15th ICCRTS
"The Evolution of C2"

Situation Awareness for Supervisory Control: Two Fratricide Cases Revisited

Topic: C^2 Information Sharing and Collaboration Processes and Behaviours

Patrick Hew
Defence Science and Technology Organisation

Edward Lewis, Penelope Radunz, Sean Rendell
Australian Defence Force Academy

Point of Contact:

Patrick Hew
Defence Science and Technology Organisation
Canberra ACT 2600
Australia
Patrick.Hew@dsto.defence.gov.au

Abstract: Supervisory control is where a machine closes a control loop, and a supervisor intermittently programs the machine. It is informally known as "on the loop", versus the human being "in" the control loop. Supervisory control is the USAF's preferred concept for the C^2 of future unmanned systems.

We show how being "on the loop" requires situation awareness of a distinctly different nature to that for "in the loop". C^2 design has yet to address this opportunity. The 1988 downing of an Airbus by *USS Vincennes* prompted the Tactical Decision Making Under Stress (TADMUS) program, but we show how TADMUS was framed under "in the loop" thinking. Similarly, the Patriot Vigilance project followed from fratricides during Operation Iraqi Freedom. While oriented to "on the loop", the research has had to retrofit a training regimen to a largely as-delivered Patriot.

We revisit these fratricide cases for "on the loop" design principles. Key results: C^2 systems should make the "loop" explicit, and configurable by the commander. The system should support the commander to rehearse the loop's future behaviour, to seek disconfirming evidence, and to trap and debug errors. Double-hatting of humans as both "in" and "on" the loop ought to be avoided.

Introduction

Supervisory control is where one or more human operators are intermittently programming and continually receiving information from a computer that itself closes an autonomous control loop through artificial effectors to the controlled process or task environment [1]. Informally known as "On the loop", supervisory control has been proposed as a C^2 paradigm for combat [2] and unmanned systems [3]. The outstanding challenge is to translate the concept into requirements for systems engineering [4]. What are the activities associated with being "on the loop", as distinct from "in the loop"? What roles are in play? And what information is gathered or exchanged?

This paper develops C^2 design principles for supervisory control. We show how being "on the loop" requires situation awareness (SA) of a distinctly different nature to that for "in the loop". SA researchers have not yet recognised that being "on the loop" has different cognitive requirements to being "in the loop". As a consequence, C^2 designers have yet to realise the opportunities from supervisory control. Indeed, in the Vincennes Incident (1988) and the Patriot-Tornado Fratricide (2003), there were cognitive needs that were not supported by the C^2 systems of the day. The absence of support was a significant factor in what transpired, and serves as an indicator to future C^2 design.

It is timely and appropriate to revisit the two fratricide cases, given the progress in research on supervisory control. The Vincennes Incident prompted a major research program into Tactical Decision Making Under Stress (TADMUS), notably to find design principles for decision support, information display and training systems. TADMUS, however, was framed under "in the loop" thinking. Similarly, the Patriot-Tornado Fratricide was one of many cases leading to the Patriot Vigilance project. While oriented to "on the loop", the Patriot Vigilance project has had to design a training regimen around a largely as-delivered system. Looking forward, supervisory control is the preferred concept for the C^2 of future unmanned systems [5].[1] If this is to become doctrine, then the processes and behaviours of supervisory control ought to feature strongly in C^2 systems design. The alternative is for warfighters to discover the gap between design, training and operational use [6], potentially while in combat.

A Model of Supervisory Control

Our goal is to distinguish between cognition "on" versus "in" the loop. This is achieved by looking at where information is gathered, and where and how it is used. The enabling theory is that of (artificial) intelligent agents and situation awareness.

Intelligent Agents and Supervisory Control

In general usage, an *agent* is defined as someone or something that acts or has the power to act [7]. In Artificial Intelligence research, an *intelligent agent* is an autonomous entity that observes and acts upon an environment (it is an agent) and directs its activity towards achieving goals (it is rational) [8]. There are no restrictions on an agent's construction

[1] The USAF Unmanned Aircraft Systems Flight Plan 2009-2047 mentions "Man on the loop" synonymously with supervisory control on p14. Section 4.6 puts it thus: 'Increasingly humans will no longer be "in the loop" but rather "on the loop" – monitoring the execution of certain decisions.'

(mechanical, electronic, biological, software …), nor whether it is unitary or a networked assemblage of components, nor whether it is mobile or stationary. Hence, for example, an unmanned aircraft system is best regarded as a collection of agents, each assembled from components (human or artificial) housed by the airframe, on the ground, or elsewhere.

To emphasise, an intelligent agent is characterized by its closing a loop from sensors to effectors. We will focus on *lethal agents*, agents that close a firing loop from sensor to weapon. Lethal agents are certainly a stressing case of special relevance to warfare, while being general enough for studying supervisory control in the large.

The definition of an artificial intelligent agent used the word "autonomous", a term we have yet to define. We do so by integrating the notion of supervisory control as follows: *supervisory control* is where one or more operators are intermittently programming and receiving information from an artificial intelligent agent [1].[2] We can thus quantify autonomy as the time between the operator providing supervision to the agent, ranging from zero to infinite as autonomy increases from low to high (Figure 1). Informally, when autonomy is low, the supervisor is "in the loop", while "on the loop" is high autonomy. Here the "loop" refers to the agent's operations, so being "in the loop" corresponds to continuous supervision, while "on the loop" is more intermittent [5].

