Cover Page for Paper 170
Submission for the 15th ICCRTS

Title:	**The Design and Implementation of a Tower Air Traffic Control System, for Rapid and Augmented Cognition**
Track:	7: C2 Approaches and Organization
Paper ID:	170

Authors:
COL (Ret) David Tan Cheow Beng (Future Systems Directorate)
Teh Shi-Hua (Defence Science and Technology Agency)
Huang Weixuan (Defence Science and Technology Agency)
Tew Tze Hong (ST Electronics (Info-Software Systems))
Tan Boon Tat (Defence Science and Technology Agency)
Oliver Tan Kok Soon (Defence Science and Technology Agency)
Lee Wei Hong (Defence Science and Technology Agency)
Seah Kai Ting (ST Electronics (Info-Software Systems))
Organizations:
a) Future Systems Directorate
303 Gombak Drive B1-51, Singapore 669645
b) Defence Science and Technology Agency
1 Depot Road, #03-01J, Singapore 109679
c) ST Electronics (Info-Software Systems)
24 Ang Mo Kio Street 65, 5th Floor Block D, Singapore 569061
Telephone: +65 6373 3909
E-mail: tshihua@dsta.gov.sg

Abstract

This paper seeks to describe the development of a prototypical air traffic control system for civilian tower controllers working in the Airport Traffic Control Tower (ATCT). This system, consisting of two main subsystems namely, Virtual Assistants and Persistent Sentinels (VAPS) and One-Glance, was designed with the intention of aiding civilian tower controllers in their visual work environment, through exploring the concepts of Augmented Cognition and Rapid Cognition in its implementation. **Augmented Cognition** is about circumventing human cognitive limitations by engineering work environments that aid people's abilities to process, store and retrieve information presented to them. The VAPS subsystem attempts to provide some degree of augmented cognition for tower controllers by designing software assistants for them. **Rapid Cognition** aims to encourage rapid information ingestion and assimilation through the use of visualization techniques. The One-Glance subsystem, consisting of a smart large-screen display and interactive panels, was designed to enable "intuitive" display of information for controllers. The paper posits that the use of the Augmented Cognition and Rapid Cognition design goals in system design can empower civil tower air traffic controllers working with heavy air traffic volume and in a stressful environment to make positive decisions in the midst of competing demands whilst avoiding unintended second order consequences.

INTRODUCTION

This paper seeks to describe the design approach and implementation of a prototypical air traffic control system for civilian controllers working in the Air Traffic Control Tower (ATCT). The prototype system, consisting of two main subsystems namely, Virtual Assistants and Persistent Sentinels (VAPS) and One-Glance, was designed to aid tower controllers in achieving improved reaction time and better quality of decision when working under heavy air traffic volume conditions. The prototype system is a bi-product of the Augmented Cognition and Rapid Cognition design goals, which directly address human cognitive limitations. A proof-of-concept was conducted to investigate how the prototype system enhanced the functional effectiveness of the controllers. The preliminary observations indicated that the prototype did improve the performance of controllers when operating under heavy air traffic volume conditions.

BACKGROUND

The Air Traffic Control Tower (ATCT) is one of the most visually prominent structures at every airport with regularly scheduled flights. It is a critical component of the air traffic control system, handling all takeoff, landing and ground traffic. The tower controllers rely heavily on visual observation from their high vantage point to control both ground and air traffic within the airport. The communications between the tower and other parties (such as pilots or departure controllers) take place via radio and landlines.

There are two main tower controllers within the ATCT – the ground controller and the local or air controller. The **ground controller** is responsible for movement on airport manoeuvring areas, which generally include taxiways, inactive runways, holding areas, and designated aprons and intersections. He handles aircraft taxiing from gates to takeoff on runways and from landing on runways back to the gates. In fact, any aircraft, vehicle or person must obtain clearance from the ground controller prior to entering and when moving about in these areas, so that he can ensure that aircraft and vehicular movement on the ground occur simultaneously in a safe and smooth fashion. The **local controller** is responsible for the active runway surfaces, as well as the immediate surrounding airspace (usually a radius of 2 to 5 nautical miles, depending on individual airports). He ensures safe distances between aircraft before giving the requisite takeoff and landing clearances. He has to coordinate with the departure controller to receive aircraft for landing and hand off aircraft after takeoff expeditiously and in a safe manner. The ground and local controllers are sometimes supported by a flight data controller, who is in charge of reviewing flight data and generating flight progress information.

A Complicated Work Domain

As described in the Cynefin Framework (Kurt & Snowden, 2003), a complicated domain refers to one where the relationship between cause and effect is knowable, though not necessarily fully known, and establishing the relationship requires analysis or some other form of investigation and/or the application of expert knowledge.

Airports are complicated systems. Depending on the amount of aircraft movement at the airport, the work load of the tower controllers can vary greatly. For instance, on a normal day, the Wittman Regional Airport might see a few hundred takeoffs and landings. However, during the week-long AirVenture aviation show, the Wittman tower controllers may have to direct more than 2000 flights daily (Paur, 2009). The world's busiest airport by traffic movements in 2008, Hartsfield-Jackson Atlanta International Airport in Atlanta, Georgia, recorded a total of 978,824 movements (Airports Council International, 2009), an average of about two planes taking off and landing every minute. Ranked the 19th busiest airport by passenger traffic in 2008 (Airports Council International, 2009), Singapore Changi Airport handles more than 4000 flights each week and over 35 million passengers each year (Changi Airport Group, 2009). Airports worldwide handled a grand total of 77 million aircraft movements in 2008 (Airports Council International, 2008). In the environment of such overwhelming load, the responsibility of coordinating the numerous moving parts within the complex environment of the airport manoeuvring areas and surrounding airspace falls to the tower controllers.

