Software Agents -
An Application to the Airline MRO Industry

By

Arthur Wai Keong Poon

B.A., Aerospace & Aerothermal Engineering
Cambridge University (UK), 2003

Submitted to the Department of Aeronautics and Astronautics
In Partial Fulfillment of the Requirements for the Degree of
Master of Science in Aeronautics and Astronautics

At the

Massachusetts Institute of Technology

June 2004

Signature of Author . . . .
Department of Aeronautics and Astronautics
May 14, 2004

Certified by . . . .
Charles P. Coleman
Assistant Professor of Aeronautics and Astronautics
Thesis Supervisor

Accepted by . . . .
Edward M. Greitzer
H.N. Slater Professor of Aeronautics and Astronautics
Chair, Committee on Graduate Students

Software Agents -
An Application to the Airline MRO Industry

By

Arthur Wai Keong Poon

Submitted to the Department of Aeronautics and Astronautics
on May 14, 2004 in Partial Fulfillment of the
Requirements for the Degree of Master of Science in
Aeronautics and Astronautics

ABSTRACT

Poor management of aircraft spare parts, human-related maintenance errors and the lack of coordinated decision-making in maintenance scheduling are some of the problems plaguing the airline MRO industry. Airlines, independent MRO providers and government agencies, particularly the Federal Aviation Administration have identified inefficiencies in MRO operations as a barrier to the realization of a safer and more effective air transportation system.

The National Airspace System is slowly beginning to address this problem and one of its initiatives involves the gradual adoption of agent-based solutions and ontologies to automate some of the business processes in MRO depots. Agents and automated processes are ideal choices for making sense of ontologies and are key components in the next generation Semantic Web. The focus lies with improving business logic and practices of air carriers' MRO activities using software agent principles.

This thesis proposes agent architectures for aircraft inspections, repairs and supply chain management of aircraft replacement parts and discusses the resulting implications of using agents as a decision-making tool for aircraft maintenance and repair in commercial airlines.

Thesis Supervisor: Charles P. Coleman
Title: Assistant Professor of Aeronautics and Astronautics

Acknowledgements

There are a number of people who have played key roles in helping me achieve my goals during my time at MIT, and I would like to thank all of them. Some have contributed directly to this work, while others have provided me with good company and lasting friendships.

Firstly, my parents deserve more thanks than I could ever give them and their support was crucial for me to complete my Masters Degree at MIT. Their love and encouragement through the years has kept me going and has allowed me to achieve more than I have ever dreamed. Without them, I could not have gotten so far.

My heartfelt thanks go to my thesis advisor Charles Coleman for giving me the opportunity to work with him. He has done a great job of keeping me focused and providing excellent feedback on my work throughout my time here. I am indebted to him for his many brilliant insights and the enthusiasm he displays when sharing his incredible knowledge. The flexibility that he has accorded to me for this research was very much appreciated. My gratitude goes to his assistant Sally Chapman for her patience and invaluable assistance concerning the administrative aspects of my work. My sincere thanks also go to Allan Goode and Russ McMahon of Agent Software for their useful inputs to this work. Based on their industrial experience, they have offered helpful suggestions and guidance especially in the areas concerning agent architectures, computer programming languages and protocols.

I would also like to thank all the graduate students who have made my experience at MIT so unforgettable. To my fellow Singaporeans at the Singaporean Students Association for easing my transition to graduate school and organizing many wonderful activities over the past year; to the soccer dudes for the adrenaline rush every Sat afternoons; to the gang at Tang Hall for all the great times and superb food at our weekly evening socials; and to my housemates Boonseh, Junjie and Jit Hin for the joy of hanging out with three of you.

Table of Contents

List of Figures

List of Tables

List of Acronyms and Abbreviations

ACL	Agent Communication Language
AI	Artificial Intelligence
DAML	DARPA Agent Modeling Language
DAG	Directed Acyclic Graph
FAA	Federal Aviation Administration
HMV	Heavy Maintenance Visit
HTTP	Hypertext Transfer Protocol
HTML	Hypertext Markup Language
JIT	Just In Time
MAS	Multi Agent System
MRO	Maintenance, Repair and Overhaul
NAS	National Airspace System
OEM	Original Equipment Manufacturer
OWL	Ontology Web Language
RDF	Resource Description Framework
RFID	Radio Frequency Identification
SHOE	Simple HTML Ontology Extensions
SOAP	Simple Object Access Protocol
UDDI	Universal Description, Discovery, and Integration
URI	Uniform Resource Identifier
VE	Virtual Enterprise
W3C	World Wide Web Consortium
WSDL	Web Services Description Language
XML	Extensible Markup Language

Chapter: 1 Introduction

Humans are experiencing an information explosion whereby original ideas, concepts and theories continue to evolve each day. We continuously strive to innovate, make new discoveries and conduct our daily lives more efficiently. To do so, we will need to manage and coordinate our growing pool of dynamic knowledge into an ordered and well-planned data structure. With this motivation in mind, many researchers among the artificial intelligence (AI) and computer science community consider software agents and ontologies as being most useful in assisting us to manage the technically complex world we are in the process of creating.

1.1 Agents and Ontologies

An agent is defined as a "computational entity that acts on other entities in an autonomous fashion, performs its actions with some level of pro-activity or reactiveness and exhibits some level of the key attributes of learning, co-operation and mobility[1]". Whereas an ontology is an organized, computer-coded hierarchy of concepts, ideas and controlled vocabularies with related meaning covering a single particular domain.

Agents can be regarded as entities that look after the interests of the parties they represent, and act on their behalf. They behave similarly to your typical travel or housing agent who manages your holiday plans by handling your flight and accommodation, or advertises and markets your property respectively. Agents are however autonomous in nature and do not require further directives in performing their tasks once their objectives have been pre-defined.

It is also relatively easy to destroy and create new agents so as to include additional features as and when they become available. The possibility of replicating agents in the

[1] http://www.cs.tcd.ie/research_groups/aig/iag/pubreview/chap2/chap2.html

thousands or even millions means a whole army of agents can be working together on one single task to form a multi-agent system (MAS).

Agents exploit embedded information inherent in ontologies to make sense of our world and understand the underlying logic, rationale and inter-relationships. In view of the fact that an increasing amount of our work is centered on computer usage, both agents and ontologies are gaining in their popularity for automating processes in the fields of workflow management, network management, air-traffic control, business process re-engineering, data mining, information retrieval/management, electronic commerce, education (Guilfoyle, 1995).

1.2 Multi-Agent System in Airline MRO

The collection of numerous agents that interact with one another to achieve a universal goal characterizes a multi-agent system. The reason behind MAS rests with the prospect of greater processing speed, system reliability and the capability to solve a problem that is too complicated for an individual agent to tackle. The key element that distinguishes between how different MAS operates lies with the way their agents coordinate, negotiate and communicate with one another.

Many MAS solutions that were previously developed in AI research labs have been implemented in areas such as healthcare, logistics and online merchandizing, however applications to the airline industry, particularly the maintenance, repair and overhaul (MRO) sector, have so far been fairly limited.

The purpose of this thesis is to highlight the current problems that are plaguing the airline MRO industry, propose agent architectures for aircraft inspection, repair and supply chain management of aircraft replacement parts and discuss the resulting implications of using agents as a decision-making tool for aircraft maintenance and repair in commercial airlines.

This chapter provides a short introduction to agents and ontologies, and explains the impetus behind employing them in multi-agent systems for airline MRO. Chapter: 2 gives an overview of the MRO industry and its characteristics while Chapter: 3 offers further details about its related problems in maintenance scheduling and spare parts inventory management. The notion of using agent-based applications together with the overlaying National Airspace System (NAS) architecture for air carrier MRO operations, and the way these will all fit together is explained in Chapter: 4.

Chapter: 5 describes the current state of Web services technology and the myriad possibilities it can offer when linked to agent software principles. Chapter: 6 discusses the evolution of the Semantic Web and how it is different from the existing World Wide Web, while Chapter: 7 provides background information on ontologies and its underlying XML technologies.

In Chapter: 8, we describe the way agents work and how it might be possible for them to solve large-scale problems collectively. The MAS solution for airline inspection and repair as well as that for MRO inventory management is proposed Chapter: 9. In Chapter 10, the implications of agent technology in MRO are discussed. Lastly, the recommendations and conclusions of this thesis are presented in Chapter: 11.

Chapter: 2 Background

This chapter describes the organizational structure of a typical airline and the functions of each sub-division. It also gives an overview of the airline MRO industry and explains why maintenance and repair form an inherent component for airlines. Section 2.3 describes the factors contributing to changes in demand and supply of MRO services. In later sections, we emphasize the importance of MRO to the safety of air transportation as well as the positive impact it has on airlines' balance sheets if performed in a cost-effective manner.

2.1 Airline Organization

There are usually 4 main sub-divisions in any commercial airline – flight operations, engineering, commercial and general administration.

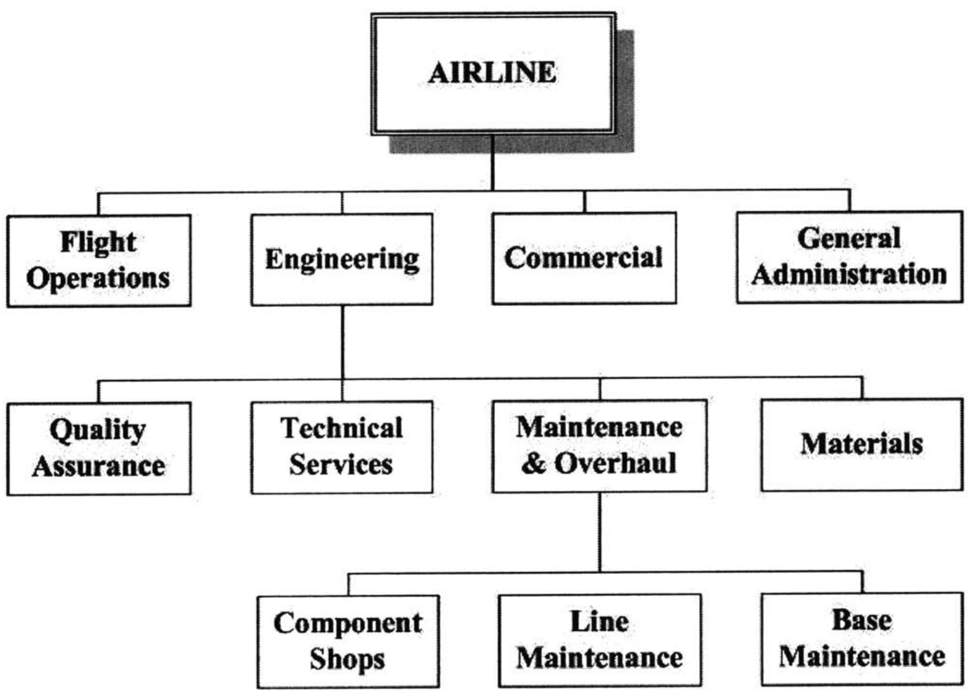

Figure 1: Organization chart of a commercial airline

The flight operations division oversees the daily flight activities and is in charge of scheduling the rosters of the flight crew and flight attendants. The commercial division handles the consumer side of the business such as marketing, ticketing, reservations and other passenger services. The usual corporate functions like finance, human resource and corporate communications come under the general administration division.

Technical support is provided by the engineering division whose quality assurance department checks that all FAA (and other regulatory bodies) regulations or guidelines are followed. It also ensures a certain level of quality standards while operating any aircraft. Technical services department provides the technical expertise and engineering support for any maintenance or modification works. They work closely with line maintenance and airline operations controllers to coordinate aircraft maintenance events, so as to minimize downtime of the aircraft. The materials department procures the required spare parts and maintains an inventory of essential components used during maintenance and repairs.

The actual maintenance work is carried out by the maintenance and overhaul department, whose scope is further divided into component shops, base maintenance and line maintenance. Component shops and base maintenance are located at the airline's home depot, while line maintenance is usually contracted out to local providers in line stations located at foreign airports (Goh, 2003).

2.2 Airline MRO Industry Overview

The commercial aviation MRO industry, which includes airframe and engine maintenance, was estimated to be worth $37.8 billion in 2002[2]. This is a $4.4 billion decline from the previous year due to the fallout from Sept 11 and the subsequent economic recession. It currently has over 100 operators running 31 million square feet of

[2] "Maintenance Repair & Overhaul (MRO) World Market Forecast," Back Aviation and Strand Associates International 2002.

maintenance facilities and employing a workforce of more than 200,000 worldwide. The industry is dependent on the size and utilization rate of commercial aircraft fleets.

MRO activities are generally divided into 4 major sectors:

1) Engine overhaul
2) Component overhaul
3) Line maintenance
4) Heavy maintenance visits and major modifications / retrofits

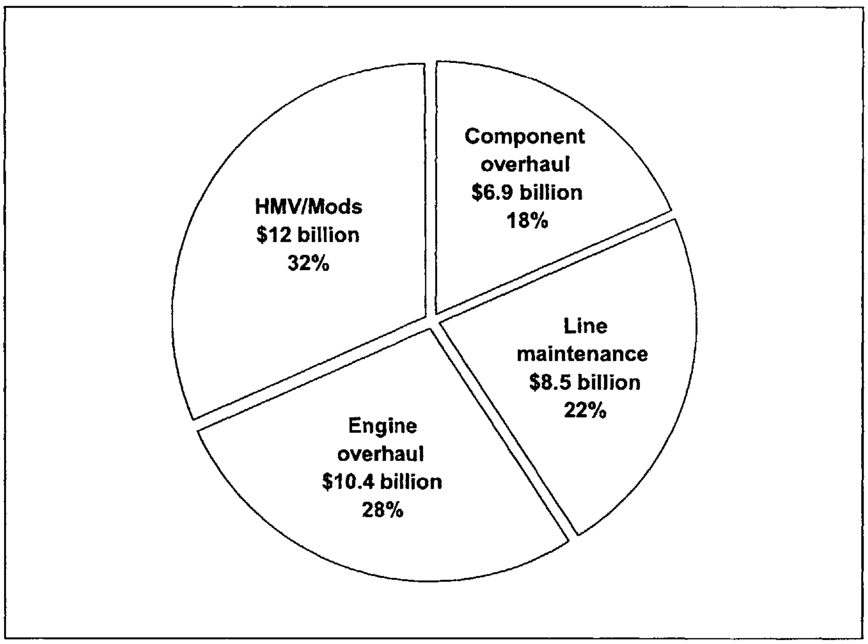

Figure 2: Airline MRO market breakdown (Jackman, 2002)

Traditionally, the key players in the MRO industry were a handful of airlines that possess in-house maintenance capabilities together with some independent MRO service providers. However, original equipment manufacturers (OEMs), especially the engine manufacturers, facing diminishing profit margins from their products and abundant profits in the MRO industry began offering maintenance services in the 1990s. Although the leading global supplier of MRO services was Lufthansa with $1.6 billion, followed by AirFrance with $1.3 billion in 2000, their market shares have been declining steadily.

In recent years, the OEMs have even started to offer overhaul services not only for their own proprietary products but other manufacturers' products as well. They were relatively successful in penetrating the lucrative aircraft overhaul market that used to be dominated by the airlines and now hold approximately 10 per cent market share in North America. 75 per cent of MRO services still remain in the hands of airlines, while independent maintenance companies take the remaining 15 per cent[3].

It used to be standard practice for airlines to provide feedback to OEMs on their products, and periodically report any problems they should encounter during flight operation. However, now that the airlines are in direct competition with OEMs when it comes to offering MRO services, they are thus more reluctant to do so. This lack of co-operation is hampering new developments in technology, especially that of aircraft engines, agent based sensors and radio frequency identification (RFID) spares.

2.3 Demand and Supply Changes of MRO Services

The commercial MRO industry is expected to expand steadily over the next 20 years and generate turnover of between $3 and $4 trillion during this period. It is estimated that this expansion will require an extra 1.4 million to 3.4 million square feet of maintenance facilities to be added over the next 5 to 10 years (Stifflemire, 2002) in the US alone. The 29 major MRO operators in the US currently account for close to 42 per cent of total maintenance capacity worldwide and 78 per cent of the reported maintenance facilities in Table 1 for the Americas.