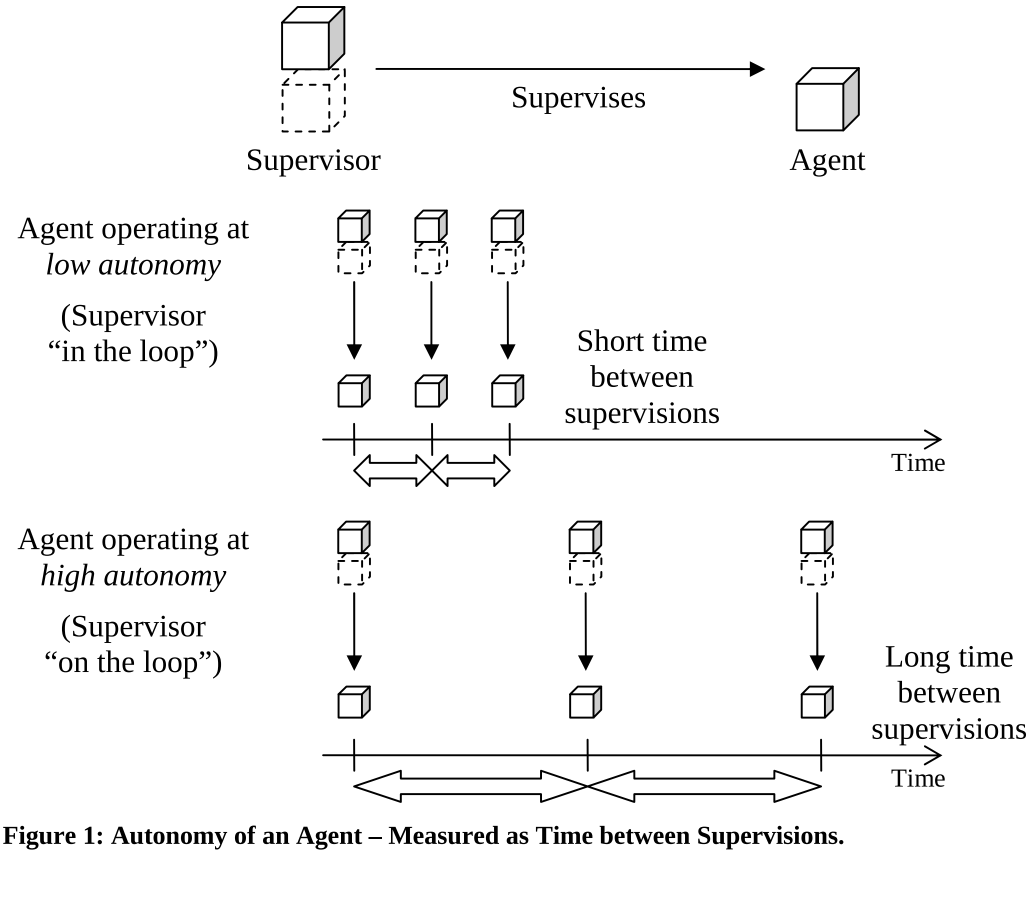

Figure 1: Autonomy of an Agent – Measured as Time between Supervisions.

[2] We take Sheridan's definition of supervisory control, delete the requirement for operators to be "human" and "continually" receiving information, and match in the definition of an artificial intelligent agent.

In defining supervisory control, we implicitly defined a *supervisor* (short for *supervising agent*). A *supervisor* is an agent that has supervisory control over a subordinate agent(s); it will intermittently reprogram its subordinates, using information that it has gathered from the environment or taken from the subordinate agents. As for all agents, a supervisor can be human or artificial, without restriction. Similarly, in line with the earlier definition of supervisory control, a supervisor's subordinates can be operating on a spectrum of autonomy, ranging from zero to infinite.

In essence, we have lethal agents that close a firing loop, from sensor to weapon. We then have supervising agents that close a supervision loop, from sensors to programming the lethal agent. To distinguish cognition "on" from "in" the loop, we need to be clear as to which loop we are concerned with.

The Loop and the "Situation" of Situation Awareness

The difference between "on" and "in" the loop cognition can be understood through the theory of situation awareness. C^2 research generally accepts Endsley's definition [9] of *situation awareness* (SA): The perception of the elements in the environment within a volume of space, the comprehension of their meaning and the projection of their status in the near future. Words such as "perception" and "comprehension" can carry philosophical overtones, vis-à-vis the notions of "I" and "me" [10, 11]. However, we propose that the closing of a loop from sensor to effector represents the generation and application of SA in some technical sense. Doing so allows us to use SA theory to distinguish "on" from "in" when looking at activities and information flows.

The key is to understand that a supervisor operates with a different "volume of space" to the lethal agent (Figure 2). For the lethal agent, the "volume of space" is the battlespace as conventionally understood. Objects are perceived in the battlespace, classified as targets or non-targets, and weapons are aimed and fired accordingly.

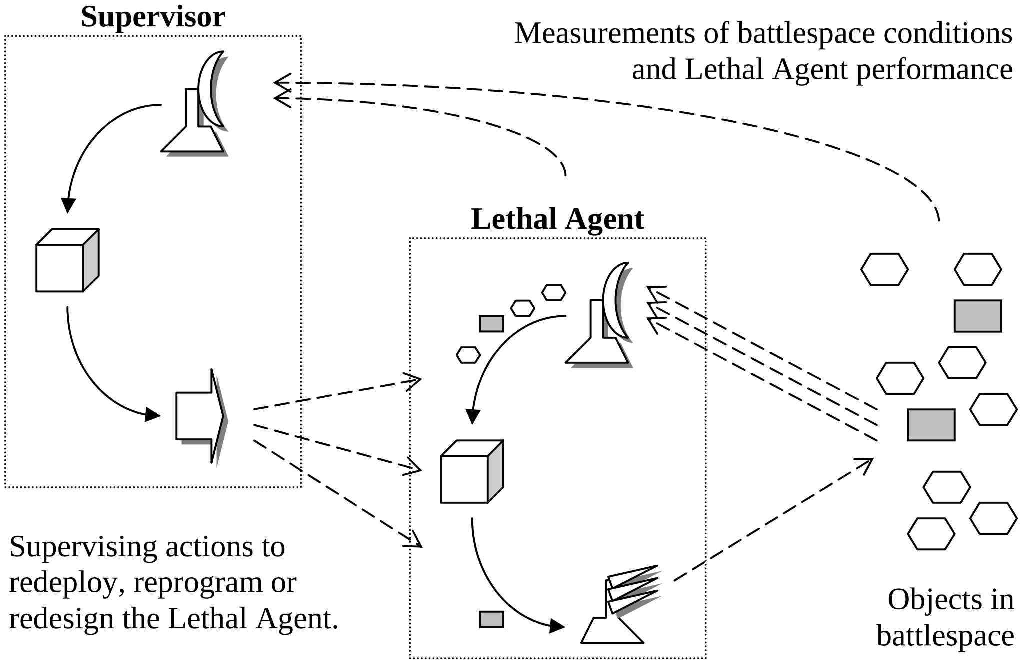

Figure 2: Supervisor and Lethal Agent.