The key priorities of air traffic controllers are to minimize flight delays, maximize airport capacity by ensuring smooth flow of traffic, while maintaining utmost flight safety. Flight delays are very costly, as they result in lost traveller productivity, as well as higher fuel and maintenance expenses for airlines. Flight delays at a single airport can also cascade and can have numerous downstream effects on its immediate region and globally. According to a report from the U.S. Joint Economic Committee from the House and Senate, U.S. domestic flight delays cost the industry and passengers $40.7 billion in 2007 (Smith, 2008). However, the costly effects of flight delays are not limited to the carriers and passengers alone. Flight delays caused by air traffic congestion at New York City's three major airports alone were estimated to have cost the regional economy more than $2.6 billion in 2008, and problems originating in New York region's airspace led to nearly three-quarters of nationwide delays (Partnership for NYC, 2009). The delays in New York City cascade and cause a domino effect, affecting flight schedules at airports throughout the global system. 32% of flight delays are caused by late-arriving aircraft from elsewhere (Bureau of Transportation Statistics, 2009), creating a vicious cycle of delay. The interconnections and interdependencies that exist within the airport and in relation to other airports contribute towards the complicated work domain for the tower controllers.

In this highly interconnected system made of numerous moving parts, a single event that occurs in the airport, a small delay, may snowball and have significant effects downstream. 78% of flight delays in 2007 occurred before take-off (U.S. Congress Joint Economic Committee, 2008), meaning that tower controllers are quite frequently placed in situations that may result in flight delays. Events that can occur in the Tower Air Traffic Control realm include pilot deviations, runway excursions, runway incursions, runway confusion due to poor visibility or inclement weather, etc.. This is on top of a number of possible aircraft emergencies that may occur, such as engine failure, high speed rejected take-off, landing gear problems and more. These events, and their corresponding responses, may have unexpected secondary effects, causing the original problem to morph into another. For instance, loose items that fall onto the runway from a vehicle may be ingested by the engine of another aircraft taking off, causing the aircraft to be stranded on the runway. If the tower controller chooses to close the runway in response, this will cause delay, which adds to the operating cost of numerous other aircraft that were scheduled to use the runway. Such unexpected situations, occurring in a complicated environment, place tower controllers at the heart of a complex adaptive mess (Bennet & Bennet, 2008).

The Motivation

In the environment of such overwhelming load, the responsibility of coordinating the numerous moving parts within the complicated environment of the airport manoeuvring areas and surrounding airspace falls in the hands of the tower controllers. Given the backdrop of such a complicated work domain, as well as the possible sudden and intense occurrence of unexpected events, a significant amount of traumatic stress is placed on the tower controllers. Compounding the challenges, the mental stress of being responsible for the safety of so many aircraft and their passengers can be immensely draining. This is in addition to the normal cognitive workload, which involves monitoring and managing routine traffic, solving conflicts, updating their spatial mental models and coordination work.

Despite such a difficult work environment, tower controllers can only rely on paper flight strips to keep track of a long line of incoming and outgoing aircraft, manually gather necessary technical information from multiple information consoles to guide aircraft and vehicles correctly, make corresponding decisions in split seconds and communicate these decisions via radio or telephone to the appropriate parties. Total concentration is required to keep track of so many entities at the same time and to make certain that everyone receives the correct instructions. Rochlin describes this as "having the bubble" (Rochlin, 1997).

> *"Those who man the combat operations centers of U.S. Navy ships use the term "having the bubble" to indicate that they have been able to construct and maintain the cognitive map that allows them to integrate such diverse inputs as combat status, information flows from sensors and remote observation, and the real-time status and performance of the various weapons and systems into a single picture of the ship's overall situation and operational status."*

In his book, Rochlin describes how this terminology was met with an "immediate and positive acknowledgement" in air traffic control centres and other similar operations, because it expressed concisely what air traffic controllers often found challenging to explain to others about their work. Skilled operators possess a repertoire of expert responses and a cognitive map that cannot be easily understood by outsiders. Thus, this poses a challenge in developing technology that can assist the human operator and not stand in the way of his ability to perform his functional tasks.

THE APPROACH

Leveraging on the advancement of technology, the idea is to introduce more automation into the current system in order to reduce the workload of tower controllers and to improve timely decision making. However, research has shown that automation does not only complement or replace human activities but may also change these activities in ways unintended and unanticipated by the system designers, resulting in new demands on the human operator. Selecting appropriate functions to automate is very crucial, as over-automation may lead to adverse effects, particularly for expert operators performing critical tasks. The thoughtless introduction of automation may interfere with the way experts obtain information and prevent them from making the important assessments that they can ordinarily make. For instance, skilled forecasters managed to do a better job forecasting the weather alone than those who were made to use computer-based systems to generate forecasts (Klein, 2009). The experts forced to use the system tended to accept forecasts that were good enough as they did not have time to enter the adjustments. Such systems are dangerous in the long run as they can even interfere with the development of expertise in novices, who develop the bad habit of blindly accepting system outputs.

Parasuraman, Sheriden & Wickens' (2000) describes a model for the types and levels of automation that provides a framework for an objective approach to the design of automation in human-machine systems. Based on the concept of "human-centred automation" , the model can be viewed as an integration of previous research on the levels of automation in decision and action selection (Sheridan & Verplank, 1978), together with the salient features of the human information processing model (Wickens & Holland, 1999). The model describes how automation can be differentiated into four broad classes - information acquisition, information analysis, decision and action selection, and action implementation. Information acquisition is the first stage of information processing, involving the sensing and registration of incoming data, while information analysis involves the integration of information that augments human operator perception. These stages are followed by decision and action selection, which involves selection from various decision alternatives, and finally, action implementation, which is the actual execution of the eventual choice of action. Parasuraman et al. (2000) recommends that different levels of automation be applied within each class, subjected to contextual requirements of the target system. In the case of ATC automation, high levels of information acquisition and analysis automation can be pursued if the system is sufficiently reliable, whereas high levels of decision and action automation should only be confined to low-risk areas. We also note that the four-stage process should be treated as a continuous closed loop, meaning that action implementation may lead to the need for additional information acquisition. Decisions made during the decision and action selection phase may be immediately proceeded by more information acquisition. Thus, the ability to seek out information, according to the specific situation or context, should also be considered an automated action implementation.

Rapid Cognition and Augmented Cognition

Rapid Cognition and Augmented Cognition are two design goals, introduced to guide the system development process and focus attention on supporting the user's cognitive work.

Rapid Cognition is about "sizing up the situation" almost immediately, relying on the unconscious' ability to process information within seconds (Cassleman, 2007). Rapid Cognition is based on a repertoire of mental models, a library of knowledge and experience, built up through extensive practice and familiarity. It is based on the belief that "overt reasoning is preceded by a non-conscious biasing step that uses neural systems other than those that support declarative knowledge" (Bechara, Damasio, Tranel, & Damasio, 1997). It recognizes that humans have perceptual and attention span limitations, which may result in lowered information ingestion rates in the event of information overload.