[3] "North American Commercial and Military Aircraft and Engine Maintenance, Repair and Overhaul Markets," Frost & Sullivan, March 2002.

Size of Facility	Number	Total Sq. Ft.
< 50,000 sf	5	93,587
50,000 < size < 100,000 sf	9	628,595
100,000 < size < 500,000 sf	9	2,231,000
500,000 < size < 1,000,000 sf	6	3,698,000
> 1,000,000 sf	5	10,018,336
TOTAL	34	16,669,518

Table 1: Reported sizes of maintenance facilities in the Americas[4]

The increased demand for MRO services can be attributed to aircraft maintenance growth from budget airlines as well as VIP aircraft. More maintenance jobs can also be expected from conventional airlines, which are looking at ways to extend the life cycle of their existing transport aircraft. Moreover, expansion of cargo aircraft fleets and the conversion of decommissioned passenger aircraft to cargo freighters is likely to fuel an increase in the overall demand for maintenance services. Freight firms such as UPS and Fedex currently spend $1 billion each on maintenance annually.

In the era of globalization, the trend in the MRO industry is to establish smaller maintenance facilities and locate them closer to the customers so as to reduce delivery costs and provide more timely service. Thus, the tenants of US based maintenance facilities are increasingly foreign owned companies, such as Jet Aviation and Evergreen, who are keen to penetrate the US market.

In order to compete more aggressively and attract a larger pool of customers, more facilities are also upgrading their capabilities to handle a mixed fleet of aircraft including both Boeing and Airbus models together with their corresponding components (e.g. engines from Rolls Royce, General Electric, Pratt & Whitney). Presently, only 8.3 million square feet (or 60 per cent) of current US capacity is capable of servicing Airbus aircraft.

[4] Air Transport World Maintenance Directory 2001

In contrast, 85 per cent of capacity is capable of servicing Boeing aircraft. However, these figures will change significantly as more US carriers choose Airbus and more facilities improve their capabilities.

Despite the more rapid growth of air travel within Asia, the MRO market in the US will still continue to grow although US share of worldwide capacity will decline over time. This is because Asian airlines operating flight routes into major US cities will have to expand their MRO capabilities in their respective home bases as well as in the US.

2.4 The Importance of MRO

The key to having a safe and reliable air transportation system lies with imposing strict standards for aircraft MRO to minimize the possibility of inspection and maintenance errors. The highest regard for passenger safety in commercial airlines, coupled with aircraft reliability issues and FAA directives[5] form the basis for all MRO activities.

The significance of MRO on commercial airlines is evident. It has been noted that 15 per cent of worldwide accidents involve a maintenance system error and we currently see about 48,000 US flights being dispatched with at least one maintenance error on board each year (Marx, 1999). In addition, 20 per cent of all in-flight malfunctions are caused by maintenance errors with each in-flight shutdown costing the airline half a million dollars (Rankin, 2000).

The first half of 2003 saw 180 unscheduled or emergency landings according to official FAA records. Of which, 80 per cent were due to maintenance inadequacies such as an engine failure or landing gear problem. In the past ten years, eight out of 14 fatal U.S flight accidents were linked directly to maintenance problems, accounting for 466

[5] FAA, 1991, "Human Factors in Aviation Maintenance – Phase 1: Progress Report," DOT/FAA/AM-91/16.

maintenance-related fatalities. There were 160 deaths in seven maintenance-related failures the decade before[6].

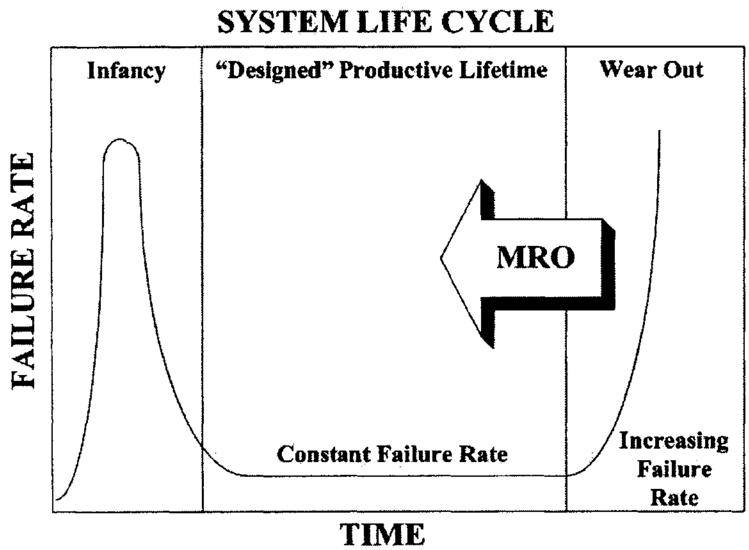

Figure 3: Purpose of MRO

During the lifecycle of a typical system, we find that the initial failure rates are high due to errors in design or production. As illustrated in Figure 3, this early period of infancy will be replaced by an interval known as the "designed productive lifetime" which is characterized by a low constant failure rate. Subsequently, when the system is worn out its failure rate will increase drastically.

The purpose of MRO is to prolong the "designed productive lifetime" of a system by routine maintenance and overhaul, as well as repairing and implementing modifications to the system to drive its operational reliability back to a higher value. Hence, the importance of MRO cannot be overemphasized.

[6] Source: The Charlotte Observer (http://www.charlotte.com)

2.5 The Need for Cost-Effective MRO

Airlines tend to be large corporations that are characterized by high overheads, high revenue but low profit margins of between 1 to 2 per cent. Since it is imperative that airlines ensure that their aircraft are safe to operate at all times during flight, an efficient and cost-effective MRO system is therefore necessary to achieve airworthiness of their entire fleet while maintaining their slim profitability (Heimlich, 2002).

Overheads from MRO activities are not likely to change in the short term, although revenues may vary due to seasonal fluctuations in passenger or cargo loads. In view of the fact that MRO is an essential component of the airline cost structure, there is an invariable need for the MRO industry to seek opportunities in order to reduce costs and render the sustainment of aircraft over a protracted lifespan affordable.

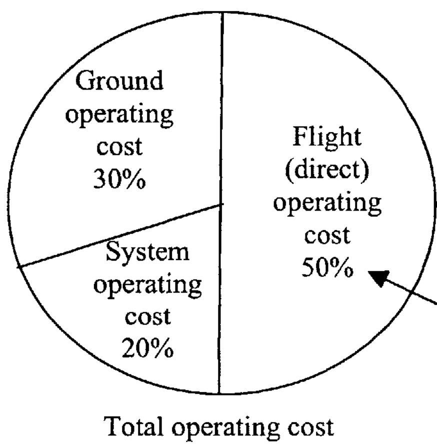

Boeing 757-200 (1999)	Cost per flight hour
Crew	$ 489 (19%)
Fuel	$ 548 (22%)
Maintenance	$ 590 (23%)
Ownership	$ 923 (36%)
Total flight (direct) operating cost	$ 2,550

Table 2: Typical aircraft cost structure (Belobaba, 2002)

Table 2 shows the typical cost structure of a Boeing 757-200 in 1999 (Belobaba, 2002). The flight operating cost relates to all costs concerning aircraft flying operations, and includes pilots, fuel, maintenance and aircraft ownership. Ground operating cost arises from servicing of passengers and aircraft at airports, and includes aircraft landing fees and reservations/sales charges. System operating cost results from marketing or administrative items, and includes in-flight services and ground equipment ownership.

The maintenance cost was given as $590 per flight hour according to Form 41 data provided by the U.S. Department of Transportation. This amount was equivalent to 23 per cent of total flight operating cost, which in turn translates to 50 per cent of total operating cost. Thus we can conclude that 11.5 per cent of total operating cost arises from MRO and therefore there is a need to make MRO activities as efficient as possible to keep overall costs down.

Apart from reducing MRO expenditures per flight hour, it is also necessary that any maintenance work performed on an aircraft to be done in the shortest time feasible so as to maximize the utilization of an aircraft during its lifetime. Any increase in the downtime of an aircraft incurs an opportunity cost because of lost revenue it could have generated if it was flying. There are also substantial operations re-scheduling costs when aircraft do not return from maintenance as planned.

2.5.1 MRO as Profit Center

In the past, MRO activities are usually viewed as cost centers and assimilated with "captive" maintenance operations of major airlines such as United Airlines, Delta Airlines, American Airlines and Air Canada. They tend to be huge maintenance facilities with large number of workers (see Table 3). When these major airlines seek to improve their bottom lines by adopting cost-cutting measures, they have a tendency to scale back their MRO operations to reduce expenditure.

Service Provider	Employment
United Services	15,800
GE Engine Services	11,300
Delta TechOps	11,000
American Airlines	9,000
Air Canada Technical Services	8,000

Table 3: 5 largest MRO employers in North America[7]

The difference in the efficiency between independent service providers such as Aviation Management, Dee Howard, Jet Aviation and a captive maintenance operator such as Delta TechOps is highlighted in Table 4. Although Delta TechOps have a slightly higher turnover compared to the other 3 independent service providers, it however employs more than 13 times the number of workers employed by Dee Howard and its total maintenance facilities measure more than 3 times larger.

Service Provider	Employment	Total Sq. Ft.	Sf per employee
Delta TechOps	11,000	2,700,000	245
Aviation Management	820	800,000	975
Dee Howard	750	850,000	1,133
Jet Aviation	650	608,966	936

Table 4: Comparison of captive and independent maintenance operation

As a consequence to their inefficiencies, some of the airlines are revamping their MRO operations along the lines of independent service providers by getting rid of their past legacies, shedding workers, reducing extensive maintenance space and offering state-of-the-art technology in maintenance products and services to a diversified customer base.

These revitalized airline MRO setups now do not have to negotiate with powerful unionized workers, which used to number in the thousands, over wages, work conditions

[7] Air Transport World Maintenance Directory 2001

and employee benefits. They tend to be more efficient and the resulting cost advantage gives them a competitive edge and many operate as profit centers for the airlines.

2.6 Maintenance

Maintenance is the act of restoring or replacing components that are faulty or have deteriorated during its service lifespan. It is crucial in order to sustain the integrity and performance of an aircraft (Hessburg, 2001) and comprises of inspections, overhauls, repairs, preservation and replacement of aircraft parts.

A whole list of activities characterizes the maintenance of an aircraft, including disassembly, repair and inspection of its engines and all other major structures and systems. Mechanics in the service shops focus on a number of devices present in the electrical, avionics, hydraulic and pneumatic sub-systems. Other personnel provide cleaning services of carpeting, seat covers, draperies, plastic and wood components used in the aircraft cabin.

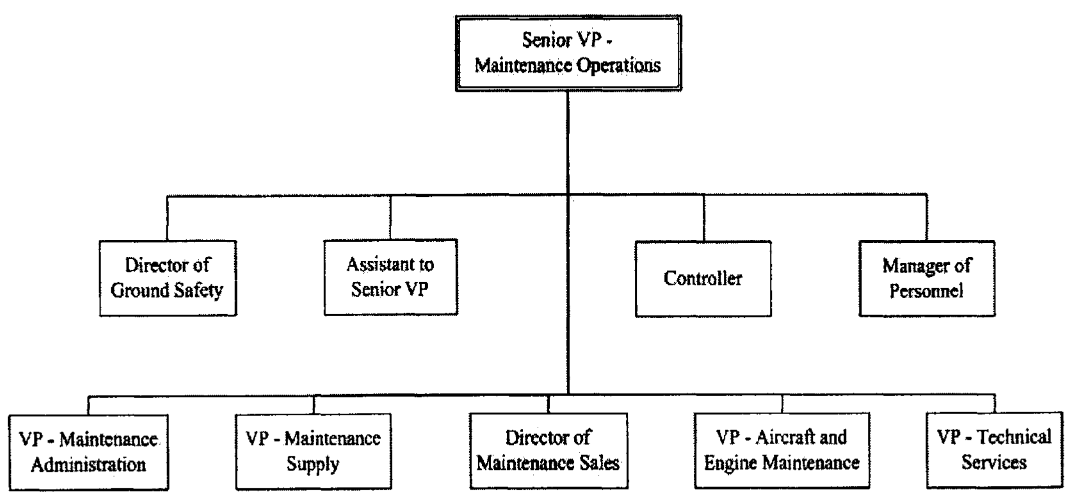

Figure 4: Maintenance organization of a typical airline

The Maintenance Administration unit is tasked to plan and formulate fleet-wide maintenance schedules while the Maintenance Supply department is responsible for the provision of spare parts, equipment and supplies across the entire aircraft fleet. Aircraft and Engine Maintenance is in charge of conducting aircraft overhauls, repair and inspection works on the airframes, engines, components, and ground power units.

Technical Services monitor and audit the quality of maintenance work done on all aircraft and ground equipment as well as offer its technical expertise. It also collaborates in R&D efforts with OEMs of aircraft, engines and components in developing new equipment. Therefore, it is accountable for technically supporting this equipment in service.

2.6.1 Usage vs Calendar-based Maintenance

The aircraft undergoes usage-based maintenance after undergoing a certain number of flight hours. For calendar-based maintenance, staff decides on the time interval between successive maintenance, and the aircraft is called into service regardless of the number of hours it has flown. Basic and intermediate levels of maintenance are classified as usage-based and these tend to be carried out near the airport where the aircraft is stationed. Advanced maintenance is usually calendar-based, and performed at specialized depots.

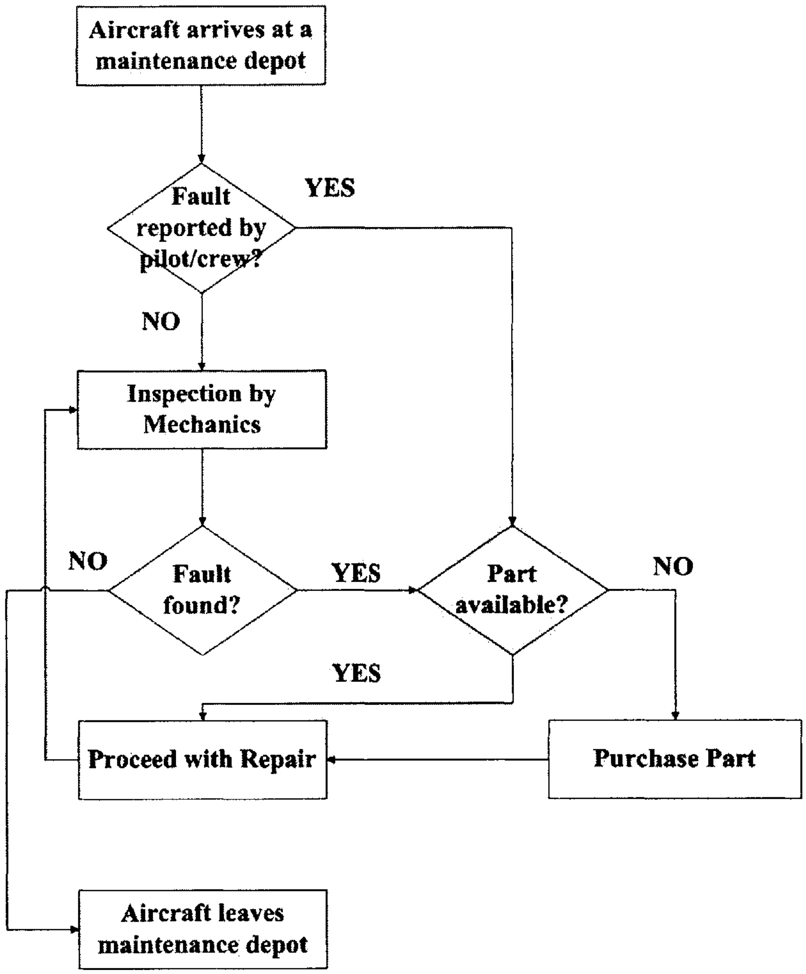

Figure 5: Flowchart of maintenance

The maintenance workflow in a typical depot is illustrated in Figure 5. Faults are detected either through inspections by maintenance mechanics or notifications from the pilot and flight crew. After assessing the repair work necessary to rectify the problem, the aircraft replacement parts will be sourced internally from existing inventory or from external suppliers. Depending on the severity of the glitch, the aircraft might have to remain grounded until the part becomes available and repair is performed. Otherwise the aircraft might be allowed to return to operational status, and the repair could be done during the aircraft's next scheduled maintenance.