In contrast, the supervisor's "volume of space" is the *state space* of the lethal agent. That is, the lethal agent cycles through some program (for example: Find-Fix-Track-Target-Engage-Assess [12]), and at any time it will be in some state. The supervisor's SA is concerned with the lethal agent's state, within the space of all possible states (the *state space*). Where is the lethal agent located within this state space? What state(s) is it moving to next? And how do those states relate to the external-world battlespace?

We can use very-simple lethal agents to see the contrast in "volume of space" between lethal agent and supervisor. A land mine exemplifies the very-simple lethal agent – it closes a firing loop between a pressure switch and explosive. In the technical sense used here, the land mine has SA of the volume of space immediately above the pressure switch. If the land mine is armed, then anything that triggers the switch is projected as being hostile and hence engaged.

The land mine can be in one of two states: armed or not. The supervisor is concerned with how the land mine's state relates to events in the battlespace: Should the land mine be armed or not? To program the mine (to be armed or not), the supervisor might use its sensors to also acquire SA of the battlespace to some quality. However, SA of the battlespace is only an input to SA and decision-making on the land mine's state.

We can view this alternately in terms of *mode awareness*. Mode awareness is about tracking an agent's *mode* (under supervisory control), so that it can be configured correctly [13]. A key requirement is for the supervisor to have a model of the agent being controlled. Here, we regard the agent's mode as being a point in its state space. Endsley's definition of SA can then be applied recursively: the modeling of an agent's mode is equivalent to tracking a point in a "volume of space" (the state space).

Cyberneticists will also recognize that the supervisor is modeling the lethal agent, and that we have reformulated the Conant-Ashby (Good Regulator) Theorem [14] in terms of SA. In particular, the state space of the lethal agent (from Endsley's "volume of space") equates to cybernetic *variety*, as per the Law of Requisite Variety [15]. Put alternately, the supervisor and lethal agent might be regarded as a single, combined system under the Viable System Model [16]. For the system to be viable, the supervisor needs to be able to audit the lethal agent, as well as (and separate from) monitoring the environment.

In general, across the lethal agent and supervisor, *three* distinct bodies of SA are formed:

1. The lethal agent maintains SA of the battlespace to close its firing loop.
2. The supervisor maintains SA of the lethal agent's state, within its state space.
3. The supervisor also maintains SA of the battlespace, so as to relate the lethal agent's state to the battlespace.

Expressed this way, it becomes apparent that the cognitive load for being a lethal agent ("in" the firing loop) is different and distinct for that of the supervisor ("on" the firing loop). In the case of a human being the "supervisor" to a land mine, the SA for the lethal agent's state is trivial (Is the mine armed or not?), and the cognitive load collapses to that of maintaining SA of the battlespace. However, we can expect difficulties if the lethal agent can take on a very large set of states, or if the state space is ill-defined or under-exposed. This is especially the case if a human is "double-hatted", into being part of the lethal agent, while also being in supervisory control of that agent.

Previous Research

Much of the extant literature is actually more concerned with humans "in" the loop, as part of (semi-) automated systems, rather than design for "on" the loop. The Tactical Decision Making Under Stress (TADMUS) program is a case in point [17], as it stands as the benchmark for subsequent research [18, 19]. TADMUS was formed as a result of the Vincennes Incident (1988), and thus focused on C^2 for maritime air defence. The centerpiece was the development of design principles for decision support, information display and training systems [17, 20], as exemplified in a prototype decision support system [21]. The work started with a cognitive task analysis, which selected the following five tasks as the primary decision-making requirements: recognizing that a track can pose a threat, identifying the track, determining the content of a suspicious track, avoiding unnecessary escalation and engaging the track if necessary [22].

We would now see these tasks as being those of the lethal agent ("in" the firing loop), rather than of the supervising agent ("on" the loop). This does not invalidate the specific design principles derived by TADMUS [23], nor the underlying models of human decision-making. However, if the goal is to support the human, then application of any model or paradigm requires clarity on which loop the human is "in". To be "on" the firing loop is to be "in" a supervision loop, with quite different cognitive requirements when compared with being "in" the firing loop.

In looking at Patriot fratricide incidents during Operation Iraqi Freedom, the US Defense Science Board (DSB) concluded that systems like Patriot ought to be designed for a more "man-in-the-loop" philosophy [2]. As the terminology for supervisory control has become established as “on the loop”, one can speculate as to whether the DSB genuinely intended “in the loop”, “on the loop” or some mixture. The subsequent Patriot Vigilance project conducted a root cause analysis, and concluded that the Patriot operator’s roles were defined as a by-product of automation [24]. That is, the operators were expected to “take care of” whatever the system could not handle. However, associated activities had not been made explicit, and hence the system was not designed to support the operators in discharging this “take care of” role.

The Patriot Vigilance project was thus invited to expose “what right looks like”, in both the design of battle command systems and the training of operators [25]. The analysis of design principles [24] raised many warnings about the introduction of automation. They especially advised again having the C^2 system coming to a decision, with the operators expected to “concur”. The proposed alternative was that the crews understand the technical capabilities of the lethal agent, and how they relate to the tactical situation [24]. The outstanding aspect is to identify the preferred crew behaviours, and foster them through the C^2 design.

The Patriot Vigilance research has particularly focused on training needs [26], to control the risk that “training failure might negate hardware promise” [27]. The need for training is not questioned, but C^2 design can and should ask, “Training in what?” An inspection of the Patriot air battle operations tasks [25] indicates that training for “in the loop” is still dominant within the US Army. In effect, the Patriot Vigilance project is seeking to rediscover and/or redevelop the roles of Patriot crew within the overall C^2 design, and thus retrofit a training regimen around existing hardware.