The key motivation of Rapid Cognition as a design goal is to assist the user in achieving comprehensive awareness quickly**.** Good situation awareness is a critical requirement for smooth and safe airport operations. Singapore Changi Airport is a good example. Singapore Changi Airport is considered a major aviation hub in Asia. On a daily basis, the tower controllers have to contend with managing an airport operation with a 1300-hectare footprint, with ground and air traffic operating from four terminals using two parallel runways (Civil Aviation Authority of Singapore, 2010). On the world-map, Singapore is a springboard for regional travel. With 83 airlines flying to and from 190 cities in 60 countries worldwide (as of

1 Jan 2009), Changi Airport is an international gateway to the Asia Pacific (Civil Aviation Authority of Singapore, 2009), handling a total of 240,360 aircraft movements in 2009 (Kaur, 2010).

In such a busy environment, comprehensive awareness within the airport, including the aerodrome, is a critical precursor to effective air traffic control. Tower controllers have to constantly regulate and prioritize aircraft and ground movements, based on planned schedules and unplanned events, on a 24/7 basis. They have to process voluminous data in order to make sense of all aircraft and ground movements while avoiding information overload. Situation awareness is not about bringing more data and information to the tower controllers. Beyond information gathering, it is about the ability to analyze and synthesize information into timely and actionable knowledge, which can be rapidly assimilated by the tower controllers, especially when confronted with rapidly changing situations such as bad weather, delays, early arrivals, ground and/or in-air demands, incidents and other disturbances.

To apply Rapid Cognition towards enhancing the information acquisition and information analysis processes, we should aim to present information to the user in a way that is tailored to his specific situational context, so as to encourage quick and meaningful assimilation of data, thus reducing the time required for the information acquisition and analysis. We want to prevent action paralysis in the event of information overload, by presenting information to the user tailored to the specific situational context. This is done through prioritization of information display according to work order and sequence, as well as by taking into account competing demands that can result in operational tradeoffs. The One-Glance subsystem, consisting of a smart large-screen display and interactive panels, was designed to enable the intuitive display of information for controllers. This was accomplished through the development of intelligent information widgets that correspond to the information needs of the users, and the careful design of a visually-appealing and intuitive user interface. Through the idea of Rapid Cognition, the One-Glance subsystem attempts to mitigate human perceptual and attention span limitations.

Augmented Cognition, on the other hand, addresses cognitive bottlenecks, which are due to the natural limitations in the number of mental tasks that humans can execute at one time (Kruse & Schmorrow, 2004). These bottlenecks include limitations in attention, memory, learning, comprehension, visualization and decision making. Augmented Cognition research is often focused on identifying and addressing bottlenecks that affect attention capacity or working memory, which primarily affect near-term decision processes (Greitzer & Griffith, 2006). It should be noted that the definition of Augmented Cognition employed here does not entail measuring or sensing the cognitive state of humans and using that information to adapt systems to humans' needs, but rather augmenting users through offloading tasks to machine automation.

The aim of Augmented Cognition as a design goal is to offload simple tasks that impose a high cognitive workload, and to assist users in better decision-making, particularly amidst competing demands. In order for the tower controllers to make the best possible decisions given the circumstances, they must make them within the operating context and not necessarily adhere to the bounded rationality of predefined standard operating procedures. The key is to understand the prevailing context and environment that the tower controllers are in, and how decisions have to be made amidst competing demands, without compromising the airport's key priorities of safety and efficiency. However, in the midst of multiple demands (both time-critical and downstream) and information overload, the tower controller may find his span of control overwhelming, which may detract from his or her ability to take action decisively. It is in this setting that Augmented Cognition is meant to create cognitive bandwidth capacity by allowing the tower controllers to concentrate on performing higher order cognitive tasks, while leaving the lower cognitive tasks to machines.

Augmented Cognition as a design goal aims to assist the human in both the decision selection and action implementation processes. It involves assisting the human through automation of simple, low-risk tasks that require high cognitive bandwidth from the user, thus reducing the number of iterations of the four-stage information processing cycle that the tower controllers have to go through. In high-risk areas, Augmented Cognition involves providing action alternatives (low automation), leaving the human to make all final decisions and actions. The aim is to increase the bandwidth and handling capacity of tower controllers in their high-tempo, high-risk and high-workload work environment. The Virtual Assistants and Persistent Sentinels (VAPS) subsystem comprise of two classes of automation: virtual assistants (VAs) and persistent sentinels (PSes). VAs aid their users by performing simple low-level information acquisition and analysis tasks (such as information retrieval and simple calculations), before providing the users with decision selection help. PSes aid the human operators by automating low-level, low-risk actions that would normally require significant cognitive bandwidth from the users. VAPS uses the idea of Augmented Cognition to increase overall cognitive performance of humans through the process of contextual reasoning and adaptive augmentation.

The two concepts of Rapid and Augmented Cognition represent an attempt to create a competitive advantage in the cognitive domain for tower controllers who have to operate in an environment where decisions have to be made quickly amidst competing demands and limited resources. There is undeniably some overlap between the two concepts, as it is quite conceivable that the ability to gain comprehensive awareness more quickly will have positive effects on the quality of decision-making and vice versa. However, the authors have chosen to differentiate the two, so as to lend emphasis to the different cognitive functions that the design goals are meant to address. Using the human processor model (Card, Moran & Newell, 1983) for illustrative purposes, Rapid Cognition is about supporting the perceptual subsystem (short-term sensory store and perceptual store), while Augmented Cognition aims to support the cognitive subsystem (working memory, cognitive processor).