2.6.2 Scheduled vs Unscheduled Maintenance

Examples of scheduled maintenance are:

- Periodic (e.g. Annual)
- On Time (Time Between Overhaul)
- Progressive (Inspection based, e.g. Cracks)
- Conditional (Monitoring based, e.g. Engines)
- Heavy Maintenance Checks

Unscheduled maintenance includes:

- Reported anomalies otherwise known as "squawks"
- Line replacement units
- Airworthiness directives, service difficulty reports

For a Boeing 767, it typically undergoes a variety of checks:

(a) Layover/overnight checks are quick inspections of the aircraft's general condition and are usually done every 3-5 days.

(b) 'A' checks are overnight line maintenance visits focused on the aircraft's general condition and limited system checks for critical systems like emergency equipment. Built-in test equipment is used for troubleshooting of any fault detected. These checks are usually conducted after each interval of 200-500 flight hours or about 30-50 days (Hansman, 2003).

(c) 'C' checks are hangar visits focused on detailed checks of individual aircraft systems and thorough inspections of specified components. Systems checks are performed to ensure the system performance is satisfactory, otherwise restoration is conducted to allow the system to operate until the next letter check. Minor structure inspections and minor cabin restoration are also done. 'C' checks are usually conducted after each interval of 6-8 months and take between 3-5 days to complete.

(d) Heavy maintenance visits (HMV) or 'D' checks are hangar visits in which the aircraft undergoes a major overhaul, including detailed inspections of the airframe with the paintwork removed. The aircraft structure is inspected in detail while the majority of the interior components, including the seats, galleries and lavatories are removed and refurbished. Facilitating the HMVs involve the opening of the wing fuel tanks, removing fairings and opening access panels that would otherwise not be done in other lesser checks. Structural repairs for corrosion and fatigue cracks are performed if necessary, as well as major modification works. 'D' checks are usually conducted after each interval of 5-8 years.

The purpose of these checks is to perform routine and non-routine maintenance of the aircraft. Known problems and defects reported by the pilot, crew and line inspectors are rectified. They also allow for the replacement of parts that are nearing the end of their lifecycles and are worn out after a certain number flight hours.

Chapter: 3 MRO-Related Problems

Given that the objective of this thesis is to propose an agent-architecture to improve airline MRO activities, it is therefore necessary to first understand the underlying problems found within the industry. Two major problems were identified and they involved resource allocation and maintenance scheduling, as well as inventory management of aircraft replacement parts.

In this chapter, we provide the reader with a background of some of the intrinsic and recurrent problems as well as those that emerged after the terrorist attacks of Sept 11. Section 3.1 covers the obstacles encountered by most MRO practitioners after the implementation of cost cutting measures by airlines facing declining passenger loads and heightened security checks in the post 9/11 period. In sections 3.2.4 and 3.2.5, the core difficulties of the airline MRO industry are explained.

3.1 Post 9/11

The commercial aviation industry has experienced a number of recent setbacks due to both the economic slowdown in 2000 and the terrorist attacks in 2001. Passenger loads fell drastically in the aftermath of 9/11 as people opt to postpone their necessity for air travel, or reduce it by looking for alternative modes of transports. Although public confidence in our air transportation system is slowly returning after end-2002, the airline industry has so far lost over $25 billion in revenues since the beginning of 2001.

3.1.1 Cost Cutting Measures

In light of these developments, some airlines that were already underperforming even before the terrorist attacks had gone into bankruptcy administration. Others that still remain in the black are barely surviving the deterioration in global air traffic numbers, and are desperately finding ways to lower their costs. Understandably, most airlines choose to cut their MRO expenditure because that will have the most immediate effect to their financial positions. As a result, total maintenance spending went down by more than $1 billion from 2000 to 2002. During this time, maintenance spending per flight was trimmed by an estimated 4 per cent across all airlines.

These cost-cutting measures in the airline industry have affected certain aspects of MRO operations. Large numbers of mechanics, inspectors and engineers have been laid off, aircraft parts inventories have shrunk, and entire depots have been scaled down or shut permanently. In doing so, the airlines have achieved their objective of balancing their books, by bringing down their costs so that they are in line with diminishing revenues. However, the fact that MRO expenditures of most airlines have decreased generally as a proportion of their total operating costs does not necessarily mean that the remaining MRO assets are operated in the most productive manner.

On the contrary, having fewer MRO bases imply that some aircraft might have to fly a further distance to arrive at the next service depot that is capable of providing the necessary maintenance support and possessing the required aircraft replacement parts. There is a need to juggle with the constraints of having fewer depots, while continuing to offer the same high standards of maintenance essential to satisfy air safety standards. This demands a good automated system to administer real-time maintenance scheduling and inventory management. The optimized solution is often yielded through numerous iterations based on dynamically changing information.

3.1.2 Longer Time Between Maintenance Checks

Along the same theme of cutting costs, some airlines are resorting to extending the period of time between scheduled maintenance checks in order to obtain greater cost savings. They choose to defer maintenance and not perform repair on non-critical components, so as to allow the aircraft to operate for a longer duration with the same level of maintenance. Having a longer interval between scheduled maintenance checks also allow airlines to operate with fewer but better-equipped MRO depots and reduce their overall inventory of parts.

However, the decision to lengthen the time between maintenance checks is often viewed as a short-term solution to the problem of costs. There is always an upper limit to how much we can increase the amount of time before a worn-out aircraft component needs to be replaced. On the other hand, if certain aircraft parts have not received regular maintenance during their working life, they can be harder to repair when they do fail. This would translate to higher maintenance costs in the long run, although short-term costs are kept deceptively low.

3.2 Perennial MRO Problems

The next few sections will describe some of the persistent problems besetting the airline MRO industry. Although these problems are not new and have constantly plague the trade for years, the recent events in the aviation industry signify a greater urgency to examine these problems more closely. There is now an unprecedented pressure to come up with innovative ideas and solutions so as to increase productivity and possibly perform MRO activities at a lower cost.

3.2.1 Varying Age of Aircraft and MRO Personnel

The typical lifecycle of a commercial aircraft ranges from 20-30 years while new aircraft models are continually introduced to the market at the rate of about one per decade - the latest models being Boeing's 7E7 and Airbus's A380. Therefore, at any one time an airline would usually have a fleet composing of a mix of aircraft types or even competing models from different manufacturers. Needless to say, the age of different aircraft in the fleet varies accordingly.

It is more difficult to detect cracks or defects in older aircraft compared to younger ones. For that reason, scheduled maintenance accounts for only 30 per cent of all MRO activities in these aircraft compared to between 60 and 80 per cent for their younger counterparts[8]. More inspections are required for older aircraft and any repair takes a longer time. These are often done outside of scheduled maintenance and are carried out at a higher cost.

It is also harder and hence more costly to source for parts in older aircraft, which are often based on more obsolete technology and such parts could be out of production. Therefore, maintenance costs will generally rise with the increasing age of aircraft in an airline's fleet. This point should be noted when making cost estimates of future MRO activities. However, younger aircraft have their fair share of additional maintenance requirements because revolutionary airframe composites necessitate more advanced inspection and analytical tools during maintenance.

Besides the aircraft, MRO personnel will also age over time. When the MRO labor force renews itself, experienced mechanics, inspectors and engineers are replaced by a set of younger but less experienced workers who may lack the technical expertise and knowledge of the previous generation. These younger workers may not have received sufficient training and thus are less skillful and adept to work on a wide range of aircraft. This is particularly true if the younger workers are assigned to maintain and repair an older aircraft that they are less familiar with. Additional costs will need to be incurred to train the younger workers so that they can work on the different models of aircraft available in the airline fleet. This "age-wave" issue requires the capture of both tacit and explicit knowledge of workers in readily repurposed form before it walks out the door.

3.2.2 Human-Related Maintenance Errors

Faulty and poor quality of aircraft maintenance can be attributed to various elements, but most notably is that of human factors which account for approximately 12 per cent of major aircraft accidents and 50 per cent of all engine-related flight delays and cancellations (Graeber, 1994). Human-related maintenance errors are considered

[8] FAA, 1993, "Human Factors in Aviation Maintenance – Phase 3, Volume 1: Progress Report," DOT/FAA/AM-93/15.

relatively common, although they are often overlooked by airlines that rely increasingly on technologically advanced diagnostics and condition-monitoring tools as means to reduce maintenance errors.

In the early 1990s, when pilot errors and weather issues frequently contributed to the cause of accidents, there was much focus placed on improving terrain following and weather tracking devices in the aircraft. Therefore, human-related maintenance errors are not accorded the same level of attention back then. Even today, there is not much focus in this area. Out of FAA's FY 2004 budget of $100 million for R&D efforts, only $87,339 is spent on improving in aviation safety (i.e. the level of inspection and maintenance technologies).[9]

A possible reason contributing to the incidence of human-related maintenance errors is the fact that many maintenance personnel are overworked, according to an Aircraft Maintenance Survey[10] conducted by the Australian Transport Safety Bureau in 1998. Since aircraft repairs and ordering of replacement parts can only be done after the initial aircraft inspection process, it is of paramount importance that inspections are completed as soon as the aircraft arrives at the maintenance depot.

Inspections are thus often carried out between the last flight of the day and the first flight of the following day. Given that the work requirements are only discovered after the inspection process, hence the overall MRO workload for each aircraft can be unpredictable and this variability of work content poses a problem. Frequent overtime and prolonged working hours coupled with poor scheduling of maintenance personnel increases the possibility of mistakes occurring.

[9] http://www.faa.gov/aba/html_budget/files_pdf/BIB-2004.pdf
[10] www.atsb.gov.au/aviation/pdf/msurv.pdf

	Maintenance Errors Resulting from Human Factors
1.	Collision of an aircraft under tow with foreign objects such as aerobridges and empennage of other aircraft.
2.	Installation procedure incomplete despite having all required parts present. For example, nuts are 'finger tight' instead of being correctly tightened by tools.
3.	Injury to personnel due to unsafe contact with hazardous material. For example, exposure to aircraft fluids, harmful chemicals and electric shocks.
4.	Stationary aircraft damaged by vehicle or maintenance equipment such as stairs or movable stands.
5.	A component not installed or assembled correctly.
6.	System operated dangerously during maintenance. For example, flaps not properly prepared for activation or unsafe to operate due to proximity of personnel and equipment.
7.	Maintenance tools or material left behind in aircraft.

Table 5: Human-related maintenance errors

The duration of a work stint most commonly reported in the survey was 12 hours, but over 10 per cent of respondents indicated that they had worked for over 20 hours at a stretch at least once during the past year. These workers are likely to lose focus and concentration over long spells of work, giving them a higher chance to make mistakes. The table above lists the types of maintenance errors resulting from human factors.

The survey also cited work pressure, fatigue and lack of coordination and training as the most common reasons why a human-related maintenance error occurs. Due to fatigue these personnel cannot effectively operate the equipment on maintenance facilities, while a significant number lack sufficient training and experience to do so. The outcome is often incorrect assembly, orientation or incomplete installation of aircraft components, and persons contacting hazards.

In addition, some workers frequently fail to adhere to procedures due to the lack of proper maintenance manuals or documentation, and as a result maintenance systems are being operated in a dangerous manner from time to time. For instance, aircraft systems such as thrust reversers were activated during maintenance when it was not safe to do so, in some cases because personnel or equipment were not clear of the area. To make things worse, there is a lack of strong focus on maintenance because mechanics, unlike pilots, are less prominent as they belong to several dispersed unions. All these factors lead to human-related maintenance errors.

3.2.3 Independent and Competing MRO Depots

The overall MRO organization of an airline consists of several maintenance depots (or hubs) that are scattered across the continent or world, depending on whether the airline operates domestic or global flight routes. There is a need for a universal presence to reach out and offer maintenance to its entire fleet of aircraft. Due to geographical practicalities, there is often limited contact among the depots in terms of personnel or parts movement. However, one of the other reasons for the lack of interaction arises from inter-depot rivalry.

Although their functions are similar, each depot however operates separately from one another and has its own supply, transportation and maintenance personnel. Supply personnel decide which parts to stock and in what quantities, while transportation personnel overseas the movement of the parts. It is the responsibility of each depot to decide its own unique operation model. A depot has to establish for itself the most suitable ordering policies, decide on inventory levels and location of storage areas, and define the kind of circumstances before it is allowed to call for outside technical expertise from the manufacturers or suppliers.

As a result, an individual depot will vary in their contracting, engineering, functional structures and financial management procedures. These divided and disjointed depots cannot achieve the potential synergy of operating collectively and significantly degrade the overall performance of the supply chain (Lee, 1997). They do not have a shared pool of parts and have to resort to keeping separate inventories, thereby expending valuable capital and resources in doing so.

3.2.4 Maintenance Scheduling

Scheduling of airline MRO activities presently involves a fair amount of guesswork and reliance on previous work scheduling practices that have proved feasible in the past. There is currently a lack of an adequate decision-support tool to manage the complexity involved in planning, scheduling and decision-making in MRO activities, while taking into account maintenance staff levels, parts inventory, location of service depots, aircraft fleet size and routing of each airline.

Part of the reason for poor maintenance scheduling is because some airlines run their information systems in an alienated manner and divide them according to their numerous functions such as ticketing, flight scheduling, part inventory, personnel roster and aircraft maintenance history. There is no common platform to collate all these different sources of data and make them available across the entire airline organization.

Common information from one system cannot be immediately reused on another, without prior replication of it into a readable form in the latter's system. These incongruent systems indicate that shared data cannot be effortlessly transferred between the different constituents of the MRO set-up. This therefore makes it extremely challenging to coordinate and integrate the various entities that exist within the MRO depots such as the tools, equipment, machines, spare parts and human resources, with the larger world outside.

Consequently, MRO decision-making is usually done at the depot level with its business model and operational philosophy not entirely orientated in a similar way to that of the airline, but rather become more focused on the wellbeing of the MRO provider itself. As a result, decisions are made without a clear idea or full picture of each and every constraint confronting the airline, and justifications for each MRO action are based on a cost/benefit analysis that does not fully represent all the airline interests. There is no mechanism to ensure that the most suitable maintenance procedures are executed in the shortest possible time, so as to maximize productivity and cost-effectiveness.

To add to the problem of maintenance scheduling, many of the existing MRO management systems are passive and are hardly predictive or automated in nature (Adams, 1993). They are not capable of conceiving future events and often require human intervention when an unexpected incident occurs, which in turn induces human errors.

Current mechanisms are also not responsive enough to react to rapid changes to business conditions and other external factors, so as to modify or update processes to reflect the current situation. Moreover, there is limited flexibility in work scheduling to cater to possible staff preference for working hours, length of works shifts, public holidays and weekends. Workers now can only choose to work fixed shifts and specific days.

3.2.5 Management of Aircraft Spare Parts

Most MRO practitioners lament the fact that parts are often not on hand when and where they are required. There are a number of reasons for this phenomenon and they broadly fall into three categories such as variability in demand, supply chain issues, and lack of proper inventory management tools.

Needless to say, these incidences of part shortage lead to maintenance delays and might even ground the aircraft for a significant period of time until the part becomes available and replacement is performed. Inventory management typically involves a trade-off. Having high levels of inventory increases holding costs, but on the other hand it improves responsiveness because more parts are available to meet unanticipated demands.

3.2.5.1 Variability In Demand

Aircraft failures are generally unpredictable and they tend to occur more frequently during flight operation than otherwise. Even though the expected time to failure of a particular part is proportional to some known and measurable quantities such as flying hours or number of take-off and landing cycles, part failures are still difficult to foresee. The advent of MRO prognostic and condition-monitoring tools, especially those applicable to aircraft engines, has improved the capability of maintenance personnel to anticipate the occurrence of a failure, however it is still tricky to accurately determine the lifespan of a single aircraft component due to inconsistency in material characteristics, varying manufacturing methods and flight conditions.