Researchers have long recognized that automation does not merely supplant human activity, but rather changes it [28]. However, there is an unmet need for frameworks that capture and model the new human activities, versus excluding them. An example of this gap can be seen in the popular Levels of Automation (Table 1) [28]. We note that Levels 1-5 have the human “in the loop”, while Levels 6-10 place the human “on the loop” to a computer that decides and acts. Furthermore, the levels are distinguished by the information provided from computer to human.

Table 1: Levels of Automation of Decision and Action Selection [28].

Level	Description
10 (High)	The computer decides everything, acts autonomously, ignoring the human.
9	informs the human only if it, the computer, decides to
8	informs the human only if asked, or
7	executes automatically, then necessarily informs the human, and
6	allows the human a restricted time to veto before an automatic execution, or
5	executes that suggestion if the human approves, or
4	suggests one alternative
3	narrows the selection down to a few, or
2	The computer offers a complete set of decision/action alternatives, or
1 (Low)	The computer offers no assistance: human must take all decisions and actions.

Unfortunately, the descriptions carry two hidden assumptions. The first is in only considering information flows at the point when the computer is (or has) made a decision and acted. The second is in implying that the computer is the sole source of information. What happens if we break these assumptions?

Hence, in the first half of this paper, we decoupled the supervisor from supervisee. Each agent operates in a loop of their own, potentially at different tempos, and they gather different information as per their particular responsibilities. The challenge, and opportunity, is to explicitly construct the roles of supervisors as part of the system's design, from which technology and training needs can follow. This situation invites a reappraisal of the incidents that led to the TADMUS and Patriot Vigilance research.

Fratricide Incidents Revisited

USS Vincennes shooting down Iran Air Flight 655 (1988)

On 3 July 1988, the guided-missile cruiser *USS Vincennes* mistakenly shot down Iran Air Flight 655, an Airbus flying over the Persian Gulf. For our purposes, the operative points are as follows [29-32] (Figure 3, Figure 4):

- *USS Vincennes* was ordered to the Persian Gulf in April 1988. The principal motivation was to enhance the surveillance coverage over the Gulf, especially over the Strait of Hormuz. The US had a growing presence in the Gulf, under operations initially intended to deter or prevent attacks on Kuwaiti shipping.
- *Vincennes* received some 7 months of training prior to deployment, and was was considered fully-capable for the mission. The training including a War-At-Sea Exercise in which threat aircraft had to be distinguished from other air contacts. There were also three exercises testing interpretation and correct response to Rules of Engagement (ROE), as current for US forces in the Persian Gulf / Middle East.

Photos courtesy U.S. Navy via Wikipedia.

Figure 3: *USS Vincennes*, 24 October 1988.

Photo courtesy U.S. Navy via Wikipedia.

Figure 4: Combat Information Center, *USS Vincennes*, 1 January 1988.

- At 1017 local time, Iran Air Flight 655 departed the Bandar Abbas joint military/ civilian airport. The aircraft was acquired by the *Vincennes*' SPY-1 radar, this being the principal sensor used by *Vincennes*' Aegis air warfare combat system. The aircraft was initially reported as Track Number (TN) 4474, later renumbered to TN 4131.
- Some 14 months earlier, the guided-missile frigate *USS Stark* had been attacked by Iraqi anti-ship missiles, prompting updates to ROE. All unknown air contacts were to be positively identified prior to designating them hostile and engaging, unless the unknown contact was displaying hostile intent or actually committing a hostile act. If radio warnings were unheeded, warships could use other means: fire-control radars, unmasking or training weapons, flares, signal/search lights or firing warning shots. However, if potentially hostile contacts persisted in closing, and if Commanding Officers believed that a threat was imminent, then they had the inherent right and responsibility to act in self-defence.
- As the Airbus departed the airport, it was interrogated by the Identification Supervisor using the UPX-29 Interrogation Friend or Foe (IFF) system. The UPX-29 saw a Mode III squawk, a code generally associated with a civilian aircraft for air traffic control purposes. At 1018, the Identification Supervisor consulted the commercial air schedule, and concluded that the contact of interest was not Iran Air Flight 655. However, Flight 655 was running 27 minutes late.
- US warships had received an intelligence briefing warning of the possibility of special operations by the Iranians to coincide with the American 4-July holiday. Iran Air Flight 655 was within the civilian air corridor, however it was 3 – 4 nautical miles off-centerline. Such a profile was abnormal.
- At 1019, the Anti-Air Warfare Coordinator ordered that TN 4131 be challenged over the Military Air Distress channel. From 1020 onwards, the first of several challenges was issued over the International Air Distress channel. *Vincennes*, however, had no way of knowing whether their radio calls had been received.
- At 1020, the UPX-29 reported an IFF Mode II squawk. At the time of *Vincennes*' deployment, IFF Mode II was regularly used by Iranian military aircraft. The Identification Supervisor saw the Mode II squawk, and reported a possible F-14 on the internal voice net. The Own Ship Display Assistant, having heard the identification, tagged the Airbus as an F-14 on the screens in front of the Commanding Officer (CO), Tactical Action Officer (TAO) and Anti-Air Warfare Tactical Action Officer. However, the UPX-29 was not actually interrogating the Airbus. The Identification Supervisor had configured ("hooked") his UPX-29 Remote Control Indicator to display the IFF results next to the TN 4131 symbol. In actuality, the IFF was interrogating the vicinity of Bandar Abbas, where there were any number of Iranian military aircraft.
- At 1022, the aircraft of interest had arrived at the critical point of 20 nautical miles from *Vincennes*. The Anti-Air Warfare Commander console operator illuminated the aircraft with a fire-control radar. A military aircraft would have known it was being illuminated, and maneuvered accordingly. However, civilian aircraft are generally not equipped to detect such emissions.