Cognitive Systems Engineering Approach

A Cognitive Systems Engineering (CSE) approach was adopted for the development of the prototype. The Decision-Centered Methodology (Crandall, Klein, & Hoffman, 2006) serves as a guide for the design and implementation process of the One-Glance and VAPS subsystems. Several Cognitive Task Analysis interviews were conducted on civilian tower controllers during the Knowledge Elicitation phase of the process. The CSE team used the Critical Incident Method and simulation-based interviews, which involved walking the tower controllers through critical incidents that they have encountered, in order to uncover key cognitive challenges they face in their daily work. A key focus of the Knowledge Elicitation phase was the elicitation of the various mental models that expert controllers held about their workspace. Effective information acquisition and information analysis aids, forming the basis of visualization design in the One-Glance Subsystem, were designed to support the formation and maintenance of these mental models. The knowledge elicitation phase also focused on the identification of the common actions and decisions that tower controllers are required to make, the level of difficulty and risk of these tasks, and the critical cues expert operators use to handle such situations. This led to various design ideas of the VAPS subsystem that were implemented during the Application Design phase. A more comprehensive discourse of the CSE process undertaken can be found in previous work (Yeoh, Tan, Low, & Teh, 2009).

THE ONE-GLANCE SUBSYSTEM

The One-Glance Subsystem (OGS) evolved from the idea of enhancing the ability of tower controllers to achieve comprehensive awareness quickly. It was intended to be capable of displaying key information needed by controllers on a large display, in a timely manner, to achieve their respective objectives. It was hoped that this could be achieved through the consolidation of relevant information on a single screen so that activities could be monitored and understood at a glance, thereby speeding up information ingestion, improving information assimilation and situation awareness for rapid decision making. Therefore, the OGS was designed to provide a visual environment that corresponds to underlying mental models held by the tower controllers, thus achieving intimacy between visualization and thought. The team nature of the tower work environment was also taken into account, and the design of the OGS attempted to address the information acquisition and analysis needs of the team as a whole.

System Architecture

The architecture of OGS is illustrated in Figure 1. There are two major components in this subsystem.

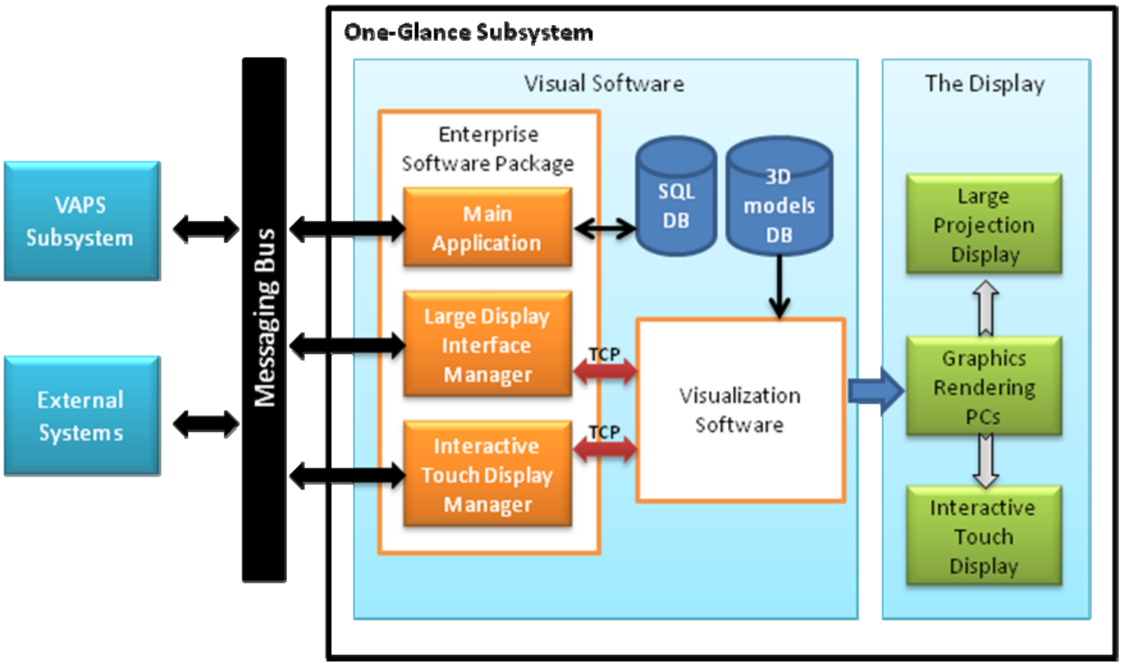

Figure 1: Architecture diagram of the One-Glance Subsystem

- The Display comprises of the large projection display, interactive touch display and the back-end graphics rendering PCs. These pieces of hardware are required to support the generation of 3D graphics and all other visual output.
- The Visual Software component consists of visualization software and an enterprise software package. It handles the processing and creation of the 3D models and also takes care of the interface and communication with other subsystems.

The Display

The Display component plays an important role in achieving Rapid Cognition. It is the key visual output that the whole system delivers to its end users. The integration of various commercial products was used to provide the visual display, as well as the human-computer interactivity.

To achieve a one-glance effect and create comprehensive team awareness, the development team adopted the use of a large display to aggregate both information and visualization. It has been shown that usage of such "Information Cockpits" as compared to standard desktop system aid in improving users' situational memory (Tan, Stefanucci, Proffitt, & Pausch, 2001). The OGS display was formed by tiled projection to produce a seamless, eye-catching wide display. This kind of setup is similar to those generally adopted by simulation systems. Compared to conventional multi-monitor display systems, which produce discernible lines between monitors and are restricted to single user operation, the use of a tiled projection solution allows the display of a wider, more seamless 3D virtual view. High-end graphics rendering engines were required to do geometry and 3D generation to support this. In addition, a transparent interactive touch display served as a secondary user console for non-time critical information and also functioned as an input device user interactivity with the system. This segregation of non-time critical information from crucial real-time information was a means to reduce information overload on the main display.

Visual Software

The Visual Software component consists of visualization software and enterprise software packages. The visualization software used was a suite of products from Vizrt, which is widely used in the broadcasting industry. It is capable of rendering 3D models and overlaying video imagery. This enabled the software developers to produce visually-appealing and intuitive graphical user interfaces. The visualization software is linked to the enterprise software package via two interface modules. The interface modules serve as the communication interface with external modules, such as the VAPS subsystem and other supporting subsystems. These interface modules, the Large Display Interface Manager and the Interactive Touch Display Interface Manager, parse incoming messages, filter them and provide logic processing for the visualization software. The OGS communicates with other subsystems via the TIBCO Enterprise Messaging Service.