A strong volatility in the demand for aircraft replacement parts accompanies such uncertainty in failure predictions. Although MRO depots typically stock a certain quantity of spares to offset seasonal fluctuations in demand, it is still highly probable for an unusual peak in demand to create a shortage for a particular part. When that happens, there is usually no viable alternative but to place a new purchase order and wait till the next delivery to take place. The key strategies in managing variability in demand of aircraft parts lie with good forecasting and having a sufficiently large inventory of parts to tide over the period of excessive demand.

Having a large enough inventory for all aircraft replacement parts can be costly. Although common items such as spark plugs, tires, instruments, hydraulic lines, etc., tend to be inexpensive and have customary use rates that are determined through work orders, purchase orders and vendor recommendations, the core dilemma lies with resolving which of the more expensive parts to stock, their quantities and location (King, 1986).

3.2.5.2 Supply Chain Issues

MRO depots are sometimes weighed down by a prolonged time to perform standard aircraft maintenance and hence incurring a backlog of repair jobs, due to shortages of relevant replacement parts. Instances of spare part shortages can be directly traced to glitches in the supply chain, which can be defined as the framework whereby resources are acquired and processed by intermediate companies or manufacturers, and subsequently delivered as finished products to the end customer.

It details the transition of raw materials to manufactured goods and the participants include vendors, manufacturing facilities, logistics providers, internal distribution centers, distributors and wholesalers. The distinction between different supply chains lies with their number of nodes and the complexity of the links.

The current supply chains for some aircraft parts are not sufficiently responsive and the waiting time to fulfill a part request or source a part can be quite long. The problem of spare part shortage is magnified by the fact that some airlines operate older aircraft that tend to require more maintenance per flight hour, hence need more parts. Furthermore, those parts that are based on out of date technology can be harder to source (Bickel, 2003).

Another reason for the low part availability stems from the fact that some MRO facilities have limited knowledge of all suitable part suppliers out there in the market, and thus purchases their parts from those they are more acquainted with. They then enter into long-term contracts with a single manufacturer and inherit the inflexibility in switching between suppliers. This over-reliance on one supplier also means that parts can be more expensive because the single supplier may be adopting unfair business practices and artificially inflating the price of its product.

In order to prevent part shortages, it is necessary for a MRO warehouse to have adequate storage capacity. It is common practice to have large inventories for low cost/high demand parts that are usually dispersed and stored in several locations. Conversely, high cost/low demand parts tend to number only a few. They are stowed away at a central site and transported to a particular depot when the need arises. To further add to the problems of MRO supply chain, aircraft parts can be costly and easily damaged without a good transportation and storage system.

The other anomalies in the supply chain are characterized by inadequate quality control from manufacturers, insufficient inspection checks on the part of the MRO personnel and late deliveries. The first two issues suggest that a large number of parts received by the MRO facility may in fact be faulty, but yet remain undetected and used to perform a repair. This gives rise to grave concerns with regards to flight safety should the fault be a source of a catastrophic failure. On the other hand, late deliveries will lead to even longer periods of part shortages so that part requests cannot be processed immediately, and compound the problem of extended repair times.

3.2.5.3 Inventory Management Tools

Aircraft parts are conventionally listed in a registry of serial numbers[11], each represented by a unique serial or Radio Frequency Identification (RFID) number. Depending on the type of sub-component the part is attached, the same part can have different serial numbers when used in different sub-systems or applications. One persistent problem of some inventory management tools occurs when one common part can be used for a number of machines, and the tool does not correctly indicate the right quantities of similar parts available and their locations.

For example, when one blade design is suitable to be fitted to four different aircraft engines models, although the actual inventory might be one but four such parts are reflected by a less competent inventory system. It fails to recognize that the part has already been designated a particular application and hence serial number. Alternatively, the blade could have been pre-assigned for use in a particular engine model. When an urgent part request arrives for one of the three other applications, the management tool does not alert the MRO personnel that the inventory actually holds a similar part that can be used potentially. A good inventory system of maintenance supplies and parts should be robust enough to handle such situations.

[11] For example, the Department of Defense assigns National Stock Numbers (NSNs) to each part for the purpose of identification.

Other flaws of some inventory management tools include not adequately accounting for part failure in the MRO shop by monitoring items with short shelf lives. Moreover, the tool might have non-autonomous and cumbersome restocking procedures that usually require a warehouse manager's approval when issuing a new invoice order for a part (Dada, 1984).

Chapter: 4 NAS Initiatives for MRO

Some of the MRO-related problems that were highlighted in the previous chapter have been identified as barriers preventing the realization of a safer and more effective air transportation system. To resolve the matter, the key stakeholders in the airline industry such as the air carriers, independent MRO providers and government agencies, particularly the Federal Aviation Administration (FAA), have been working closely together to examine these issues and propose new strategies (Schofield, 2003).

Under the FAA, the last couple of years have seen the National Airspace System (NAS) exploring new ways and ideas to upgrade its infrastructure and undergoing a series of modernization programs. Among the many concerns of the NAS, there is a strong urgency to improve airline efficiency by enhancing integration of normal flight activities with its MRO operations (Phillips, 2003).

In this chapter, we present some of the NAS initiatives including the adoption of agent-based solutions and ontologies to automate some of the business logic and processes in MRO depots. Section 4.1 provides an overview of the NAS and its function while section 4.2 describes in greater detail how the proposed agent architecture will fit into the overall scheme of things to come in the NAS.

4.1 National Airspace System

The NAS is a complex collection of systems, facilities, aircraft and personnel, including pilots, tower controllers, terminal area controllers, maintenance personnel and airline dispatchers. Together with their computers, communications equipment, satellite navigation aids and radars, they form the NAS. Its task is to ensure the safe and smooth traffic of aircraft through US airspace.

The rationale for modernization of the NAS is to cope with the gradual upswing in demand for air travel predicted in the years to come, following its sharp decline after the events of Sept 11. There are currently 18,292 public and private airports within the NAS that handled close to 640 million emplanements during last year alone, and this number is expected to rise to over 1 billion in the next 5 to 10 years[12].

NAS modernization encompasses work involving the airport facilities, air traffic control systems as well as the airlines so as to upgrade its reliability and capability. Its work also includes improving the business practices and logic of maintenance, repair and overhaul (MRO) activities of air carriers using software agent principles, which is among the many new NAS initiatives.

4.1.1 NAS Architecture 5

As part of its modernization efforts, the NAS has published a blueprint known as Architecture 5 that describes in detail the objectives, tasks and actions to be taken by various parties belonging to the airline community. Released in November 2003, the proposal is a wide-ranging, multi-year plan for advancing the capability of NAS to collaborate with airlines in planning and conducting flights with greater safety, flexibility and efficiency. The plan takes into account the decision-making processes and finances of the FAA, along with integrating many of the different entities' interests, plans and programs.

Scheduled for implementation in the 2012 to 2015 timeframe, the architecture together with its associated systems and facilities will form the "National Airspace System Concept of Operations and Vision for the Future of Aviation". It provides the baseline for future planning to meet the challenges of providing the public safe, secure, and efficient air travel. The NAS Architecture will also guide the evolution of automation systems, surveillance systems, navigation systems, communication systems, and infrastructure replacement to ensure future modernization challenges are met.

[12] http://www.faa.gov/education/resource/National Airspace System (NAS).htm

4.2 MRO Overlay Agent Architecture and Ontologies

Many of the applications and databases within the NAS, particularly those in the airline MRO industry, exist in isolation. There is an imperative need to connect these disparate entities and automate enterprise business logic across different applications so as to provide broad and overarching collaboration. The commercial aviation industry and NAS are slowly beginning to address this problem. As mentioned earlier, one of the initiatives put forth by the NAS entails the gradual adoption of agent-based solutions and ontologies to automate some of the business processes in MRO depots. Agents and automated processes lend themselves to on-demand real-time scheduling and skills matching, bringing about the optimized use of resources and enhanced throughput.

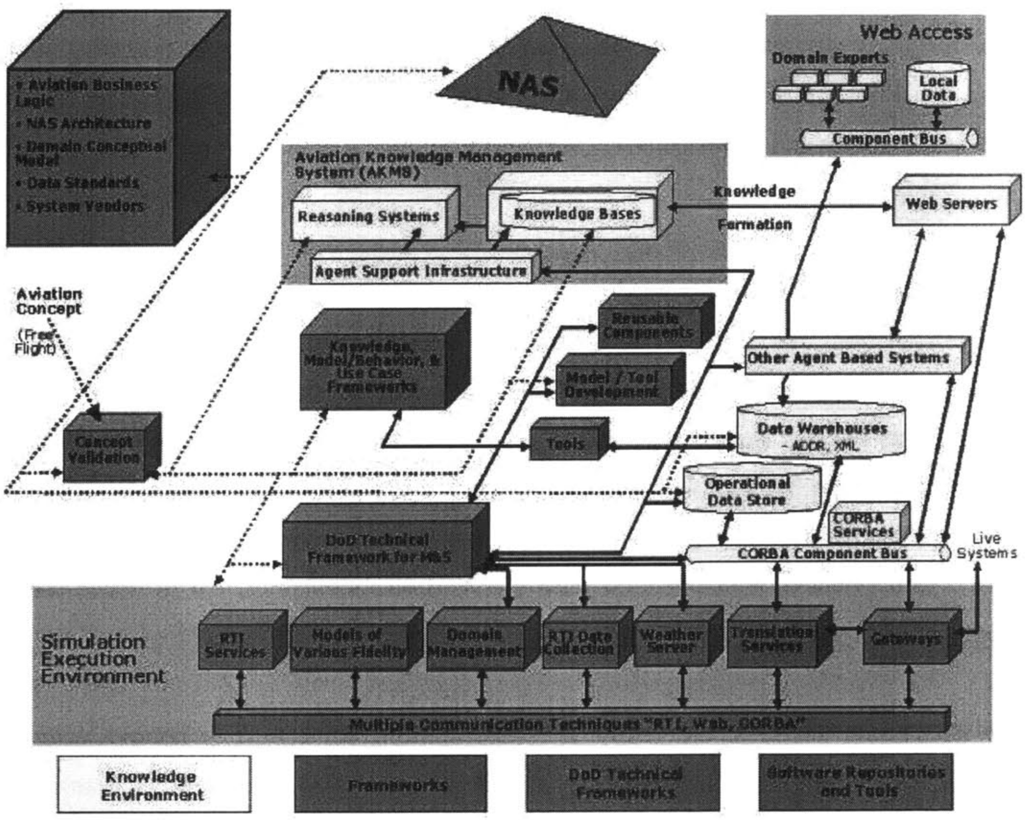

Figure 6: Overlay agent architecture interoperates with the NAS vision[13]

[13] http://www.nas-architecture.faa.gov/Architecture/hierarchy.cfm

Essentially, we are looking to detail an overlay MRO agent architecture illustrated in Figure 6 that will interoperate with the NAS vision of having fully-integrated air carrier activities in order to drive optimization and make MRO activities more efficient. This next generation system can ensure accurate and timely information exchange between all flight and aircraft maintenance personnel. In order to do so, we will require the development of a number of ontologies to accurately model the MRO business, which can be more rapidly achieved through the modification of other existing ontologies.

4.3 MRO IT

Ontologies provide a platform for agents to access information and perform their respective tasks. For this reason, they play a crucial role for the creation of an agent-architecture to automate some of the airline MRO business processes by providing the IT infrastructure to do so. The use of ontologies and an appropriate rule/logic base, coupled with data mining and analysis, allow significant improvements to airline efficiency to be achievable[14].

With ontologies, participants need not change existing information systems to plug into the information tree through Extensible Markup Language (XML) namespaces and ontology relationships. The advent of XML that serves as the *lingua franca* of modern information systems has the potential to significantly reduce errors by provisioning data to all value chain partners that is both consistent and current. This is an area where we can establish the inter-relationships of safety inspections, finance, budgets, audits, routes and reports/squawks with retention schedules and documentation. All of which can be approached in an interoperable modular E-Records fashion.

By developing universal ontologies, it is also feasible for agents to interact with one another in solving large and complicated problems in MRO activities. Such an agent-

[14] Working with Agent Software (a California-based company with expertise in software agents), we understood the importance and principles behind the development of a comprehensive ontology which details and captures information on aircraft scheduling, constraint ranges, business logic, mechanic skill sets, logistics linkages as well as inventory arrangements.

architecture will have extensive applications, both on and off the Internet. The potential of the evolving "cognitive Web services", which is described in the next chapter, bodes well for the possibility of collaborative reasoning and decision-making.

Some of the potential benefits include:

- Real time prediction and analysis of aircraft maintenance
- Highly relevant maintenance alerts and notifications
- Real time maintenance costing analysis
- Advanced decision support system
- Integration of external information including in-flight sensors or flight logs
- National Airspace System (NAS) integration could provide a rich set of information sharable with other systems and value chain partners
- Information Capture and central repository of distributed information and business logic
- Workflow Management (prediction, alerts, decision support, supplier integration)
- Workforce Collaboration (forums, calendaring, alerts, information sharing)

Chapter: 5 Web Services

The creation of universal ontologies within an agent-architecture will initiate the growth of "cognitive Web services" for airline MRO businesses. Web services can be thought of as a means to buy or sell goods and services on the Web, while the term "cognitive" implies some form of collaborative reasoning and decision-making. Web transactions involve at least two parties, namely the supplier of a particular service or product and its end-users. In our context, requests for MRO services as well as aircraft spare part purchases can be made through Web services in the future once the infrastructure for its implementation is well developed.

In this chapter, we will define what are Web services, the role they play and their limitations. A description of the Web service framework is offered in section 5.2. In section 5.3, we discuss how current standards of Web services can be improved while section 5.5.1 allows us to discover the ways in which Web services can be tied up with software agents to open the door leading to the formation of virtual enterprises.

5.1 Definition and Background

Web services are one of the recent developments in the field of information technology. They are considered to be an amalgamation of the Web, XML, isolated databases and other distributed systems. It is among the many ways to efficiently bring together transaction systems, electronic data swapping and business-to-business (B2B) services. As defined by the World Wide Web Consortium (Lea, 2003), Web services are

- comprehensible using normal Web protocols
- recognized by Uniform Resource Identifiers (URIs)
- allow XML-based messages to be sent, received or acted upon
- capable of interacting with applications and programs that are not directly human-driven user interfaces

The ingenuity behind Web services is the use of HTTP-like protocols to make network processing available to programs. These include one-off services that execute a command or return some data, in addition to more complicated services involving a string of commands executed in a particular order and timing.

Prior to Web services, programming languages such as Corba and Java RMI are used to build network applications. However, these architectures lack Web services in two aspects. Firstly, a pool of knowledge is required between the service programmer and the client programmer. This meant that altering a service would entail changing the shared knowledge together with the client's profile, even though the alteration does not impact on any feature of the client. Secondly, the HTTP-like protocols used by Web services have the ability to bypass firewalls unlike Corba or Java RMI-based technologies (Filman, 2003).

Web services were made possible by the ability of XML tools to deliver requests and response messages across servers. Supported by all software companies, it is becoming increasingly attractive as a mode for data swapping and transaction systems, which can be universally accepted and approved by the IT community involving application designers and developers.

They are a key component of Microsoft's .Net initiative and are deemed essential to the business directions charted by IBM, Sun and SAP (Huhns, July 2002). Web service frameworks can also be used to improve current infrastructure and applications by combining and assimilating them to the extent that was not previously possible due to opposing and often non-compatible operating standards or systems.

5.2 Web Services Framework

There are three broad areas in the Web services framework: communications protocols, service descriptions and service discovery. Each of these areas has their own specifications, namely Simple Object Access Protocol (SOAP), Web Services Description

Language (WSDL) and Universal Description, Discovery, and Integration (UDDI) respectively. Collectively, they are viewed as the three foundation standards of Web services.

5.2.1 Communication: SOAP

The SOAP[15] language defines a format for passing messages between Web services and describes how components can interoperate even if the components are written in different languages. Before SOAP, the main method of cooperation between distributed components was to use remote procedure calls (RPC). When a local process needs the assistance of a remote process, RPC is used to make the connection and takes care of passing the values the processes require.