- Also circa 1022, the CO asked for a status report on TN 4474. He was unaware that the track number had been changed (to TN 4131). Moreover, subsequent analysis indicated that TN 4474 had now been reassigned to an A-6 Intruder in the Gulf of Oman, at that time descending and accelerating. This is the probable origin of consistent verbal reports aboard *Vincennes* of a descending contact.
- Meanwhile, the guided-missile frigate *USS Sides* was in company with *Vincennes*, and was also monitoring the action. TAO *Sides* noted that TN 4131 was continuing to climb, but the *Sides* air tracker was unable to gain the attention of TAO *Vincennes*. CO *Sides* had evaluated TN 4131 as a non-threat, but did not pass this evaluation to CO *Vincennes*.
- Circa 1023, CO *Vincennes* was searching for any kind of electronic emission that might help identify the “unknown-assumed hostile” contact that was steadily closing in range. He had acknowledged an earlier comment that the contact may have been a commercial airliner.
- At 1024, the CO *Vincennes* initiated the firing sequence. Two SM-2 Standard surface-to-air missiles were launched, destroying the Airbus.

Lest the depiction be over-simplified, we recognise that the above 7 minutes followed on the heels of a gun battle between *Vincennes*, *USS Elmer Montgomery* and Iranian small-armed boats. Also, an Iranian P-3 patrol aircraft was nearby, in a position and orientation consistent with providing targeting information to any aircraft preparing to attack the *Vincennes*. Finally, we acknowledge that we have a consolidated information picture, one with the benefit of hindsight. Within these bounds, we can consider the lethal agent and supervisor, and hence extract engineering requirements to support supervisory control:

1. **Data transport and consolidation is *not* supervisory control**. It is a long-standing principal of systems engineering to not employ humans where machines could suffice (for instance, the Fitts List [33]). However, in this case study, the lethal agent consisted of: the Aegis air warfare combat system, shipboard weapons (SM-2 missiles), shipboard sensors (especially the SPY-1 radar and UPX-29 Interrogation Friend or Foe system) and members of the *Vincennes*’ crew up to and including the Commanding Officer (CO). That is, the lethal agent was *not* a pure machine under the supervisory control of a human, but was constructed from both human and machine components. In the altitude reports over the voice net or the actions of the Own Ship Display Assistant, we see data and information being transported within the lethal agent via verbal and manual means. A modern system would seek to replace this verbal and manual transportation with machine-to-machine communication.
2. **A chain of command does not automatically equate to supervisory control**. The *Vincennes* chain of command was *not* a supervisory control chain. We see this in the handling of data from the UPX-29. Within the lethal agent, the measurements made by the UPX-29 were transported, relayed and summarized for the CO to action. In contrast, supervisory control would be to ask whether the UPX-29 was interrogating the contact of interest. This functionality was not apparent.
3. **Make the lethal agent’s program explicit**. Recall that the lethal agent can be modeled as following a program, and that the supervisor seeks to relate the program’s

state to the battlespace. This is difficult if the program has not been made explicit. The *Vincennes*' program was implicit: A non-hostile aircraft was one that matched the schedule for commercial flights, kept to the centre of the civilian air corridor, squawked Mode III, responded to radio calls, kept clear of *Vincennes* and held or increased altitude. If an aircraft diverged from schedule or the air corridor, squawked Mode II, did not respond to radio calls, closed with *Vincennes* (less than 20 nautical miles) and/or descended, then it was deemed to have hostile intent. Only when we have made the program explicit can the supervisor ask: Is this program robust enough to distinguish hostiles from non-hostiles at the level of discrimination that I want?

4. **Help the supervisor to rehearse/preview the lethal agent's behaviour in future scenarios**. The lethal agent needs to be made comprehensible to the supervising agent, prior to the lethal agent seeing combat. Note that this may need to occur every time the supervisor intervenes to reshape the lethal agent. We see this in the *Vincennes'* ROE. If the ROE envisaged using fire-control radars to warn off aircraft, then ROE promulgation was the time to establish that civilian aircraft could detect such radars, and hence respond. Similarly, if the ROE envisaged using warning means that were (essentially) visual-range limited, then ROE promulgation was the time to find means that would operate beyond 20 nautical miles.

5. **Supervisory control builds upon the options for shaping the lethal agent**. If the supervising agent has little or no options for reprogramming or redesigning the lethal agent, then the window for effective control will close. In the case study, we see CO *Vincennes* searching for electronic emissions for identifying the "unknown-assumed hostile" contact, in what turned out to be 2 minutes before launch. At this point, the options were severely limited; they were those emissions that *Vincennes* could detect, and that the Airbus could transmit. CO *Vincennes* did not have the option of obtaining an entirely new sensor within those 2 minutes.

6. **Supervisory control looks for disconfirming evidence**. The supervising agent needs to have ready and easy access to evidence that the lethal agent is incorrectly programmed. The case of conflicting information about an unknown aircraft is more cognitively stressful than other situations in air defence [34], and through our framework we can see why – the lethal agent has entered a part of its program that is ill-defined, forcing the supervisor to intervene at high tempo in a diminishing time window. The solution does not lie within the lethal agent, but from information feeds that the lethal agent has not taken into account. If CO *Sides* had come to a threat evaluation that differed from CO *Vincennes*', then the system ought to have enabled those evaluations to be shared and compared, to mutual benefit.

7. **Supervisory control looks to trap and debug errors**. The supervising agent should be alerted to instances when the lethal agent has reached an internally inconsistent state. In computer programming terms, when CO *Vincennes'* requested information on TN 4474, this track number had become a "dangling pointer" (an attempt to retrieve data, where the data has been moved to a new location) [35]. Dangling pointers are hard to trap and debug unaided – the program doesn't crash immediately, but continues to execute, with unpredictable results. The system needs to aid the supervisor to catch errors, locate their source, and thus resolve the problem.

8. **Avoid double-hatting of humans as both the supervisor and lethal agent**. The role of being a supervisor has a cognitive workload, separate and distinct from the demands on a human as (part of) a lethal agent. Hence, if a human being is double-hatted as both supervisor and lethal agent, the supervisor hat could well be dropped as the lethal agent ramps up. In the case study, the candidates for the supervising agent are the Anti-Air Warfare Tactical Action Officer and CO *Vincennes*. That there is ambiguity is from their also being components of the lethal agent.

The observations are consistent with findings that effective supervisory control is about training to "think outside the box", and to handle anomalies from systems being fallible [2, 26]. However, these activities are neither ad hoc nor anomalous in themselves; they are expressions of the supervision "loop", and stand in their own right.