System Features

The main focus of the OGS was to achieve the effect of Rapid Cognition by designing the user interfaces to correspond to the mental models of the tower controllers. This was further enhanced through the precise usage of graphics, shapes, information widget placement and colour to distinguish between the different levels of criticality of real-time events. Careful use of colour in the display is important as the effectiveness of colour use deteriorates with an increase in the number of colours on display (Xing, 2006). Excessive use of colours may reduce text readability and increase visual fatigue.

The following features were implemented to enhance information assimilation by presenting information to the user tailored to the specific situational context.

Embedding 2D widgets into 3D virtual world

The tower controllers work as a team to ensure safe aircraft departures and arrivals. Each Controller handles his own slice of work in ensuring aircraft safety under his charge. To maintain comprehensive awareness, the tower controllers need to know the locations of all the entities under their charge. As such, the tower controllers must construct and maintain “the bubble” by being alert and aware, constantly updating their spatial models and cognitive maps with incoming information.

The combined use of 2D and 3D displays has led to new possibilities for explorative analysis of spatial information (Bleisch & Nebiker, 2008). The OGS adopted this idea and employed the use of 3D models to create an augmented virtual 3D airport. Aircraft information in the form of 2D widgets was overlaid onto 3D aircraft entities, modelled based on actual aircraft types and textures. Also, a 2D airport layout overlay on top of the virtual 3D airport was introduced to allow tower controllers to quickly identify aircraft’s location by providing a bird’s eye view of the airport. This avoids sacrificing the tower controllers’ ability to build their mental 3D map and allows them to easily comprehend the traffic situation in the airport manoeuvring areas. Real-time information about the movement of various entities, such as distance, time and speed, was also displayed using 2D widgets.

Picture-in-Picture (PIP) view

Even though the tower controllers operate from a high vantage point, there are areas that are too far to be visible from the control tower. Cameras are typically used to view obstructed areas. However, this presents additional workload to the controllers, as they need to figure out where and what exactly is shown on the camera videos (imagine looking at the world through a straw). What the OGS provides is a PIP view of live video feeds, overlaid onto the virtual 3D airport, from cameras that are deployed in the airport. This enables the tower controllers to more easily identify the location of the video shown, thus enhancing their ability to assimilate information from the video content.

Critical Cues Display

Tower controllers spend a significant amount of time looking at and updating flight strips and panel mounted monitors showing information from a number of systems to achieve situation awareness. Such a layout scatters their attention and forces them to take their eyes off the ground situation and airspace, presenting a barrier to quick information assimilation. The OGS addresses the importance of tower controllers increasing “heads-up” time by providing an augmented large display where important information about departing and landing aircraft is consolidated. This information is separated into two levels of detail. Critical information, like callsign and speed, is displayed at the first level. Controllers are able to interact with the display through the touch panel to retrieve second level details, like flight route, time, aircraft type, altitude, etc.. Flight strip information is periodically refreshed on the screen to keep tower controllers updated of the status of each aircraft.

When guiding aircraft to depart, local controllers are required to line up aircraft prior to takeoff. With aircraft schedules represented using icons and tiled according to their takeoff sequence, the local controller is able to quickly identify the next aircraft that should depart. The schedule icon automatically changes in colour when the associated aircraft is behind schedule, serving as a visual indicator to the local controller to perhaps reshuffle the takeoff sequence, thus reducing the likelihood of delays. The aim is to present data on the augmented large display in a holistic and well thought-out manner, so to encourage effective information assimilation.

Enhanced Radar Display

The enhanced radar display (see Figure 2) was developed from a series of Cognitive Task Analysis interviews that were carried out to extract the tacit knowledge and critical cues used by the tower controllers.

The enhanced radar display consists of zoned areas that represent the distance in nautical miles from the threshold of the runway. Based on the information gathered from CTA interviews, the distance of an aircraft is an important piece of information during its approach. This data can be especially critical during emergency situations where the decision to hold or land can change in seconds, depending on ground traffic conditions. The zoned areas help the local controllers to see at one-glance the distance of aircraft, so that they can quickly decide what instructions to give to pilots. Reference lines that indicate the approach route also help local controllers to ensure that the pilots have lined up for the correct, allocated runways.

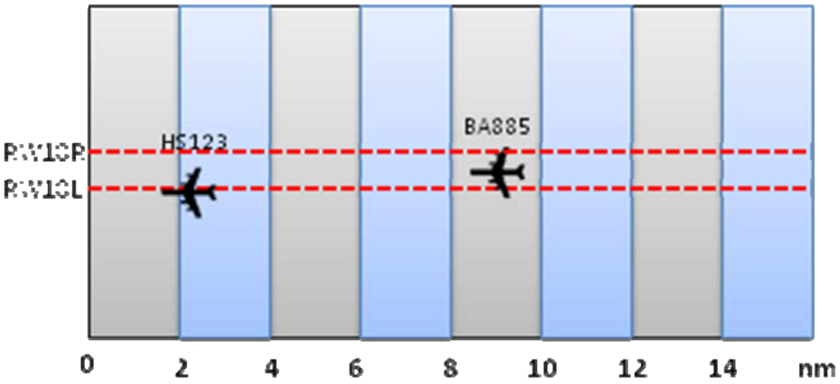

Figure 2: Enhanced Radar Display

THE VAPS SUBSYSTEM

To fulfil the aim of Augmented Cognition, the VAPS subsystem was built to provide software virtual assistants and persistent sentinels. The software assistants were designed to perform low-risk tasks and provide decision support, while leaving the user to make key decisions and carry out actions related to high-risk responsibilities. The VAPS subsystem was built as a multi-agent system, which corresponds closely to the concept of software assistants. The VAPS subsystem was not built with a separate user interface, but relied on the One-Glance subsystem for display and user interaction.

System Architecture

The architecture of the VAPS subsystem is as shown in Figure 3. The VAPS subsystem, similar to the OGS, communicates with other subsystems through the TIBCO messaging bus. All information to and fro the VAPS subsystem and other subsystems (such as flight plans databases, radar subsystems, etc.) pass through the Communication Interface. The Communication Interface sends and receives information to and from external subsystems (such as flight plans databases, radar subsystems, etc.) through the messaging bus and relays the data internally to and from the various VAPS components. Incoming data is first passed through the Situation Representation Module, which interprets and processes the data, before forwarding it to the appropriate agent tier module. This is a crucial step, as the data processing is tailored according to how tower controllers would themselves interpret the incoming data. It is key for augmented cognition systems that “there are mechanisms to provide an ongoing awareness of the operator’s context so that adaptations may be initiated in a contextually appropriate manner” (Forsythe, Kruse, & Schmorrow, 2008).