RPC has some limitations that led to the SOAP protocol. RPC works best when the processes are developed in the same language and using the same RPC tool kits, but more important it was difficult to have RPC work across firewalls and gateways common in networks. SOAP addresses the issues of RPC and also allows a good way to pass message-based data.

This can be accomplished because all the parameters and functions that need to be called in a remote process are described in the body section of the SOAP envelope (Box, 2000). To get across firewalls, the SOAP envelope is carried in the data section of a HTTP packet. Since all but the most secure networks will allow web (HTTP) traffic, this means that SOAP services have no problems running inside and outside firewalls. SOAP satisfies the main requirements of being platform-independent, secure and lightweight.

SOAP is a minimal set of conventions for invoking code using XML and HTTP. It is simply an application of XML to HTTP and invents no new technology. As such, SOAP leverages the engineering effort already invested in HTTP and XML technologies by codifying the application of the two in the context of remote method invocation.

[15] For more on SOAP, check out www.w3.org/TR/SOAP/ or www.develop.com/soap or www.soapware.org

5.2.2 Description: WSDL

Jointly developed by IBM and Microsoft, WSDL[16] is an XML format used for describing a Web service based on an abstract model of what the service offers. It characterizes the way to access a Web service, specifies the location of the service and the type of operations that will be performed.

Services are defined as collections of network endpoints, or ports in a WSDL document. These ports are capable of exchanging messages via the SOAP language. WSDL service definitions provide documentation for distributed systems and act as a guideline for automating the elements involved in applications communication.

5.2.3 Discovery: UDDI

It is a mechanism for discovering service providers similar to having a computerized registry or "phone directory" of Web services. UDDI describes the framework and operation of an index of services, by defining the way to encode information of each service, and how this information can be accessed and updated.

There are two common forms of directories - white pages and yellow pages, which are established by how the entries are sorted in the directory. White pages have their entries sorted by name, while the entries in yellow pages are sorted by their characteristics and capabilities. Less common are green pages that include technical data about the services.

UDDI[17] is based on XML and SOAP, and is a web service in its own right. It offers white, yellow and green-pages services, but is not advanced enough to provide any brokering service.

[16] More information on WSDL can be found at www.w3.org/TR/wsdl or xml.coverpages.org/wsdl.html
[17] Resources for information on UDDI – www.ibm.com/services/uddi or www.uddi.org or www.uddicentral.com

5.2.3.1 Simple Directories

Directory services provide an example of a Web service. They allow participants and components to find and detect one another. Components and participants might be agents, Web service providers, Web service requestors, applications, people, objects or procedures. A simple directory might contain relevant explanations or descriptions on the type of services each participant can offer. These data and information are stored in the directory, and can then be accessed using search queries by other participants wishing to make use of that particular service.

5.2.3.2 Complex Directories

A more sophisticated directory might be more proactive by providing more than just a search service. Instead, it could be a brokering or supporting service. For example, a participant might want to employ one or more agents to handle a particular query. It can do so by using a brokering service that is capable of recognizing the aptitudes of registered service providers, so as to decide which providers are suitable to forward a query to. It would then send out the query to those providers and return their replies to the initial requestor, at the same time discovering the characteristics of the answers returned.

5.3 Current Standards for Web Service

Web services denote the publication of machine-readable information detailing the use of a service and the way to utilize it using other programs. They depend on the functionalities of find, publish, and bind, and all interactions are via an agent-communication language (ACL). See Figure 7 for illustration. ACL is the common language software programmers use to communicate and swap information, and many of them are based on XML technology.

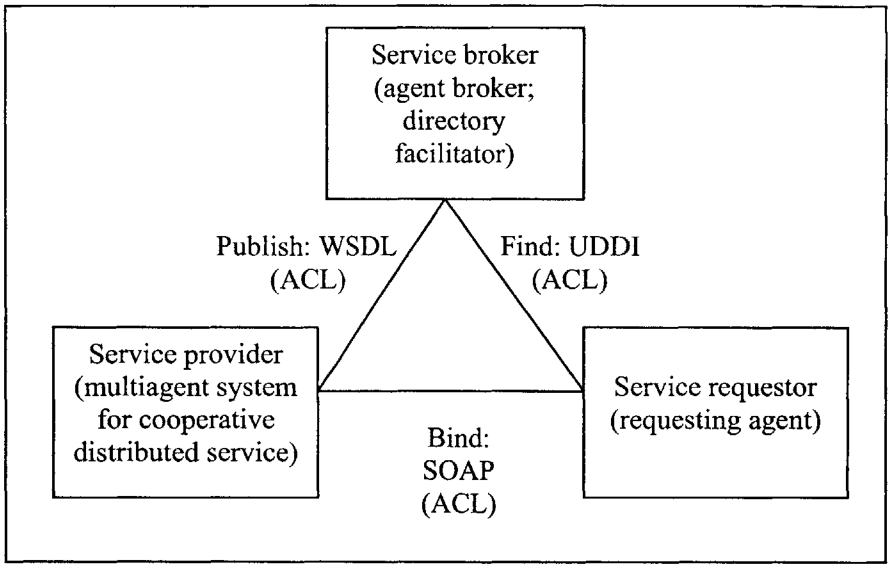

Figure 7: Current practices for Web services (Huhns, 2002)

SOAP provides the common protocol systems need to communicate with each other. It can be used for services request such as to schedule appointments, order parts and deliver information (Box, 2000). Describing the services in a machine-readable form is WSDL, which specifies the names of functions, the required parameters and results. UDDI provides a way for users and businesses, known collectively as clients, to find needed services by specifying white or yellow pages of services.

Apart from accepted norms for XML, SOAP, WSDL and UDDI, Web service providers have to conform to the semantics of specific domains. This is achieved by software such as Resource Description Framework (RDF), the DARPA Agent Modeling Language (DAML) and, more generally, ontologies. More details on the above can be found in Chapter: 7.

5.4 Shortcomings of Web Services

Despite Web services becoming more prevalent, they still pose a significant challenge to the many software designers to develop the necessary tools to support them. For instance, they will need to improve on the space, bandwidth and processing time of existing XML-based technologies. In addition, the sophistication of the software components has to be reconciled with its dependability so that future work can be done using better tools to generate messages, bindings without having to program in the underlying "assembly language" of raw XML.

Efforts to make the Web more computer-friendly via Web services have so far met stiff resistance. The limitations of Web services meant that they are only aware of themselves and not their surroundings (i.e. users, clients or customers). Web services are also unable to resolve or reconcile incompatible ontologies used by different parties and are dormant until being called upon. They will not offer alerts or updates when a new situation arises.

The reason why Web services cannot behave independently is due to a lack of explicit semantics. They cannot comprehend responses and messages in each other's service or what duty each service carries out. Web services also cannot represent business logic, rules and relations in a computer-readable format.

5.5 Benefits of Improving Web Services Infrastructure

Web services standards are determined by their interoperability standards and protocols for performing business-to-business (B2B) transactions. One way to improve standards of Web services is to have them act as independent objective-based agents, that dynamically seek out fellow agents in order to complete a particular task as well as capable of collaborating with each other on their mode of interaction and contact.

By improving current standards of Web services, we can establish a set-up in which programs serve as agents within networks of Web services. These programs will facilitate automated business transactions, by procuring and contributing information via these Web services.

Moreover, supplementing semantics to Web services infrastructure will enable business rules and logic to be explicitly described and allow the meanings of exchanged messages to be understood. This makes it feasible to comprehend the responsibilities of individual Web services and awareness of such services can be made possible by automated service discovery. Web services can also be combined to form more complex services.

5.5.1 Virtual Enterprise

With improved standards of Web services, it is possible for the formation of a Virtual Enterprise (VE) at a short notice that can perform a wide range of tasks and responsibilities. For instance, a group of participating companies and organizations can come together to offer their respective core competencies such as to devise an inventory management system or bid for a construction project. Capabilities are assembled as and when they are required and there is no idle time. This nimbleness and flexibility is the main attraction behind VEs, which allows it to cut substantial costs and time to market.

As the name suggests, VEs are intangible organizations that are formed by means of software based on a set of selection principles and criteria. Buyers, suppliers, markets and associates are found spontaneously to accomplish a certain mission without earlier accord. For that reason, companies that are part of VEs can draw on its worldwide contacts and connections to compete for ad-hoc special projects. This system of real-time dynamic VEs will mean business supply chains would be faster and projects can be completed more quickly and often at a lower cost because of a lack of intermediaries.

The concept of the VE is not new. As early as in the late 1980s, Marty Tenenbaum had proposed a "sea of services" on the Internet that will make VE a reality (Petrie, 2003). As such, the recent emergence of highly advanced technology in Web services has led many to ponder how this newly developed knowledge can be best applied to realize the full capability of VE as a business model. Although it may still require further research into the viability of software agent principles, there are some who propose using them to build on Web Services standards and ultimately realize the potential of a VE.

Chapter: 6 The Semantic Web

Web services are essential to the idea of a Semantic Web, which is what the World Wide Web is converging towards. The goal of the Semantic Web is to automate Web-document processing by creating a collection of computer-readable data. It forms a platform for agents (we shall further discuss the definition, roles and potential of agents in Chapter: 8) to operate, whereby they can contribute their competency and expertise to the Web.

This chapter will highlight the key differences between the Semantic Web and the World Wide Web. The various languages used in Semantic Web creation are described in section 6.3 while the associated complexity involved in such creation is explained in the following section of 6.4.

6.1 Semantic Web vs World Wide Web

The main difference between the World Wide Web and the Semantic Web is the fact that both of them are primarily designed for different set of intended users. The World Wide Web was envisioned for human use, and its human-readable pages are thus incoherent to computer programs. As a result, the Semantic Web was developed so that digitized information on the World Wide Web can be made available for computer use and understanding. The Semantic Web provides tools for explicit markup of Web content and is thus more accessible to agents (Berners, 2001).

The Semantic Web has a wide range of ontologies (defined in Chapter: 7) that other information resources can commit to. When two or more sources commit to the same ontology, they are obliged to use the same terms and relationships as defined by the ontology. A source can either commit to an existing ontology or create a new one. The Semantic Web is therefore a decentralized way of creating standard vocabularies.

Up to now the Web we have is a model centered on having data and information, in the form of text and graphics, posted on web pages. Each page has a number of links that guide the human user from one page of information to the next. All pages can be accessed and viewed in turn. The Hypertext Markup Language (HTML) is used for encoding, determines the manifestation and appearance of the page but not its substance. In comparison, software agents are concerned about the substance but not the appearance.

Despite its limitations, the current World Wide Web is still used by some agents today. The shopbot, which is an agent that browses through an online merchandisers' product catalog and return the sale price of similar items in the search request is one such example. Shopbots operate by a form of screen-scraping in which they first seek out the name of an item of interest from the downloaded catalog.

They then proceed to search for the adjacent set of numerals that has a dollar sign, which is assumed to be the unit price of the item in question. Finally, the name of the item and its corresponding price are returned to requester. The use of Semantic constructs, such as ontologies represented in well-established languages, will allow agents to comprehend and identify what is on a page more easily.

6.2 Semantic Web Vision

The vision is to have a Semantic Web capable of inferring data from Web content and thus converting the Web into an Internet-wide information representation system. Researchers of the Semantic Web have endeavored to realize the dream of having dynamically interoperating nodes capable of B2B e-commerce, instead of the current static pages in the World Wide Web.

The Semantic Web will need agents that are capable of understanding the content and obtaining logical inferences, with ontologies forming the basic structure for interpreting such information provided by Web pages. In the future, it is likely that the Semantic Web will offer Web services the information it needs through formal languages, and the logic about business rules and service descriptions through ontologies.

6.3 Semantic Web Language

Semantic Web languages are used to annotate Web pages, making its content accessible to both agents and computer programs. They allow the creation of ontologies that specify standard terms and machine-readable definitions. They also present methods of extending existing ontologies, giving their creators the option of including domain-specific information.

There are a number of Semantic Web languages, each with their own attributes. These range from earlier versions such as Simple HTML Ontology Extensions (SHOE) and Ontobroker to newer ones like DARPA Agent Markup Language+Ontology Interchange Language (DAML+OIL) and Ontology Web Language (OWL).

DAML+OIL is the end product of synthesizing the best attributes of previous Semantic Web languages and is essentially an expressive description logic with a RDF syntax.

6.3.1 DARPA Agent Markup Language for Services (DAML-S)

DAML-S offers a possible answer to the problem of providing semantics for distributed search of Web services, and by doing so provide the missing link between Web services infrastructure and the Semantic Web. Basically, it exploits a standard taxonomy (or ontology) of common terms and semantically annotated Web pages in order to provide semantics for service operations and messages.

DAML-S is both a language and ontology, and Web service providers can use the markup language constructs contained in DAML-S to explain the characteristics of their services in a machine-readable form. The last version, DAML-S 0.9, was released in 2003 and subsequent releases will be based upon OWL that follows from a further simplification and clarification of DAML+OIL. OWL is currently developed by the Web-Ontology Working Group at the World Wide Web Consortium (W3C) and has recently become an official W3C recommendation[18].

6.4 Difficulty in Semantic Web Creation

The complexity of Semantic Web creation lies with annotating Web pages with Semantic markup languages, and the resulting ability of the pages to deduce consequence and be knowledgeable. It is widely regarded that a standardized Semantic Web language and the Semantic Web itself will be very useful in helping us get work done. However, achieving both of these will not be easy nor can it be done in a short time.

There are a number of challenges that we face which includes converting information into a suitable format, and using the Semantic Web to describe and compose Web services. We also need to be able to scale the Semantic Web technology to handle "Web size" data, devise ways to deal with incoherent information, as well as to build up the users' trust in performing business transactions over the Semantic Web.

Among these obstacles lie the most critical problem of them all – the creation of ontologies. The skills necessary to design an ontology is not readily available in our labor market, and current tools such as Protégé (Noy, 2001) have their limitations. We will need to set up and maintain huge ontology libraries to promote sharing in the community and reduce duplication of efforts[19].

Another problem we have is the fact that various users often adopt different ontologies to describe the same domain. To ensure interoperability of these ontologies, we will need to develop the technology to translate, align and merge them. Finally, Semantic Web ontologies have to be updated and modified from time to time, and it is essential for a good mechanism to manage and track such changes.

[18] http://www.daml.org/services/owl-s/

[19] The DAML Web site (www.daml.org) already has 200 or more DAML ontologies, but more can be done in terms of the sophistication of the libraries' search capabilities.

Chapter: 7 Ontologies, XML and its Extensions

Ontologies form the backbone and building blocks of Web services that constitute the Semantic Web. Since our daily lives involve immense amounts of data and information, it will be valuable if we can somehow categorize and order this information in a meaningful way. Ontologies serve this purpose well and make it possible for knowledge to be shared and re-used among multiple parties in a highly efficient manner (Noy, 1997).

We provide an explanation of what constitutes an ontology in this chapter and offer an explicit example of an ontology that is related to aircraft specifications in section 7.2. In section 7.3, we give an overview of XML and its extensions of RDF, because XML is the core upon which most standard ontologies are based.

7.1 What is an Ontology?

An ontology is an organized hierarchy of concepts, ideas and controlled vocabularies with related meaning covering a single particular domain. A good definition would be that "it is a logical theory that accounts for the intended meaning of a formal vocabulary" (Guarino, 1998). Although there are other definitions of an ontology applied in other fields and some even differ in their manifestations, but many of the concepts overlap and the essential view of a hierarchy with beliefs is almost universal (Heflin, 2000). Most ontologies have a taxonomy of terms and an expression to describe the terms and their relationships.

A well-organized ontology is able to fulfill this role by supplying useful meanings and associations to the content present in the ontology itself. When adopting a new ontology in order to enhance business processes, the most intricate part is the development of the ontology itself. However, once it is defined and accepted as the operations standard by all parties involved, it straightaway becomes a common language that can streamline communication across the various entities. It is clearly advantageous to have an excellent system for organizing information as well as for all trading partners and separate department units to synchronize with one other. In other words, having ontologies will allow everybody to "read from the same page".