US Army Patriot Battery shooting down a RAF Tornado (2003)

On 22 March 2003, a US Army Patriot Surface-to-Air-Missile (SAM) Battery destroyed a Royal Air Force Tornado over Kuwait. For our purposes, the operative points are as follows [36-40] (Figure 5, Figure 6):

- Patriot Battery C/5-52 was deployed to the north of Ali Al Salem Air Base, as part of the air and missile defence for Coalition forces fighting into Iraq. The principal mission was to defend ground troops from Iraqi tactical ballistic missiles. Extant Rules of Engagement (ROE) allowed the Battery to fire in self defense.
- At 2331 Greenwich Mean Time,[3] RAF Tornadoes with callsigns "Yahoo 75" and "Yahoo 76" were returning to Kuwait from operations over Iraq. They were acquired and tracked by the Kuwaiti Air Operations Center on radar. At 2345, Yahoo 75 was cleared to descend, and was directed to contact Approach Control at Ali Al Salem Air Base for approach and landing instructions. The aircraft followed the published speed and height procedures for a return to Ali Al Salem.
- At 2347, a contact was acquired by C/5-52's MQM-53 radar, and reported to the MSQ-104 Engagement Control Station. The contact was tracked as descending and closing with the Battery, so the MSQ-104 interrogated the track with the TPX-46 Interrogation Friend or Foe (IFF) system. From the trajectory and response, the MSQ-104 classified the contact as an anti-radiation missile (a missile designed to attack air-defence systems by homing in on their radar emissions). While the TPX-46 had been loaded with Mode 4 IFF codes (encrypted military), it had not been loaded with the Mode 1 codes (unencrypted).
- When the contact was reported to the MSQ-104, it was displayed to the Tactical Control Assistant (TCA). He, in turn, alerted the Tactical Control Officer (TCO). Normally, C/5-52 would have been tied into a Patriot Information and Coordination Central, or a Control and Reporting Centre, for situation awareness of the airspace beyond MQM-53 range. However, with much equipment still in transit from the US, C/5-52 only had hand-held voice communications with Patriot Battery B/2-1 and limited communications with higher headquarters.

[3] Greenwich Mean Time is 2 hours behind Kuwait Local Time. The article in *Jane's Missiles and Rockets* claims to use Local Time, but this does not reconcile with the Military Aircraft Accident Summary.

Photo courtesy U.S Department of Defense via Howstuffworks [40].

Figure 5: Patriot Battery MQM-53 Radar, incorporating TPX-46 Interrogation Friend or Foe.

Photo courtesy U.S Department of Defense via Howstuffworks [40].

Figure 6: Terminal at the MSQ-104 Engagement Control Station.

- The TCA had seen, on his MSQ-104 scope, symbology indicating that the contact was within 50 km of the battery and heading towards them. The TCO looked at her scope, and also saw symbology indicating a contact heading straight towards them. The TCO reinterrogated the track with the TPX-46. She also made a radio report to B/2-1. B/2-1 passed on the call for information to the Control and Reporting Centre.
- With no response to IFF, the TCO gave permission to the TCA to place the Patriot launchers in "operate" configuration, and authorized him to engage the track. A single Patriot PAC-2 was launched, destroying Yahoo 76.

Subsequent investigation looked closely at whether the Tornado's IFF beacon was operating and configured correctly. The ground engineering check on Yahoo 76's Mode 4 IFF was completed satisfactorily pre-engine start. However, at the time that the Tornadoes were acquired by the Kuwaiti Air Operations Center, there was only one IFF response, most-likely Yahoo 76. Indeed, at 2333, an E-3 airborne surveillance aircraft had identified the two radar contacts in the area as being "friendly", from having one set of IFF returns and a single, valid Mode 4 response.

We again have a consolidated picture, with the benefit of hindsight. However, we see many of the same engineering requirements as we did in the previous case study:

1. **Data transport and consolidation is *not* supervisory control**. In this case study, the lethal agent was again constructed from human and machine components: the MSQ-104 Engagement Control Station, the anti-air weapon (PAC-2 missile), sensors (notably the MQM-53 radar and TPX-46 Interrogation Friend system) and the crew of Battery C/5-52. Moreover, the TCA and TCO appear to be duplicating functions of the MSQ-104; the TCA in being a verbal repeater of a visual display, and the TCO in reinterrogating the track with the TPX-46.
2. **Make the lethal agent's program explicit**. What was the lethal agent's program? It was embedded in the threat logic within the MSQ-104. In particular, if a contact was closing with, and descending towards, the Battery at sufficient speed, and if it did not respond to IFF, then it was an anti-radiation missile (a high threat). In retrospect, we see that the lethal agent did not take into account the existence of the friendly air corridor into Ali Al Salem. While subsequent investigation concluded that the crew had a "rudimentary" understanding of the Patriot's target discrimination algorithms [39], there is finite budget and time for crew training. Systems design needs to lighten the cognitive load on the supervising agent, by making the lethal agent's internal logic readily comprehensible.
3. **Help the supervisor to rehearse/preview the lethal agent's behaviour in future scenarios**. That C/5-52 was operating with limited situation awareness and connectivity did not appear to be apparent to the airspace management planners. However, the threat-response logic (including Rules of Engagement) is known, and is promulgated as part of airspace planning. Hence, for instance, we could contemplate running a coarse-grained simulation of an airspace management plan, to rehearse it and hence detect problematic interactions between units (friendly and/or hostile).
4. **Supervisory control builds upon the options for shaping the lethal agent**. The MSQ-104 used generic criteria for anti-radiation missile, based on missiles deployed

worldwide. The UK investigation felt that it should have been configured to a tighter threat library [36]. Software-configuration threat libraries and algorithms are well within the capabilities of modern computer systems, and ought to be exploited.

5. **Avoid double-hatting of humans as both the supervisor and lethal agent**. The design expectation may have been that the TCA and TCO were actually double-hatted, as both components of the lethal agent and as the supervisor. However, it is likely that the crew placed their emphasis on being part of the lethal agent, as seen in an emphasis on mastering routines and a tendency to automation bias [41].

