. Especially the repeated re-evaluation of the involved risk posed by the OP-objects and the resulting changes in priorities could create overload situations in these types of scenarios. This is less likely to happen when traditional sensors are used because then the OP is compiled from VSR observations that may be augmented by IFF data or information from other sensors.

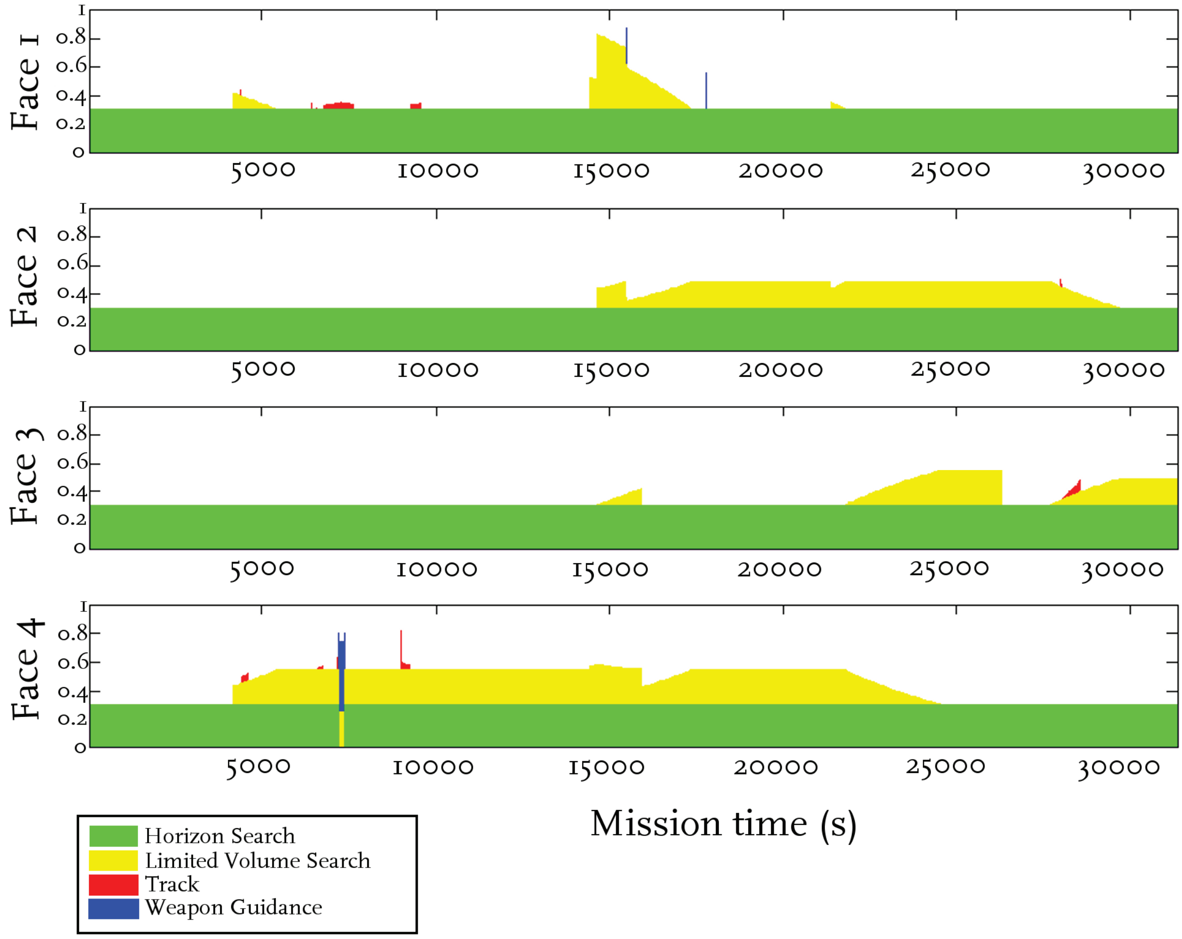

Figure 7. The allocation of the MFR budget for each of the four antenna faces

Threatening objects are automatically assigned to one of the two or three available STTs based on their classification and calculated Time on Top. In this case, the picture is compiled by means of several different sensors, where each sensor has its own controller or operator, provides only limited functionality and has its own TEB. Because an MFR is able to track and illuminate many more targets than two or three STTs while it is also capable of performing several different types of surveillance functions (semi) simultaneously, many more monitor and control actions are required in comparison with traditional sensors. As each function draws its TEB from the same source, the allocation of this budget among the different sensor functions must be controlled by a single operator/controller in order to avoid conflicts. This situation provides many more opportunities for the creation of overload situations causing sub-optimal sensor deployment. Because the MFR has only been recently introduced, no management guidelines exist and only limited MFR management experience is available and therefore it is difficult to compare the performance of the sensor manager with a human operator. Nevertheless, the general feeling of the operational experts was that the sensor manager was very well capable of outperforming a human operator, especially in terms of dynamic adjustment of the different control parameters and integral management of the complete sensor suite.

Other use

The object-oriented C2 concept was also tested in a prototype mission manager that controls the deployment of Unmanned Arial Vehicles (UAVs) and in an experimental combat management system that was developed to prevent firing at own forces (so-called 'blue on blue engagements') [De Jong et al. 2008]. In both systems, the applied concept resulted in an enhancement of the situational awareness and showed promising results with respect to the (re)deployment of the available resources (UAVs, soldiers) and the prevention of 'friendly fire'.