The break up of data and process is a result of current trends in computing and information systems. This means that data-driven applications and architectures (such as ontology) are becoming more detached from process, especially with the advent of XML technologies. Data that changes often are separated from process, so that the process code can remain the same. This reduces code bugs and makes the application more flexible and adaptable. An ontology that is a major part of a data driven architecture, is typically expressed using XML.

An automated process such as an agent can operate on the contents of an ontology. Such autonomous agents may broker among themselves to carry out various tasks based on the resources available in the ontology. More on how agents can tap on the resources available in ontologies will be emphasized in Chapter: 8.

7.2 Example of an Ontology – Aircraft Specification

In a simple ontology of concepts related to an aircraft, we might find the following terms: craft-type, commercial jet, 767, craft-configuration, range and fuel-capacity. Organized in a hierarchical manner, the more general concepts in an ontology are placed in the higher levels while more specific concepts are assigned to lower levels. Hence, "craft-type" would be at the top of the hierarchy since it is the most general concept. Directly under that will be attributes of the class (i.e. commercial jet, 767), followed by the "craft configuration" (i.e. range and fuel-capacity).

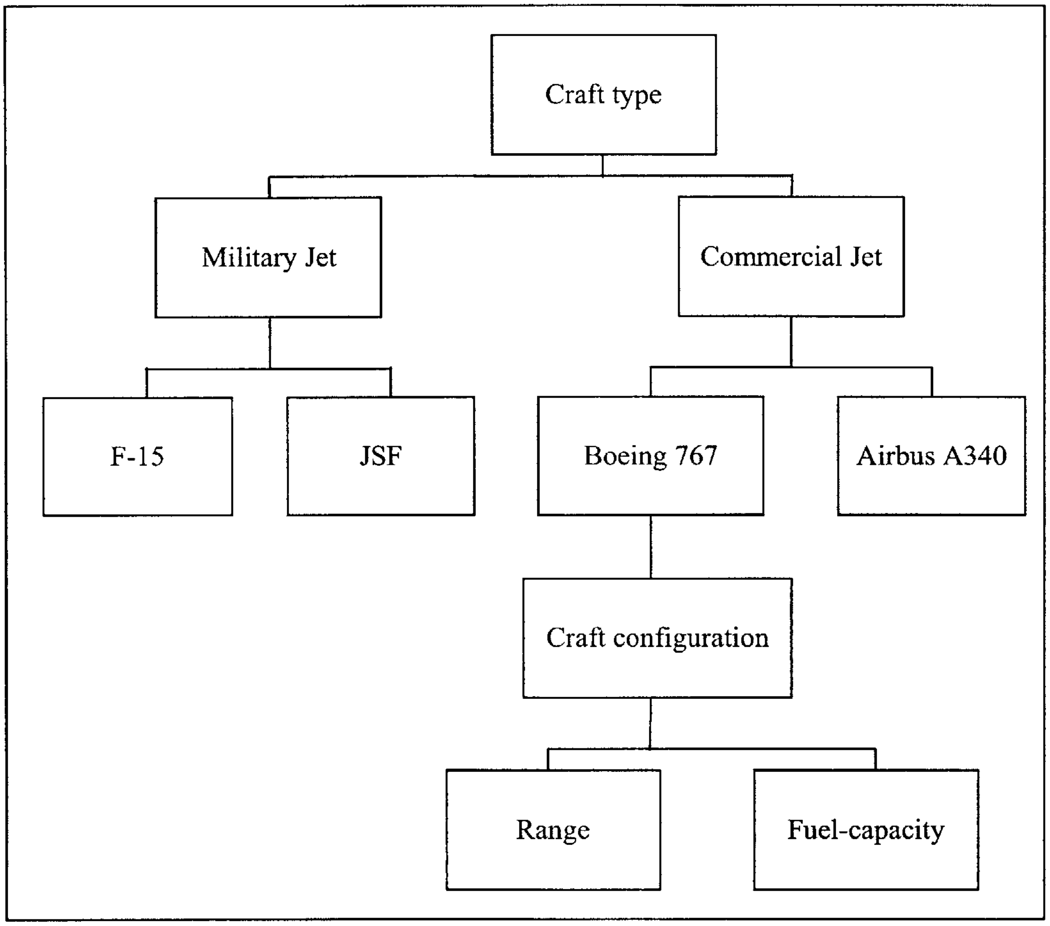

Figure 8: An example of an aircraft ontology

The above organization in Figure 8 illustrates an important feature of ontologies – concepts have relationships between them. In the above case, 767 would be a child of "commercial jet" and is a more specific concept than its parent. Both would be of a craft-type class, exhibiting what is called an "is-a" relationship with craft-type. The craft-type class may also include a "has-a" relationship to "craft-configuration" with attributes such as range and fuel-capacity.

Due to the complexities of ontologies, the simple outline-like hierarchy used in the example is usually insufficient to capture the full domain of knowledge, so the rules in ontology creation allow for concepts to have multiple parents along with multiple children. To expand the original illustration, 767 may have a parent military jet as well as commercial jet. This is allowable as long as the concepts increase in specificity as they go lower in the hierarchy and that no concept is the parents of its own ancestor. That constraint confines these ontologies to what is called a directed acyclic graph (DAG) structure.

As a hierarchy of related concepts, ontologies provide an ideal framework onto which data and information can be organized. In our example, the term 767 is linked to the concepts. An ontology with its associated data can be processed to verify the concepts and also to make inference about the content. Sophisticated tools can be loaded with the ontology to provide for a level of autonomous reasoning. These tools are called rules engines and form the basis of a knowledge base.

7.3 XML-based Ontologies

Since XML is the core upon which most standard ontologies are based, it is important to give an overview of XML, its extensions (i.e. RDF) and their purpose.

7.3.1 XML Background

A trend in software design over the last decade or so has been a shift from developing everything in one bucket or monolithic programming, to development using discrete parts or component based programming. The advantages of using parts to assemble an application are many fold but most importantly, the individual parts can be reused using parts that are well-tested. This will lead to less buggy applications and any changes in the code can be isolated to only a few parts of the entire application (Walsh, 1998) unless the bugs are in design. So having an application built up from discrete parts is a good thing, but there needs to be a glue layer that holds them together.

Although any interface or communication protocol could act as the glue, a standard system would obviously be better. By this time most developers were used to working with HTML and its simple tags that tell a web browser how to render a page. Hence, it was only natural that a close relative but more flexible solution be chosen as the glue to connect the various parts of an application. XML is more flexible and process-orientated than HTML, which is good for components that need to work together and do not necessarily provide user interfaces. XML also serves as an excellent data transport and data storage format.

7.3.2 XML Technology

XML uses elements (tags) and attributes in the same way as HTML, whereby tags like <body>Some Text</body> have defined meaning to a browser. However, there is no definition of what must go in the tags (Wilde, 2003). Developers are thus free to describe their own languages. There are RDF, SOAP and WSDL just to name a few. The XML model that is described by the language is even verified and shown to conform through the use of Schema which itself is XML.

With all XML's powerful features it is still just a data describing language and often needs to be coupled with a procedural programming language like Visual Basic or Java to build applications. However having a universal and flexible language allows and supports the use of standard tools to process and work with XML data. These tools essentially operate the same way but are available on the programming platform of choice. The latest XML (1.0 Third Edition) recommendation by the World Wide Web Consortium (W3C) was released in February 2004, and can be found at http://www.w3.org/TR/2004/REC-xml-20040204. This document contains examples of XML codes and its purpose is to draw attention to the specification of XML syntax and to promote its widespread deployment.

7.3.3 RDF Background

Resource Description Framework (RDF) is a XML vocabulary and schema that provides methods of describing data such that its meaning is fully described within a domain. RDF is the main mechanism through which the Semantic Web and the efforts of the W3C will be realized (Mitra, 2000).

As described above, developers using XML are free to develop all kinds of languages and data formats. This is great and certainly a powerful feature but it makes it hard to make any assumptions about what the language means or what the information is trying to convey. Applications and processes thus need to become tightly coupled to the data in order to function.

Also it is often the case that the focus of our inquiry is not the data itself but data about the data. For example, maybe a comprehensive document was created about a particular aircraft's flight characteristics and of particular interest is when was this document created and by whom. The "when and by whom" is data about data, and in this case it is data about the flight characteristic document. RDF was developed to address these types of issues, mainly clarifying information through the use of metadata.

7.3.4 RDF Technology

RDF views the world as a series of statements containing a subject, a predicate and an object. Using the example above, the flight characteristics document would be the subject and created-by would be the predicate, while the object would be the name of whoever created the document. The statements are called a triple for obvious reasons.

The statements can be combined in various configurations but the triple is always maintained. A nice feature of RDF is that the statements are easily displayed as graphs with nodes for subject and object, edges for predicate. A powerful feature of RDF is that it presents an object-oriented view of XML data through types that are classes or property types (Bray, 2001). The latest RDF Vocabulary Description Language (1.0 RDF Schema) recommendation by the W3C was released in February 2004, and can be found at http://www.w3.org/TR/2004/REC-rdf-schema-20040210. This document contains examples of RDF codes and its purpose is to specify RDF vocabularies, classes and properties so that RDF's role as a general-purpose language for representing information in the Web can be enhanced.

DARPA Agent Markup Language (DAML) and Ontology Web Language (OWL) are extensions to RDF that provides greater meaning to the contained ontology content. The extensions provided by DAML and OWL are important parts of enabling distributed agents and services. OWL has recently been adopted by W3C and now replaces DAML and DAML-S.

Chapter: 8 Agents on the Web

Having described ontologies and agent architecture in the earlier chapters, it is important for us now to understand the characteristics of the agents themselves. The creation of agents, which are visualized to roam freely on the web, is heralded as a major breakthrough in artificial intelligence. It is believed that their combined logic and analytical ability could eventually be utilized for decision-making. Coupled with their enormous computational muscle, it may be possible for agents to handle huge amount of information and solve problems that entail logic, calculation, and knowledge. Research work done at Stanford and Berkeley universities had illustrated that agents can be used in this capacity[20] (Turkett, 2003).

In this chapter, we offer an insight into the way agents interact with one another and how collaborative reasoning comes about. In addition, the issue of agent security and scalability are discussed in sections 8.4 and 8.5 respectively. Examples of some existing agent applications such as in the healthcare and logistics sectors are described in section 8.6.

8.1 How Agents Work?

It is conceivable in the future to have thousands if not millions of agents to be engaged simultaneously in order to solve complex problems. These agents might have unique preferences and different reasoning abilities, hence making centralized planning impractical. In such a situation, there must be some form of distributed planning so that the many distinct agents can work together to accomplish a common task.

Humans' proficiency in solving distributed planning problems stems from our shared values and principles. We use our common beliefs to maintain social harmony, law and

[20] Projects such as SETI@home (http://setiathome.ssl.berkeley.edu/), which processes radio telescope signals and Folding@home (http://folding.stanford.edu/), which analyzes protein folding.

order in our civilized societies. Equivalently, agents can adopt a similar approach to govern their agent societies based on analogous norms and ethics. To do so, they will need a methodology to convert these theoretical societal values to utilities that can be communicated in a computational framework.

In Figure 9, we see how it is possible to translate an agent's shared abstract principles into a set of preferences so that the agent can then use decision and utility theories to rationalize its available options. It will subsequently proceed with a course of action that best emulates its preferences. Hence, colossal distributed planning and implementation at different levels of granularity (Rose, 2001) can be achieved by incorporating distributed planning and with an agent's inherent social values.

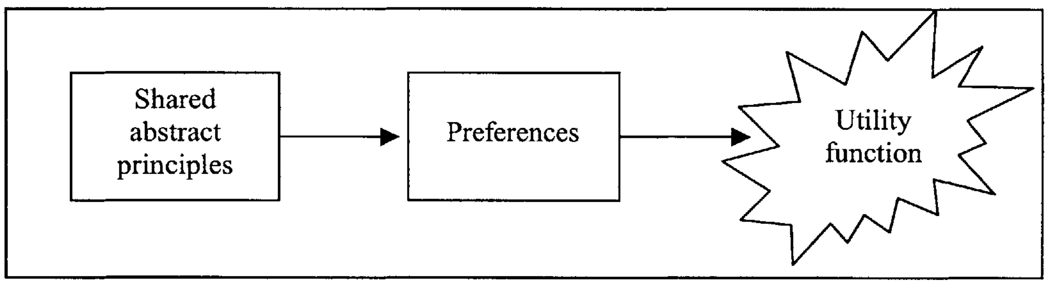

Figure 9: Mapping from shared abstract principles to utility functions (Turkett, 2003)

8.2 Collaborative Planning and Execution

There are three phases in an agent-architecture, namely coalition formation, coalition plan development, and commitment management. Coalition formation arises when an agent is unable to carry out a job that calls for an ability that the particular agent lacks, or a certain resource that is not available to that agent. This phase involves sub-stages such as discovery, proposal and reply. In the discovery stage, a requesting agent will seek out a task agent through several means such as a directory service or broadcast messages.

Upon finding a suitable task agent, a proposal will be sent to it. Evaluating the job description and constraints with its own inherent abilities, the task agent will either accept the job if it deems favorable to do so or reject the job otherwise. If the job is rejected, the stages of discovery, proposal and reply have to be repeated until the requesting agent is able to locate a task agent willing to take up the assignment.

Once a task agent agrees to take up a job, it enters the coalition plan development phase whereby the plans of numerous agents are consolidated and examined. A cost-benefit analysis is done to ensure the global and individual goals of the agent community are met with the minimum cost (opportunity cost of the task agent fulfilling this job). This analysis is done using decision and utility theories, coupled with the agents' social principles.

When done properly, the analysis acts as a mechanism to inhibit unproductive actions that will not lead to overall social improvement. Before carrying out the plan, agents prioritize their tasks to avoid bottlenecks and form teams to deal with dynamic changes in information, tasks, number of agents and their capabilities. The last phase of commitment management also forms part of the benefit-cost analysis.

Figure 10 shows a typical flow chart depicting the various processes involved in an agent-architecture. The coalition formation phase may require several iterations of the discovery, proposal and reply stages, depending on the ease of finding a suitable task agent. After which the phases of coalition plan development and commitment management follow.

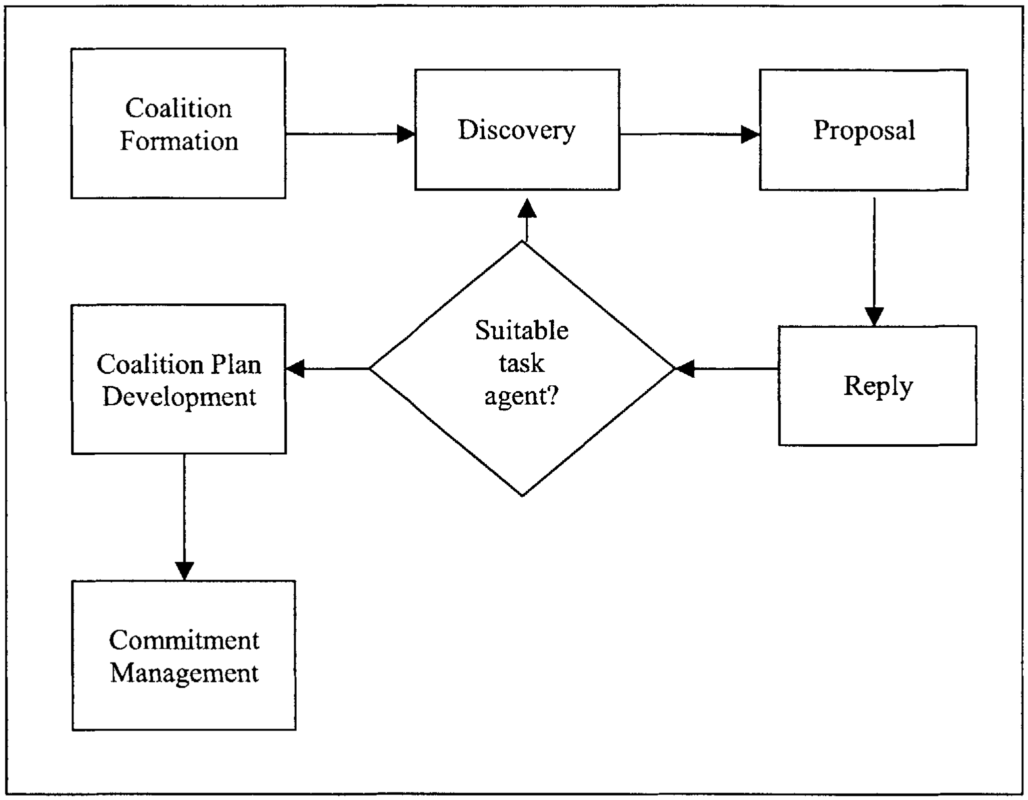

Figure 10: Flowchart for collaborative planning and execution by agents

8.3 Difference between Agents and Web Services

There are certain similarities between agent architectures and Web services. Like Web services, agent architectures offer both yellow-page and white-page directories which allow the agents to announce their services and for other agents to subscribe to them. However, agents broaden the scope and functionality of Web services.