The Patriot-Tornado fratricide has been described as a case where the operators only had a short time to confirm or veto an automated system [42]. Certainly, the operators had about one minute to decide whether to engage or not [36]. However, the firing loop was closed by the operators, *not* by the Patriot. There was no "confirming" or "vetoing" as such; the operators were "in the loop".[4]

Moreover, the interpretation of a short time interval is an artifact of "in the loop" thinking. The supervision "loop" runs at its own tempo, with its own triggers. The window for supervisory control was in the long hours before the contact. C^2 systems designers could exploit this time, in fostering the "on the loop" situation awareness.

Implications for Future Research and Future Systems

Research into Situation Awareness

The impact on SA research can be considered in a number of different yet equivalent ways. Perhaps the simplest is to recognize that supervisory control puts a premium on research into mode awareness. To recall, mode awareness is about what the supervised agent is doing now and what it is going to do next, as distinct from awareness of happenings in the external world. As developed previously, the Endsley-based theory of SA can still be applied, if we broaden our notion of what constitutes the "situation".

Put alternately, when a researcher studies an agent's SA, they need to establish what kind of agent they are looking at. Does the agent perform actions into the external world (such as a lethal agent), or is it a supervisor? The key is to use the activities and information flows to characterize the agent in question, versus the outward trappings – a person might have been designated as a "supervisor", but their behaviour may indicate otherwise. To incorrectly establish the kind of agent is to fit the wrong model to data.

SA researchers thus need to be prepared to assert that a supervising agent ought to be present, somewhere, and to seek that agent out. If we conclude that someone is double-hatted to both "in" and "on" the loop, then so be it! In tracing and studying agents, it may be necessary to expand the boundary of the system in question. High-capacity communication systems can expand the spatial boundary beyond the walls of any single room. Similarly, in the temporal domain, an agent may be operating at a comparatively glacial tempo, but its actions may nonetheless be critical. The emphasis ought to be on the end-to-end "loops", and not the constituent components.

[4] Compare with the Levels of Automation at Table 1. In the case study, the Patriot system was operating at Levels 3-5. "Confirming" or "vetoing" is at Level 6.

Design and Engineering of C^2 Systems

For the C^2 designer, the key obligation is to make the supervisor's role explicit. A clear recognition of the supervisor's role is crucial if C^2 system is to account for their needs, across the full inputs to capability (hardware, software, training, doctrine, …). The supervisor is *not* a new encumbrance on C^2 design – the requirement was always there, but was neither recognized nor addressed. However, this doesn't make the workload go away; it merely means that it has to be handled by ad hoc and unbudgeted means, with the consequences that we have seen.

The design principles from the case studies could serve as guidelines for future C^2 systems. They stand on their own merits, but are not exhaustive, being manifestations of a single phenomenon: the distinctness of roles for being "on" versus "in" the loop. It is for the C^2 designer to find an appropriate allocation of roles to skilled personnel, as enabled via technological systems.

The opportunity is to reinvigorate the use of robotics and automation within C^2 systems. In both of the fratricide cases, the failures were not in the automated components, but in how they had been configured by their supervisors. Indeed, the personnel were performing functions "in the loop" functions that could be readily automated, freeing them up for their "on the loop" responsibilities. A system could be engineered for better "in the loop" SA, but this will do nothing to address the "on the loop" issues that were identified above.

Conclusion

The situation awareness needed for being "on the loop" is different and distinct from that needed to be "in the loop". The supervisor centres their situation awareness on the state of the "loop", now and into the future, and how those states relate to the battlespace. In the two fratricide cases, the humans were double-hatted as both "in" and "on" the loop; as the "loop" ramped up, it is perhaps unsurprising that the humans dropped their "on the loop" hat. The roles thus need to be made explicit as part of C^2 systems design. Further, C^2 systems ought to make the "loop" program explicit, and configurable by the commander. They should support the commander in rehearsing/previewing how the "loop" will behave in future scenarios, to seek disconfirming evidence, and to trap and debug errors.

Acknowledgement

Colin Allen, Susan Blood, Richard Brabin-Smith, Lydia Byrne, John Canning, Gerhard Dabringer, Tony Dekker, Mary Cummings, Christian Enemark, Stephan Fruehling, Veronica Gray, John Hawley, Alex Kalloniatis, Gina Kingston, Elizabeth Kohn, Stephanie Koorey, Patrick Lin, Jeffrey Malone, Jeremy Manton, Gary Millar, Jeffrey Nachem, John O'Neill, Jon Rigter, Jason Scholz, Noel Sharkey, Robert Sparrow, Alan Stephens, Mike Sweeney, Richard Taylor, Paul Whitbread. The views and opinions expressed in this article are those of the authors and do not necessarily represent those of the Australian Department of Defence or the Australian Defence Force Academy.

References

[1] T. B. Sheridan, *Telerobotics, Automation, and Human Supervisory Control* Cambridge: MIT Press, 1992.

[2] J. K. Hawley, “Patriot Fratricides: The Human Dimension Lessons of Operation Iraqi Freedom,” *Field Artillery*, pp. 18-19, January-February, 2006.

[3] M. L. Cummings, S.Bruni, S. Mercier *et al.*, “Automation Architecture for Single Operator, Multiple UAV Command and Control,” *The International C2 Journal,* vol. 1, no. 2, pp. 1-24, 2007.

[4] "DoD Architecture Framework Version 1.5," 23 April 2007.

[5] "*Unmanned Aircraft Systems Flight Plan 2009-2047*," United States Air Force, 2009.

[6] J. K. Hawley, A. L. Mares, and C. A. Giammanco, *The Human Side of Automation: Lessons for Air Defense Command and Control,* ARL-TR-3468, Army Research Laboratory, 2005.

[7] *Macquarie Dictionary*, 2009.

[8] S. J. Russell, and P. Norvig, *Artificial Intelligence: A Modern Approach*, 2nd Edition ed., Upper Saddle River, NJ: Prentice Hall, 2003.

[9] M. R. Endsley, “Toward a Theory of Situation Awareness in Dynamic Systems,” *Human Factors,* vol. 37, no. 1, pp. 32-64, 1995.