Conclusions

In this paper, sensor management was approached from an operational perspective and was described as the process that determines *what* sensor functions are required, *which* sensor(s) should be selected to execute these functions and *how* these sensors should be controlled to compile a *complete* and *accurate* OP.

An important new concept resulting from this research lies in representation of the OP as a set of objects that represent both the *observed* and the *expected* mission-relevant elements in the environment. The information about the properties of these elements in combination with the related uncertainty is stored in the attributes of these objects. This representation of the OP allows the picture compilation processes to be defined as the processes that seek to determine the properties of these objects and reason about the threat they pose. These processes now form half of a novel sensor control cycle. The other half of the cycle consists of a newly designed object-oriented, three-stage sensor manager. The *first stage* of this sensor manager analyses each of the objects in the OP to determine what tasks need to be executed to reduce the uncertainty related to the properties of the object in order to improve the *accuracy* of the OP. The *second stage* then selects the most appropriate sensor for this task based on the QoI delivered by each of the available sensors. Finally, the *third stage* controls the settings of the selected sensor or hands the task back to the second stage if the resources of the sensor are depleted. To initiate the sensor control cycle, *virtual objects* were introduced that enable the allocation of surveillance functions necessary to detect the expected objects, thus contributing to the required *completeness* of the OP. The objects that are detected by these surveillance functions are now tracked, classified, and identified to ensure the *accuracy* of the OP using sensor observations.

The prototype C2 system and sensor manager that were developed from this design were tested by managing an MFR model in a simulation environment. The execution of the simulation showed that autonomous deployment of complex sensors like the MFR by means of the information stored in the OP is feasible.

Because the sensor manager assumes the existence of an object-oriented OP, the integration with existing C2 systems either involves the reengineering of those C2 system components that execute the picture compilation process or the development of an interface between the existing C2 system and the sensor manager. These interfacing issues have to be resolved before the sensor manager can be successfully integrated into existing C2 systems.

References

Blackman, S. and Popoli, R. (1999) *Design and Analysis of Modern Tracking Systems*, Artech House, Norwood, MA, pp 1004-1018.

Bolderheij, F. and Van Genderen, P. (2004) Mission Driven Sensor Management.: *Proc. 7th Int. Conf. on Information Fusion*, Stockholm, pp 799-804

Bolderheij, F., Absil, F.G.J. and Van Genderen, P. (2005) Risk-Based Object-Oriented Sensor Management, *Proc. 8th Int. Conf. on Information Fusion*. Philadelphia.

Bolderheij, F. and Absil, F.G.J. (2006) Mission Oriented Sensor Management.: *Proc. Cognitive systems with Interactive Sensors,* Paris.

Boyd, J.R. (1987-1992) A Discourse on Winning and Losing. Unpublished briefing notes. Various editions. Available from: http://www.d-n-i.net/richards/boyds_ooda_loop.ppt

De Jong, J.L., Burghouts, G.J., Hiemstra, H., Te Marvelde, A., Van Norden, W.L., Schutte, K. and Spaans, M. (2008) Hold Your Fire!: Preventing Fratricide in the Dismounted Soldier Domain, *Proc. 13th Int. Command and Control Research and Tech. Symp.* Bellevue, WA.

Huizing, A.G. and Bloemen, A.A.F. (1996) An Efficient Scheduling Algorithm for a Multi Function Radar, *Proc. IEEE int. symp. on Phased Array Systems and Technology*, Boston, MA, pp 359-364.

Klein, G.A. and Crandall, B.W. (1996) *Recognition-Primed Decision Strategies*, ARI Research Note 96-36, U.S. Army Research Institute for the Behavioral and Social Sciences, Alexandria, VA.

Komorniczak, W. Kuezerski, T. and Pietrasinski, J.F. (2000) The Priority Assignment for Detected Targets in Multi-Function Radar, *Proc. 13th Int. Conf. On Microwaves, Radar and Wireless Communications*, Mikon, pp 244-247

McIntyre, G.A. and Hintz, K.J. (1999-I) A Comprehensive Approach to Sensor Management, Part I: A Survey of Modern Sensor Management Systems, *IEEE Transactions on SMC*.

McIntyre, G.A. and Hintz, K.J. (1999-II) A Comprehensive Approach to Sensor Management, Part II: A new hierarchical model, *IEEE Transactions on SMC*.

McIntyre, G.A. and Hintz, K.J. (1999-III) A Comprehensive Approach to Sensor Management, Part III: Goal Lattices, *IEEE Transactions on SMC*.

NATO. AJP - 01(B) (NATO/PfP unclassified): Allied Joint Doctrine.

Strömberg, D., Andersson, M. and Lantz, F. (2002) On Platform-Based Sensor Management, *Proc. 5th Int. Conf. on Information Fusion*, Annapolis, ML, pp 1374-1380.

Van Delft, J.H. and Schuffel, H. (1995) *Human factors onderzoek voor toekomstige commando centrales KM*, TNO-TM 1995 A-19 (in Dutch).

Van Norden, W.L., De Jong, J.L., Bolderheij, F. and Rothkrantz, L.J.M. (2005) Intelligent Task Scheduling in Sensor Networks, *Proc. 8th Int. Conf. on Information Fusion*, Philadelphia.