A Web service is only acquainted and familiar with itself, and not the environment or its associated users, clients and customers. In contrast, agents not only know about themselves but are also aware of other agents that they happen to come into contact with. Through such interactions with other agents, an agent could possibly build up and develop its skill and competency. Lacking the ability to learn and be aware of its surroundings, the Web service can never become more intelligent or able to take advantage of new developments in the environment. It is also not possible for Web service to adapt its service to a client, such as providing enhanced service to repeat users.

The Web service, unlike agents, cannot resolve differences in ontologies adopted by the customer, service provider or any third party. The incompatible ontologies will impede any useful form of information exchange or data transfer. Moreover, Web services are non-autonomous and passive in nature whereas agents can be programmed to provide alerts and periodic updates when new information becomes available. Independent agents can jointly work together in groups to accomplish tasks and services that are beyond the capabilities of individual agents.

8.4 Security of Agents

The advent of the Semantic Web has led to the emergence of new security risks on the Web. Agents and ontologies are vulnerable to attacks because malicious users and their agents have the ability to divulge classified information or disrupt operations of Web services. The Web usage of normal agents can be monitored and analyzed, and have their communications subsequently tampered by other ill-intending agents.

Due to the openness of multi-agent systems, ordinary agents are susceptible to other destructive agents whose aim is to cause damage. For example, an agent responsible for the acquisition of office equipment from a regular supplier may be "killed" by another unscrupulous agent who might then pose as the original agent and begin accepting orders and payments. This constitutes a direct attack. An indirect attack might mean scrutinizing the activities of a competitor's marketing agents and steal information about a competitor's customer profiles and spending habits.

Hence, in order for the Semantic Web to thrive, its components must be safeguarded to provide a sense of security and assurance to users, service providers and program developers (Farkas, 2002). The way many developers are securing multi-agent systems involves coming up with the means to sanction and verify individual agents and platform components. In addition, the protection of communication protocols, agent privileges models and trust management systems are vital (Tan, 2002). The supporting technologies such as XML, RDF and DAML-OIL also need to be safeguarded.

8.5 Scalability

Scalability is the ratio between performance and resources. Agents need to be scalable if they are to exist in very large numbers, be discerning of their surroundings and network between themselves and their human-users. These agents will need not only scalable infrastructures and network support, but also scalable social services encompassing ethics and laws (Huhns, Jan 2002).

An agent-based system is considered scalable if the system performance increases when more resources are committed to it. This is possible because a scalable system is able to harness the extra capacity and computational power provided by the increased number of agents. In addition, the system should operate as per normal regardless of the level of resources (i.e. number of agents). Scalability can be classified as either static or dynamic. Static systems have to be re-synchronized and rebooted when the number of agents available change. On the other hand, dynamic systems can manage with such changes in resources while still running.

In practice, reactive agents are not concerned with scalability because they use the system resources only when activated with a new message. Increasing the number of reactive agents merely poses a storage complexity and slower communication rate of messages. However, proactive agents utilize resources constantly even when they are idle. There are a number of ways to achieve scalability:

- Using distributed computing technologies such as DCOM, .NET, and JINI to allocate system components across numerous physical machines.

- Program the components in an optimal way to perform the task while minimizing runtime.
- Organize the components in structured manner or into hierarchies to make them run more efficiently

Besides having scalability in our agent-based system, the infrastructure and the inter-agent services also need to be scalable. Agent services include name services, location services, directories, facilitators, and brokers. Infrastructure services include message transports, human interfaces, and CPU cycles (Brazier, 2001).

Looking ahead, the lifespan of future agents hinges on scalability. Currently, agents are not designed for long extended lifespan with most if not all of them dying after they have accomplished their immediate task. For instance, an agent might be programmed to obtain a price quote for a flight from New York to London from a list of travel agencies' web pages. After doing so and returning the information, the agent ceases to exist. However, future agents (especially those which represent items or customers) will have to live for a longer period of time.

Such agents will build up an information and knowledge base during its lifetime, as they come across new environment and situations. This makes it expensive and impracticable to just let them die. Likewise, the relevant infrastructure services must live for many years to offer the necessary support for such long-living and adaptable agents.

8.6 Other Application of Agents

Presently, applications of agent technology are still very much in the development stage. However, earlier agent architectures and models have already been implemented in certain diverse sectors ranging from healthcare, online merchandizing and logistics. Some of these pioneering works have enjoyed varying degrees of success, fueling the belief that further research into this field will yield even better results.

8.6.1 Healthcare Sector

Patients in a hospital setting can be denoted by agents, whose aim is to maximize the well-being of the patient they represent. The agents will issue electronic reminders or alerts should the patient fail to turn up for his/her medical appointment, or notify people of any new medical drugs pertaining to their condition should such a new treatment procedure become available (Davis, 1998).

Presently, systems handling medical information and patients' data-storage are detached and cater to only doctors and medical practitioners who are in the position to access them. There is not one system that directly assists the patients themselves, simply because of the fact that they are just too many of them. However, with agent technology, it is financially feasible to attach an agent to each and every individual patient throughout his/her lifetime.

The agent can track the patient's medical condition (blood pressure, heart rate, glucose level) and alert doctors should an abnormality arises. It can also readily produce records of the drugs the patient had consumed, the type of allergies that the patient may have as well as the details of each medical appointment attended. Such agent technology will greatly improve the efficiency of the healthcare sector, and ensure better reliability in the management of medical records.

8.6.2 Online merchandising

Agents can also represent customers at online merchandisers such as Amazon, Barnes & Noble, as well as at online auctions like eBay. The role of the agents is to ensure that the purchase and sale of products and services are done correctly and competently. For example, a buyer represented by an agent wishes to make an online purchase of a particular book at the lowest price possible. The agent will relay the specific information about that book (title, author, publisher and publication date) to a number of online merchandisers such as Amazon and Barnes & Noble, which will then process the agent's request.

Consequently, the agent's job will be to return the information on the price quotes to the buyer and let him decide if he wants to purchase the book at the stated price. If so, the agent will perform the business transaction and payment on behalf of the buyer who will subsequently be issued a transaction code for his purchase.

Apart from the relatively simple tasks of buying and selling a product, agents can also cater to the unique taste and preference of the customer they represent by interacting with other agents to determine which products are more popular. This has profound effects on consumerism because shoppers and retailers have current and up-to-date information on which products sell better than others. Vendors can also improve their management of inventory if they know where and when to re-stock their products.

Parallel to the agent technology employed by these online stores, Yenta[21], from MIT Media Lab, brings together agents with similar inclinations (Foner, 1997). Stanford's LIRA research system uses collaborative filtering to find the things liked by people who are similar to you (Balabanovic, 1995); while AT&T's ReferralWeb help find people with a particular knowledge and skill (Kautz, 1997).

8.6.3 Logistics

Another application of agent technology is illustrated by the logistics sector such as the postal service, whereby each individual agent represents either a parcel or mail. The agent can be used to trace the exact route taken during delivery of the package, its current location as well as to provide an estimated time/date of arrival at its destination. Such an application is extremely helpful considering the postal service typically handles millions of packages each month.

Similarly, agents representing logistical items in a military operation can be programmed to ensure safe delivery to friendly forces. Each item can be tagged with an agent-based smart card, which allows it to be tracked while navigating in a hostile environment. The

[21] Yenta – http://foner.www.media.mit.edu/people/foner/yenta-brief.html

smart card would have some form of logic and a data-base of routes, conveyances, and conflict-resolution strategies; and the agent's intention, purpose and goal. Therefore, an agent representing an ammo crate that was heading towards enemy-held territory might send out an alert signal to its army battalion logistics officer, warning him of the impending danger (Staats, 1996).

Chapter: 9 Multi-Agent System Solution for MRO

The aircraft MRO industry is both time-sensitive and cost-conscious, with great emphasis placed on efficiency and creating value for the customer. The main objective of the MRO process lies with executing the most suitable maintenance procedures in the shortest time possible. Therefore, the delivery of a cost-effective MRO will require a system that is capable of integrating the many entities within the MRO set-up. This will range from the maintenance depots, machines, spare-parts, information systems, right down to the personnel such as the mechanics, inspectors and engineers performing the actual maintenance. We thus propose a solution for MRO based on agent technology.

In this chapter, we propose a MAS solution for MRO. Sections 9.2.1 and 9.2.2 discuss the types of agents and data inputs required for our proposed agent architectures. The MAS architectures for aircraft inspection, repair and restocking aircraft parts are laid out in sections 9.2.3, 9.2.5 and 9.2.6 respectively.

9.1 Overview of MAS

An agent can be regarded as a unit that possesses domain knowledge and capable of using it to accomplish objective-orientated tasks and goals. Groups of agents can come together to form multi-agent systems (MAS) which are useful in providing the interface necessary in coordinating and managing the interactions between various parties or factions. The MAS is autonomous and there are no centralized moderators (Kalakota, 1995).

The basic idea behind MAS is to harness the capabilities of each individual agent in order to execute the overall system functionality. Each agent has to cooperate (possibly compromise) with one another and contribute its unique skills and competencies. MAS can help to mitigate the complexity in dealing with multiple (possibly conflicting) goals and facilitate information exchange between the different entities.

MAS also include protocols for inter-agent interaction and relationship, and the technology required to create new agents. There is no restriction to the number of agents a system can have, because agents are free to enter and exit at any given time. The multi-agent system also holds multi-modal information that is required by multiple users and interoperable business processes.

The numerous agents in MAS comprehend and make use of data contained in ontologies. A certain degree of autonomy is present in agents who can decide among themselves if they want to take up a particular task alone, assign it to others or collaborate with one another to complete it. Agents interact and communicate with one another to share information, determine the level of information required at each point as well as negotiate a compromise if their goals differ.

9.2 MAS Solution for MRO

In aircraft MRO, the airline flight operations division, maintenance depot and warehouse, and aircraft part suppliers are the three main business units that interact with one another. See Figure 11. Each of them can be conceptualized as a virtual organization and represented as individual cells.

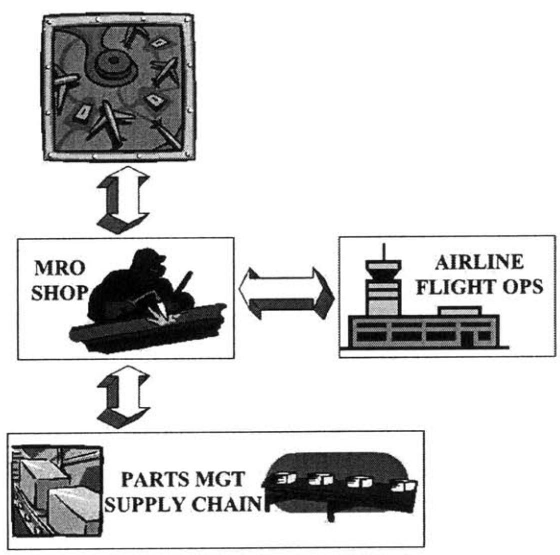

Figure 11: Three basic cell entities

The overall MAS solution for MRO that we propose can be broadly divided into three sub-types of agent architecture. There is a different architecture for each of the aircraft inspection and repair procedures, and another that handles the supply chain management of parts.

9.2.1 Types of agents

Within each sub-type in the proposed MAS solution, there will be a distinct set of agents. These agents can be autonomous or semi-autonomous and have varied roles and responsibilities according to the way they are classified. The proposed solution calls for five types of agents, namely information agents, task agents, negotiation agents, broker agents and resource agents. Collectively, they perform the duties of collaborative planning and execution as shown in Figure 10.

An information agent performs the role of managing, manipulating or collating information from many distributed sources. It provides intelligent access to a heterogeneous collection of information sources and has models of information resources and strategies for source selection and filtering, information access, and information fusion. They can receive requests for information and reply with the required information.

A task agent has pre-set objectives and purpose defined by its user or owner. It comes up with plans and strategies, and executes them in order to achieve a particular aim. The agent is aware of the task domain and the existence of other constructive information and task agents that are useful in helping it achieve its goal. It then decomposes the plan and co-operates with appropriate task and information agents for plan execution, monitoring and results composition (Sycara, 1999). A task agent can also be used to represent a customer through an order placed into the system. It defines the processes needed to complete the order and monitors the order status.

A negotiation agent interacts with two or more task agents to help resolve their differences and possibly reach a compromise between their conflicting goals. It obtains the objective specifications, shared abstract principles and preferences of the task agents and models them as a utility function so as to devise an optimized solution.

A broker agent helps customers find providers of services and products. It brings together information on the location, availability and capabilities of other agents, and possibly additional information about them. A customer registers his demand with the broker agent who then advertises and broadcast this information to prospective service or product providers.

A resource agent represents a physical resource within the cell. The physical attributes and capability of the resource (machine, tool, equipment, and worker) are reflected. This allows the status tracking of the resource, and also permits quantity and allocation of resources, their utilization rate, profit and loss figures of the entire cell to be reported. Table 6 summarizes the types of agents and their characteristics.

Cell agents				
Information	Task	Negotiation	Broker	Resource
Provides information about resources within each cell Responsible for source selection, information access and fusion	Co-operates with appropriate task and information agents for plan execution Represents a customer through a order placed into the system Defines the processes needed to complete the order and monitors its status	Resolve differences and reach a compromise between agents' conflicting goals Communication interface with the outside world on behalf of the cell	Helps customers find providers of services and products Provider registers with the broker Brokers are organized hierarchically based on geographical regions	Represent a physical resource within the cell (machine, person) Attributes and capability of the resource Status tracking of the resource Profit and loss figures

Table 6: Types of agents

9.2.2 Necessary Inputs

There will be four main sources of information within the entire MAS – aircraft, depot, flight schedule and inventory. As shown later in Figure 12, agents rely heavily on information stored in these sources in order to make decision choices. For example, the combination of these information sources will allow agents to find out the condition of the aircraft and arrange appropriate maintenance or repairs during the aircraft's non-flight hours, based on the depot resources and availability of parts.

The information contained in these sources is dynamically changing and subject to human inputs and intervention. In this way, humans can affect the running of the whole MAS by altering the data sources or setting constraints on their inter-relationships and interactions.

The aircraft information source will contain data inputs on the exact specifications or blueprints of the aircraft. Information such as the aircraft's manufacturer, model type, year of production, wing-span, height, fuselage diameter, engine identification number will all be furnished in great detail. Each aircraft will have its own unique repair history and maintenance schedule, listing the type of faults detected and relevant repair performed during the aircraft's lifecycle.

The depot information source will hold information pertaining to the maintenance staff and resources at the MRO depot. Size of the maintenance staff and its specialization, job specifications and availability of each worker, personnel roster will be reflected. Likewise, the same level of information on the machines, tools and equipment necessary for performing repairs and maintenance will be presented. The description of inspection or maintenance works to be carried out in response to a particular observation, event or failure will be made accessible together with the resource (tools and personnel) requirements for each task and the way to sequence the tasks.

The flight schedule information source includes details on the airline's fleet of aircraft, its size, type, usage hours and most importantly the airline flight roster. The roster states the routes taken by aircraft when flying between cities operated by the carrier, the type of aircraft operated in those various routes, the flight timings and frequencies of each route, duration or flight time of each round trip.