[10] D. Hofstadter, *I am a Strange Loop*, New York: Basic Books, 2007.

[11] R. Kurzweil, *The Age of Spiritual Machines: When Computers Exceed Human Intelligence,* New York: Penguin Books, 1999.

[12] J. A. Tirpak, “Find, Fix Track, Target, Engage, Assess,” *Air Force Magazine,* vol. 83, no. 7, pp. 24-29, July, 2000.

[13] N. B. Sarter, and D. D. Woods, “How in the World Did We Ever Get into That Mode? Mode Error and Awareness in Supervisory Control,” *Human Factors*, vol. 37, no. 1, pp. 5-19, 1995.

[14] R. C. Conant, and W. R. Ashby, “Every Good Regulator of a System must be a Model of that System,” *International Journal of Systems Science*, vol. 1, no. 2, pp. 89-97, 1970.

[15] W. R. Ashby, “Requisite variety and its implications for the control of complex systems,” *Cybernetica,* vol. 1, no. 2, pp. 83-99, 1958.

[16] S. Beer, *Brain of the Firm,* London: Allen Lane, The Penguin Press, 1972.

[17] S. C. Collyer, and G. S. Malecki, "Tactical Decison Making Under Stress: History and Overview," *Making decisions under stress: Implications for individual and team training,* J. A. Cannon-Bowers and E. Salas, eds., pp. 3-15, Washington DC: American Psychological Association, 1998.

[18] B. A. Chalmers, R. D. G. Webb, and R. Keeble, “Cognitive Work Analysis Modeling for Tactical Decision Support,” in 7th ICCRTS, Quebec City, 2002.

[19] M. J. Liebhaber, and B. Feher, “Air Threat Assessment: Research, Model, and Display Guidelines,” in CCRTS, Monterey, California, 2002.

[20] J. A. Cannon-Bowers, E. Salas, and J. D. Grossman, “Improving Tactical Decision Making Under Stress: Research Directions and Applied Implications,” in 27th International Applied Military Psychology Symposium: A Focus on Decision Making Research, Stockholm, Sweden, 1991, pp. 49-71.

[21] J. G. Morrison, R. T. Kelly, R. A. Moore *et al.*, “Tactical Decision Making Under Stress (TADMUS) Decision Support System,” in IRIS National Symposium on Sensor and Data Fusion, MIT Lincoln Laboratory, Lexington, 1997.

[22] G. L. Kaempf, G. Klein, M. L. Thordsen *et al.*, “Decision Making in Complex Naval Command-and-Control Environments,” *Human Factors*, vol. 38, no. 2, pp. 220-231, 1996.

[23] S. G. Hutchins, *Principles for Intelligent Decision Aiding*, Technical Report 1718, Naval Command, Control and Ocean Surveillance Center, 1996.

[24] J. K. Hawley, and A. L. Mares, *Developing Effective Human Supervisory Control for Air and Missile Defense Systems*, ARL-TR-3742, Army Research Laboratory, 2006.

[25] J. K. Hawley, A. L. Mares, and J. I. Fallin, *Reconfigurable Tactical Operations Simulator (RTOS) Operational Demonstration: Post Deployment Build 6 Follow Up*, ARL-MR-0703, Army Research Laboratory, 2008.

[26] J. K. Hawley, A. L. Mares, and C. A. Giammanco, *Training for Effective Human Supervisory Control of Air and Missile Defense Systems*, Army Research Laboratory, 2006.

[27] J. K. Hawley, “Patriot Vigilance Project - Training and Leader Development for the Future Force,” *Fires Bulletin*, pp. 36-39, January-February, 2009.

[28] R. Parasuraman, T. B. Sheridan, and C. D. Wickens, “A Model for Types and Levels of Human Interaction with Automation,” *IEEE Transactions on Systems, Man, And Cybernetics-Part A: Systems And Humans*, vol. 30, no. 3, pp. 286-297, 2000.

[29] "Formal Investigation into the Circumstances Surrounding the Downing of Iran Air Flight 655 on 3 July 1988," Department of Defense, 1988.

[30] K. A. Dotterway, “Systematic Analysis of Complex Dynamic Systems: The Case of the USS Vincennes,” Masters Thesis, Naval Postgraduate School, 1988.

[31] S. A. Kelley, “Better Lucky than Good: Operation Earnest Will as Gunboat Diplomacy,” Masters Thesis, Naval Postgraduate School, 2007.

[32] N. Polmar, *The Naval Institute guide to the ships and aircraft of the U.S. fleet*: Naval Institute Press, 2001.

[33] C. W. Geer, *Human Engineering Procedures Guide*, AFAMRL-TR-81-35, 1981.

[34] L. Adelman, M. Christian, J. Gualtieri *et al.*, “Examining the Effects of Communication Training and Team Composition on the Decision Making of

Patriot Air Defense Teams,” *IEEE Transactions on Systems, Man, And Cybernetics-Part A: Systems And Humans,* vol. 28, no. 6, pp. 729-741, November, 1998.

[35] Z. A. Alzamil, “Application of redundant computation in program debugging,” *Journal of Systems and Software*, vol. 81, no. 11, pp. 2024-2033, November, 2008.

[36] "Aircraft Accident to Royal Air Force Tornado GR Mk4A ZG710," Military Aircraft Accident Summary, Directorate of Air Staff, Ministry of Defence, 2004.

[37] "Patriot System Performance Report Summary," Defense Science Board Task Force, 2005.

[38] "MIM-104 Patriot," *Jane's Land-Based Air Defence*, 2009.

[39] “Patriot confused RAF Tornado with a false target, say US academics,” *Jane's Missiles and Rockets*, 2004.

[40] M. Brain. "How Patriot Missiles Work," HowStuffWorks.com (28 March 2003), http://science.howstuffworks.com/patriot-missile.htm (15 January 2010).

[41] J. K. Hawley, *Looking Back at 20 Years of MANPRINT on Patriot: Observations and Lessons,* ARL-SR-0158, Army Research Laboratory, 2007.

[42] M. L. Cummings, “Automation and Accountability in Decision Support System Interface Design,” *Journal of Technology Studies*, vol. 32, no. 1, pp. 23-31, 2006.