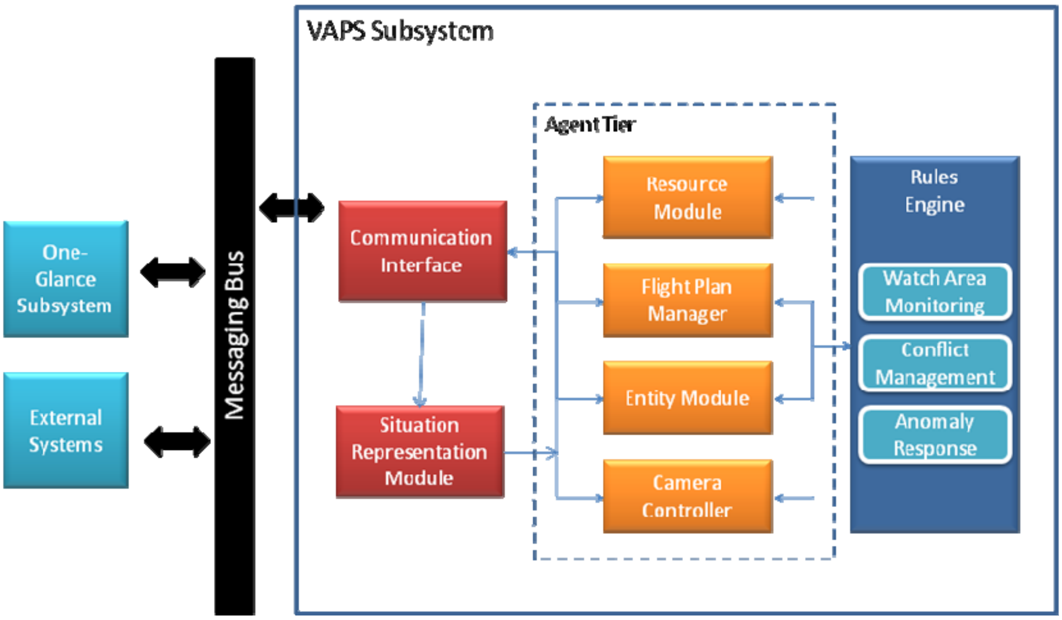

Figure 3: Architecture diagram of the VAPS Subsystem

The Agent Tier is designed to represent various aspects of the tower controller work domain. The Resource Module contains agent representations of the various resources in the tower control domain, including locations (such as runway, taxiway, intersections, airspace sector) and actual resources (such as gate availability, fire engine availability, etc.). The Flight Plan Manager keeps the updated flight plans of aircraft that are scheduled to use the airport within the day. The Entity Module contains software agent representations of real-life entities, such as aircraft and vehicles. The Camera Controller maintains the status of the Pan-Tilt-Zoom cameras that are available for tower controller usage. These four components receive processed information from the Situation Representation module and update their internal states accordingly. For instance, upon receiving notification that a new aircraft has been detected on the radar subsystem, the Entity Module will proceed to create a software agent that is aware of various pieces of information associated with that aircraft (such as its location and callsign).

The various agents in the Agent Tier will then submit information about their internal state to the Rules Engine, which executes configurable rules and outputs results back to the agents. Based on these results, the agents will carry out appropriate responses. To elaborate, continuing from the previous example, the created aircraft entity agent may submit its location and entity type to the Rules Engine. Assuming the Rules Engine is configured to trigger a camera response whenever an aircraft enters the runway, it will then trigger the Camera Controller to send out commands to external camera sensors to pan to the location where the aircraft is entering the runway.

The Communication Interface, Resource Module, Flight Plan Manager, Entity Module and Camera Controller components were implemented using an open-source agent framework called the Java Agent Development Framework (JADE). The various modules exchange relevant information (asynchronous message based communication) with each other through the internal message transport service provided by the JADE agent platform. The Rules Engine is implemented using a combination of Prolog engines and the IBM's WebSphere ILOG JRules. ILOG JRules provides a scalable execution engine and a means to store and manage rules from a centralized repository, while the Prolog engines enable the use of more advanced algorithms that are impossible to implement in ILOG JRules and can provide more advanced responses and suggestions to the user. Together, all the modules contribute to the ability of the VAPS subsystem to receive information from other sensor subsystems, populate and update the internal context of the various agents, and respond according to the rules configuration.

System Features

The three key system features that are described below are powered by sets of rules that are deployed to the Rules Engine module. However, the intelligent processing of incoming data by the Situation Representation Module serves as an important precursor, such that the rules may be written and configured in a more intuitive manner. For instance, the Situation Representation module may receive latitude/longitude information from a sensor subsystem, and convert that to a logical location, such as "Intersection E". This enables users to write rules referring to these logical locations, thus corresponding more closely to the tower controllers' mental models of their work space.

Watch Area Monitoring

The Watch Area Monitoring feature is a Persistent Sentinel function that allows users to define specific areas of interest where movement should be monitored. This feature achieves the function of Augmented Cognition by aiding the users in the low-level, high cognitive bandwidth task of constantly scanning the numerous areas under their watch. For instance, the ground controller may wish to be alerted of any vehicles that enter the active runway. Using the Watch Area Monitoring feature, he would then configure the Rules Engine to trigger a desired response (such as a text alert or video display) to help him monitor for any intruding vehicles. A more refined version of this rule could be to look up the clearance status of the vehicle (if such data is available) to determine if it is in fact a runway incursion. By automatically triggering system responses based on the conditions in specific areas of interest, the Watch Area Monitoring feature helps to reduce the cognitive workload of the tower controllers.

Conflict Management

The Conflict Management feature is a combination of a Persistent Sentinel function and a Virtual Assistant function. Since the state of the runways and taxiways are core to an aircraft's ability to take off or land, the Conflict Management feature was created to warn tower controllers of possible conflicts in the near term. For instance, if there are vehicles temporarily occupying any of the platforms, flight plans of aircraft scheduled to depart or land in the near-term that are affected by the obstruction will be highlighted on the One-Glance subsystem. This function involves participation from several modules of the Agent Tier, as the vehicle entity agent in the Entity Module must update the Resource Module of its occupation of a section of a runway, and the change in the resource agent status will trigger the Rules Engine to sieve out flight plans that are affected by the obstruction. This information is then passed to the Flight Plans Manager, which will generate the visualization display changes for the respective affected flight plans. This feature achieves Augmented Cognition for the tower controllers, by preventing them from having to sequentially acquire several pieces of information manually and alerting them of the possible disruptions to existing taxi plans. However, the change in taxi plans can have compound downstream effects, rendering it a high-risk area. Thus, this feature stops short of automatically revising the taxi plans in response to such conflicts.

Anomaly Response

The Anomaly Response feature provides recommended courses of action (COAs) in response to anomaly events. Anomalies can be triggered in two ways in the VAPS subsystem. Some anomalies can be automatically detected based on information available in the system. For instance, a runway incursion can be identified if a vehicle is found to be trespassing the runway while an aircraft that is scheduled to take off soon is already on it. Alternatively, users may tag an area using a predefined list of incidents. When an anomaly is triggered, the Anomaly Response feature will provide possible COAs that can help to resolve the anomaly. Many of the COAs are predefined and obtained from the standard procedures that tower controllers are supposed to follow. The timely display of the COAs is aimed at assisting the controllers in the decision and action selection phase. The agents add value by checking the availability and adequacy of resources required for the recommended actions. For example, if a fire accident occurs and the airport fire station's information system indicates no fire engines available, the original "Activate Airport Emergency Services" COA may be replaced by "Inform External Fire Services".

PROOF OF CONCEPT TRIAL

A proof-of-concept trial was carried out to assess the overall effectiveness of the prototype ATCT system. To determine the individual degrees of improvement in achieving comprehensive awareness and better decision-making would be difficult, as they are very much inter-related. Thus, the team chose instead to determine the extent to which the prototype system enhanced controller effectiveness in terms of reaction time, quality of decision and quality of awareness.

Six civilian tower controllers individually trialled the system in simulated scenarios that replicated the high tempo, high risk and high workload environment, with emphasis on two possible types of runway incursions: pilot deviation and vehicle deviation. Pilot deviation is when a pilot moves an aircraft into a position, without air traffic control approval, that leads to a loss of separation. Vehicle deviation is one where a vehicle enters a runway without air traffic control approval that leads to a collision hazard (Federal Aviation Administration, FAA, 2002).

A Subject Matter Expert (SME) from the ATCT, whose responsibility was to note interesting observations relating to decision-making processes during the trials, was invited to assume the role of an observer. Surveys and interviews, which not only served as a subjective means to garner user feedback on the prototype system, but also to obtain the level of self-perceived timeliness and awareness, were also administered to the controllers after each trial.

Observations and Feedback

One of the observations noted by the SME was that the prototype system seemed to aid the controllers in his decision and action selection phase, for instance, better decisions were made in terms of the re-schedule of the flight schedules to avoid delays due to disrupted flight plans. Further observations by the SME revealed that the Conflict Management and Anomaly

Response features in VAPS were useful to the controllers in achieving avoidance of flight delays and maximising airport capacity without compromising safety.

The six tower controllers generally felt that their level of awareness was improved with the ATCT prototype system, which had provided them a better sense of the situation, as well as increased bandwidth and handling capacity in their high risk, high workload environment. During the interviews, the controllers complimented that the One-Glance features helped them in detecting a loss of separation of aircraft more quickly, and assisted them in rectifying the problem in-time to avoid possible collision hazards.

CONCLUSION

In conclusion, given the complex and time-critical work environment of tower air traffic controllers, there is a need to be mindful of the degree of automation that is provided to expert operators and the specific processes that the introduction of technology is meant to improve. Based on preliminary observations from the proof-of-concept trial, which indicated that the implemented prototype did provide positive assistance to the tower traffic controllers, the choice of Rapid and Augmented Cognition as design goals have appeared to be useful in guiding the development of various features in the One-Glance and VAPS subsystems to enhance the work effectiveness of tower air traffic controllers. Future possible work could include more rigorous validation of the two subsystems to investigate their effect on other work effectiveness parameters for tower controllers, as well as the expansion of the range of test scenarios that the controllers are put through.

REFERENCES

Airports Council International. (2009, July 28). *Passenger Traffic 2008 FINAL*. Retrieved January 24, 2010, from Airports Council International: http://www.airports.org/cda/aci_common/display/main/aci_content07_c.jsp?zn=aci&cp=1-5-54-55_666_2__

Airports Council International. (2009, July 28). *Traffic Movements 2008 FINAL*. Retrieved January 24, 2010, from Airports Council International: http://www.airports.org/cda/aci_common/display/main/aci_content07_c.jsp?zn=aci&cp=1-5-54-57_666_2__

Airports Council International. (2008). *Worldwide Airport Traffic Report.* Geneva, Switzerland: ACI World.

Bechara, A., Damasio, H., Tranel, D., & Damasio, A. R. (1997). Deciding Advantageously Before Knowing the Advantageous Strategy . *Science* , 1293-1295.

Bennet, A., & Bennet, D. (2008). The Decision-Making Process for Complex Situations in a Complex Environment. In F. Burstein, & C. W. Holsapple, *Handbook on Decision Support Systems 1* (p. 854). Heidelberg: Springer.

Bleisch, S., & Nebiker, S. (2008). Connected 2D and 3D Visualizations for the Interactive Exploration of Spatial Information. *ISPRS Congress.* Beijing.

Bureau of Transportation Statistics. (2009, November). *Airline On-Time Statistics and Delay Causes*. Retrieved January 24, 2010, from Bureau of Transportation Statistics: http://www.transtats.bts.gov/OT_Delay/ot_delaycause1.asp?type=21&pn=1

Cassleman, E. A. (2007, June 19-21). Reclaiming Rapid Cognition. Paper presented at *the 12th International Command and Control Research and Technology Symposium*. Newport, RI.

Changi Airport Group. (2009, June 1). *Air Traffic Statistics*. Retrieved January 24, 2010, from The Official Site of Changi Airport Group: http://www.changiairportgroup.com/cag/html/the-group/air_traffic_statistics.html

Civil Aviation Authority of Singapore. (2009, January 28). *Singapore Changi Airport Registers Growth in Passenger Traffic in 2008*. Retrieved January 24, 2010, from Welcome to Changi Airport: http://www.changiairport.com/Journalist/Press_Releases/2009/news_0008.html?__locale=en&site=null

Civil Aviation Authority of Singapore. (2010, January 24). *Some Facts on Changi Airport*. Retrieved January 24, 2010, from Changi Airport Singapore: http://www.changiairport.com/changi/en/about_us/fact_sheets/facts_changi_ap.html?__locale=en

Crandall, B., Klein, G., & Hoffman, R. R. (2006). *Working Minds: A Practitioner's Guide to Cognitive Task Analysis.* MIT Press: Cambridge, MA.

Engineering Systems Division. (2001, October). *Working Paper Series.* Retrieved January 24, 2010, from MIT Engineering Systems Division: http://esd.mit.edu/WPS/esd-wp-2002-01.pdf

Endsley, M. R. (1990). *Predictive utility of an objective measure of situation awareness.* Proceedings of the Human Factors 34th Annual, 41-45.

Endsley, M.R., (1995). *Measurement of Situation Awareness in Dynamic Systems, Human Factors,* Vol. 37, pp. 65-84.

Federal Aviation Administraion, FAA (2002) *Chapter 4, Air Traffic Incidents,* retrieved from http://www.faa.gov/air_traffic/publications/atpubs/ATQ/CHAP%204.htm

Forsythe, C., Kruse, A., & Schmorrow, D. (2008). Augmented Cognition. In C. Forsythe, M. L. Bernard, & T. E. Goldsmith, *Cognitive Systems: Human Cognitive Models in Systems Design* (pp. 99-136). Mahwah, New Jersey: Lawrence Erlbaum Associates, Inc.

Greitzer, F. L., & Griffith, D. (2006). A Human-Information Interaction Perspective on Augmented Cognition. *Foundations of Augmented Cognition (2nd Edition)* (pp. 261-267). Arlington, VA: Strategic Analysis, Inc. .

Kaur, K. (2010, January 20). *Changi sets record traffic*. Retrieved January 24, 2010, from The Straits Times: http://www.straitstimes.com/BreakingNews/Singapore/Story/STIStory_479723.html

Klein, G. (2009). *Streetlights and Shadows: Searching for Keys to Adaptive Decision Making.* Cambridge, MA: MIT Press.

Kruse, A. A., & Schmorrow, D. D. (2004). Augmented Cognition. In W. S. Bainbridge, *Berkshire Encyclopedia of Human-Computer Interaction: When Science Fiction Becomes Fact* (pp. 54-59). Great Barrington, MA: Berkshire Publishing Group LLC.

Kurtz, C. & Snowden, D. (2003). The New Dynamics of Strategy: Sense-making in a Complex-Complicated World. *IBM Systems Journal, vol. 42, no. 3,* pp. 462-83.

Parasuraman, R., Sheridan, T. B., & Wickens, C. D. (2000). A Model For Types and Levels of Human Interaction with Automation. *IEEE Transactions on Systems, Man, and Cybernetics, Part A, 30(3),* 286-297.

Partnership for NYC. (2009). *GROUNDED The High Cost of Air Traffic Congestion.* New York: Partnership for NYC.

Paur, J. (2009, July 30). *Inside the World's Busiest Air Traffic Control Tower*. Retrieved January 24, 2010, from Autopia: http://www.wired.com/autopia/2009/07/airventure-control-tower/

Rochlin, G. I. (1997). *Trapped in the Net: The Unanticipated Consequences of Computerization.* New Jersey: Princeton University Press.

Salmon, P.M, Stanton, N.A, Walker, G., Green, D. (2006, March). *Situation Awareness Measurement: A review of applicability for C4i environments* Retrieved January 25, 2010, from Brunel University Research Archive: http://bura.brunel.ac.uk/bitstream/2438/1422/1/Situation_awareness_measurement_Salmon_et_al.pdf

Sheridan, T. B., and Verplank, W. L. ,"Human and Computer Control of Undersea Teleoperators," MIT Man-Machine Systems Laboratory, Cambridge, MA, Tech. Rep., 1978.

Smith, A. (2008, May 22). *Flight Delays cost $41B in 2007 - study*. Retrieved January 24, 2010, from CNNMoney.com: http://money.cnn.com/2008/05/22/news/economy/airline_delay/index.htm?postversion=2008052215

Tan, D. S., Stefanucci, J. K., Proffitt, D. R., & Pausch, R. (2001). The Infocockpit: Providing Location and Place to Aid Human Memory. *ACM International Conference Proceeding Series; Vol. 15* .

U.S. Congress Joint Economic Committee. (2008, May 22). *Schumer Releases Report Detailing Over $40 Billion in Costs of Flight Delays To Passengers, Airlines, and U.S. Economy As Memorial Day Travel Season Approaches*. Retrieved January 24, 2010, from Joint Economic Committee: http://jec.senate.gov/index.cfm?FuseAction=Press.PressReleases&ContentRecord_id=11115e27-d6ed-aa95-0871-fa86e86cd44a

Wickens, C. D., and Hollands, J. G. (1999), *Engineering Psychology and Human Performance (3rd Edition)*, Prentice Hall.

Wikipedia. (2007, December 3). *Linate Airport Disaster*. Retrieved January 24, 2010, from Wikipedia, The Free Encyclopedia: http://en.wikipedia.org/wiki/Linate_Airport_disaster

Xing, J. (2006, June). Color and Visual Factors for ATC Display. Retrieved from http://www.faa.gov/library/reports/medical/oamtechreports/2000s/media/200615.pdf

Yeoh, L. W., Tan, O., Low, K. S., & Teh, S.H. (2009, June). The Design and Implementation of a Cognitive Command and Control System for Tower Controllers. Paper presented at *the 9th bi-annual International Conference on Naturalistic Decision Making* . London, UK.

Yuditsky, T., Sollenberger, R. L., Rocco, P. S., Friedman-Berg, F., & Manning, C. A. (2002). *Application of Color to Reduce Complexity in Air Traffic Control.*