The inventory information source identifies each aircraft part by its serial number/RDIF tag. The serial numbers and RDIF tags are often sorted by which sub-system the part belongs to, which component it is fitted with, as well as the type of part itself. Therefore, a bearing on a compressor shaft in an aircraft shaft might have a different serial number corresponding to the same bearing on the airframe. Information on the quantity of each aircraft replacement part currently available in the inventory, its shelf-life (if any), and delivery dates of new stock are all stored in the inventory information source.

9.2.3 MAS Architecture for Aircraft Inspection

Figure 12 below shows the proposed MAS architecture of a typical inspection process for an aircraft on arrival at a maintenance depot. There are a number of sequential stages leading from the onset of inspection to the final state whereby a task agent eventually decides that a repair job is inevitable. It consequently employs the assistance of a broker agent who is then responsible to advertise for such a repair service.

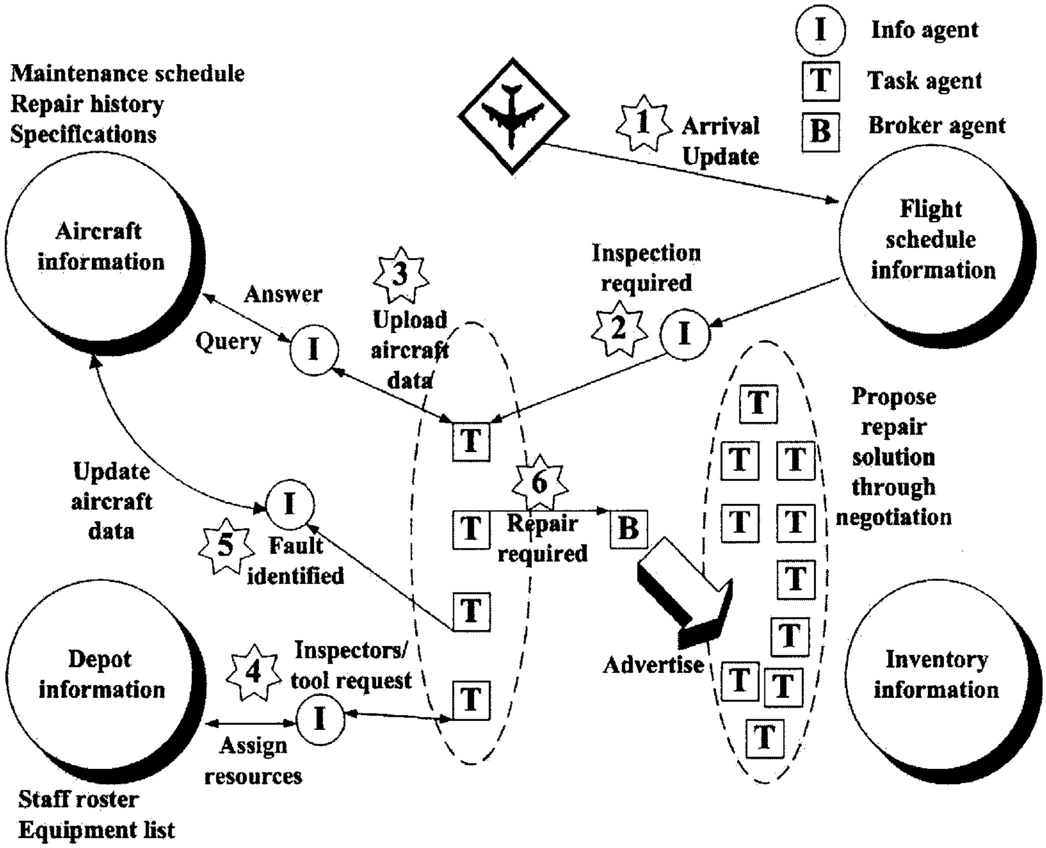

Figure 12: MAS for inspection process

The following steps occur during the inspection process –

1) The flight schedule information source is notified of the aircraft's arrival at a particular MRO depot, and this can be checked for consistency alongside the flight roster.
2) An inspection request is communicated to a task agent via an information agent.
3) One of the task agents among the group of agents nominated to carry out the inspection sends a query to the aircraft information source to find out the current state of the aircraft and any relevant maintenance history. It then

receives an answer and uploads this information to be made available to the rest of the group.

4) Based on the level of inspections to be done, another task agent makes a request to the depot information to seek the necessary personnel and tools. The resources will be subsequently assigned according to how best to optimize utilization of resources.
5) After inspecting the aircraft, any fault detected will be updated to the aircraft information source for follow-up action.
6) A repair is sought to remedy the fault identified. A task agent will convey the repair request to a broker agent who will then be responsible to advertise this service demand to a large group of task agents, stating the details of the fault and the repair job on hand.

9.2.4 Formation of Repair Team

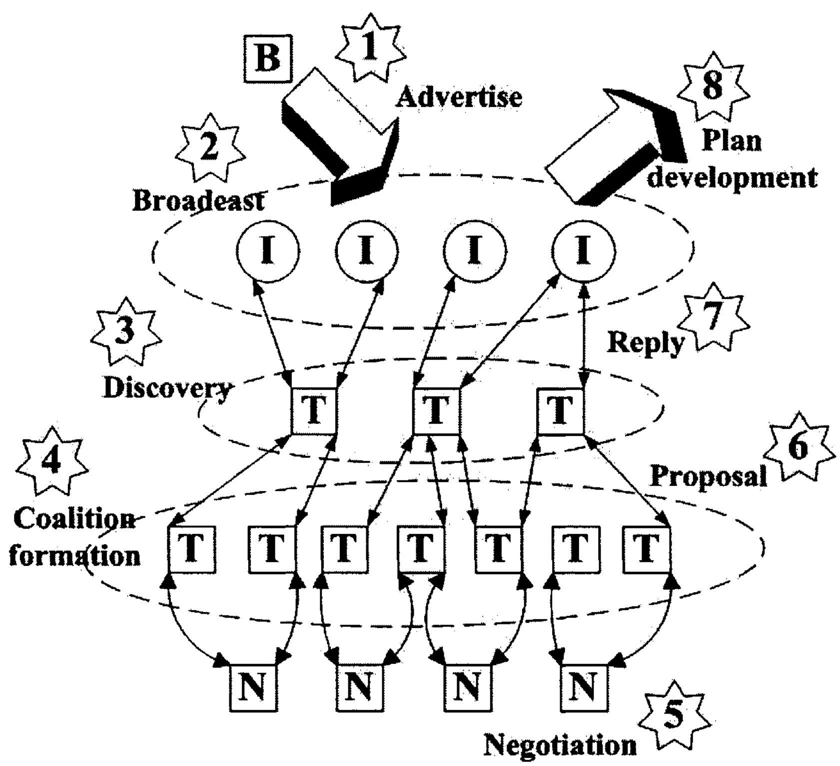

Figure 13: Negotiation process in team formation

The following steps occur during the formation of a repair team (refer to Figure 13) –

1) A broker agent advertises a repair request resulting from the inspection process.
2) The specific details concerning the repair job (i.e. part to be replaced, estimated duration, resources required) are broadcast to a group of information agents.
3) During the discovery phase, there will be a group of task agents evaluating the key features of the repair job. They will decompose the problem into smaller tasks and try to find or locate other suitable agents capable of performing these tasks.
4) A broad solution emerges after the formation of a coalition of agents, whereby each task will be allocated to an individual agent. A favorable assignment of tasks means the most appropriate agent is selected for each task.
5) Negotiation takes place to ensure that there are no conflicts of interest among the agents and their respective goals are met.
6) The final solution is presented as a proposal.
7) Proposal is conveyed to the information agents in a reply.
8) Solution is ready for implementation.

9.2.5 MAS Architecture for Aircraft Repair

Figure 14 below shows the proposed MAS architecture for a typical aircraft repair, after the team of agents has been selected to perform the repair. Every phase of the repair is carefully monitored and the fault, repair solution and corresponding outcome are updated in the maintenance record (stored in the aircraft information source) of each aircraft. The flight roster is also amended when the aircraft is restored and returned to normal duties.

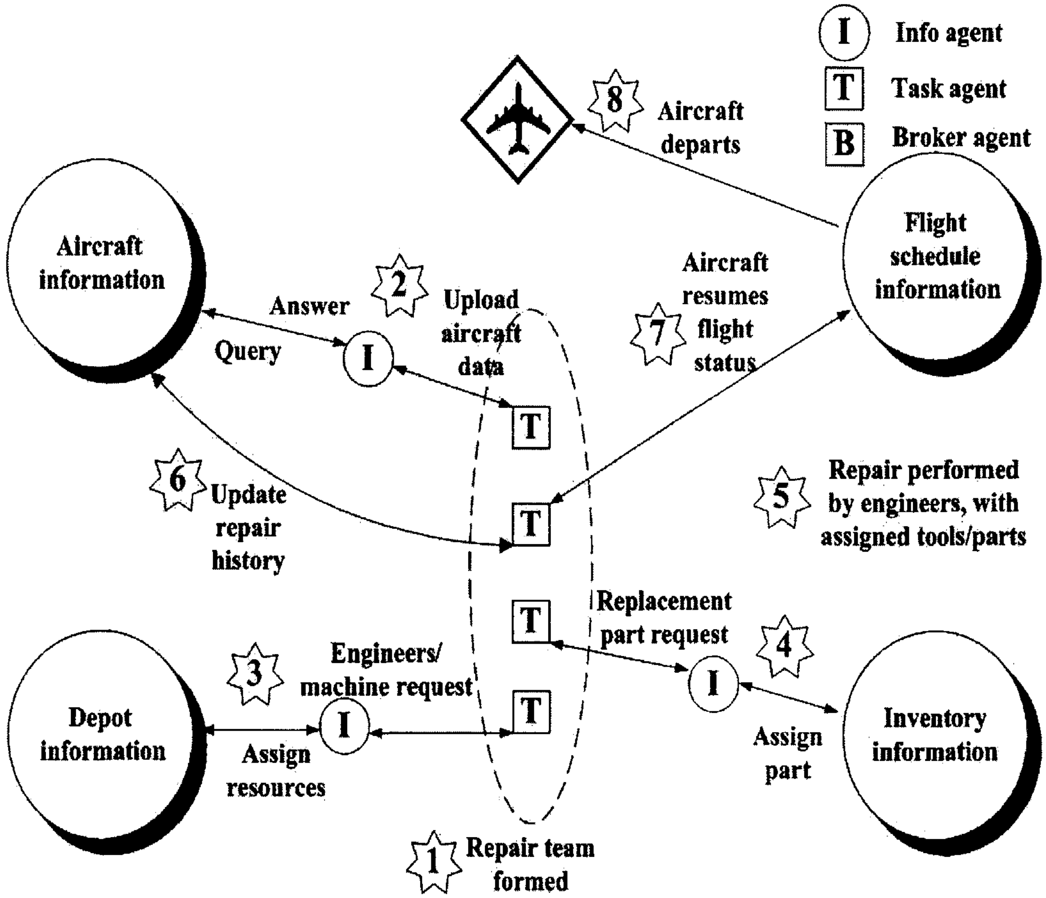

Figure 14: MAS for repair process

The following steps occur during the repair process –

1) The repair team is formed through a negotiation process, as described in section 9.2.4.
2) One of the task agents uploads the aircraft information and its maintenance history, and makes it available to the rest of the group.
3) A request is made for the engineers and machines required to perform the repair. These will be assigned depending on their availability.
4) Likewise, a separate request is made for the aircraft part (if any) to be replaced. Subject to records in the inventory information source, the part

will be issued if it is currently in stock. Otherwise, a purchase order will be made to the supplier to restock the item in question.

5) Maintenance engineers execute the repairs using the assigned machines, tools and replacement parts.
6) The maintenance history of the aircraft, contained in the aircraft information source, is updated after repairs are done.
7) The aircraft resumes flight operations, and this new event is reflected in the flight schedule information source.
8) Aircraft leaves MRO depot.

9.2.6 MAS Architecture for restocking parts

The essential parts required for aircraft maintenance might not always be readily available when they are called for. This is influenced by both the storage capacity within the MRO facility as well as the cost and hence size of inventory. Some parts can be very expensive or have a short storage life, therefore they are only acquired upon a requisition request.

Figure 15 below shows the proposed MAS architecture for restocking aircraft parts when stocks run low. It involves interactions between the inventory information source, MRO warehouse and the suppliers in the supply chain. The MRO supply chain management will be further expounded in section 9.3.

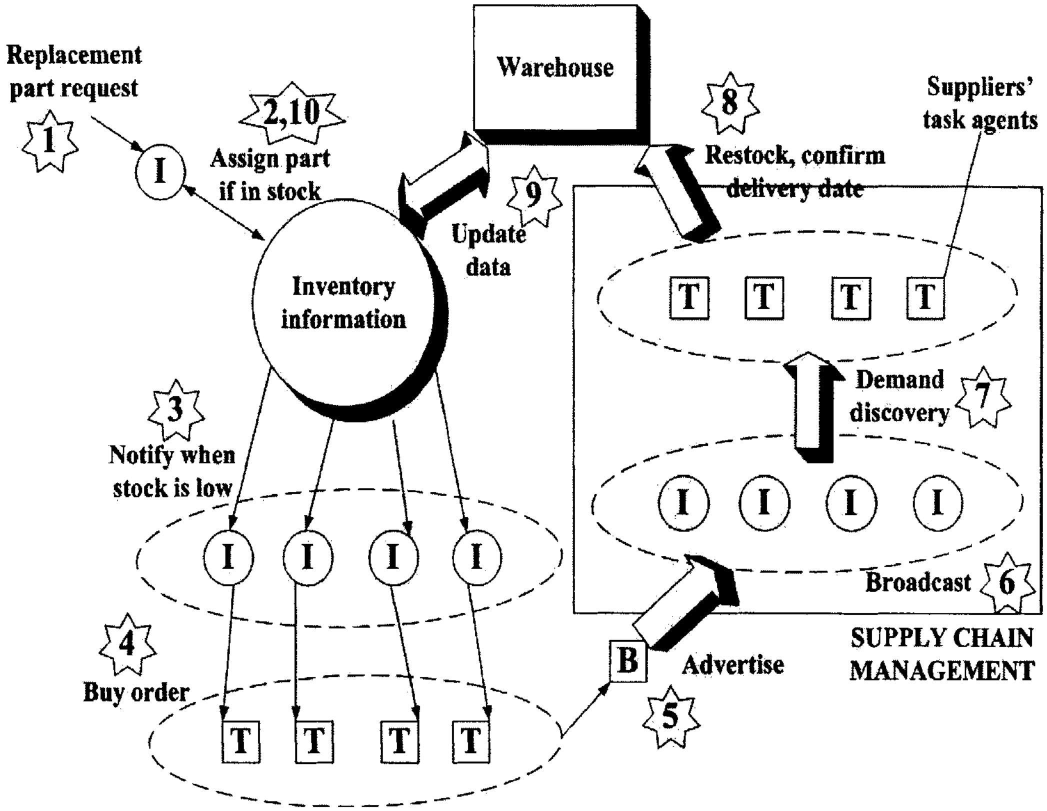

Figure 15: MAS for restocking parts

The following steps occur during the parts restocking process –

1) A part request arrives, as a result of the part being a requisite for a repair job.
2) The part is assigned if it is currently available in the inventory. Otherwise, the repair job has to wait until the part has been purchased.
3) Each part can be coupled to an information agent, which is in turn tied to a particular task agent. In the event that the quantity of a part is depleted (level can be set by human user), the information agent assigned to the part will be alerted.
4) The associated task agent will issue a buy order to replenish the stock.

5) A broker agent will have to communicate the specifications of the part, offer price and terms of delivery to prospective suppliers in the supply chain.
6) Information agents will broadcast the order request to a separate MAS (see Figure 16) contained in the supply chain management.
7) Upon discovery of a demand for the part, the suppliers' task agents will submit their respective proposals and attempt to win the contract.
8) The successful bidder will arrange a delivery schedule and restock the part.
9) The inventory information source will be updated once the stock for the part arrives.

9.3 MAS for MRO Supply Chain

To exploit MAS as a means to manage the supply chain, each entity in the supply chain can be represented as separate cells. The need for agents in each cell to compromise with each other arises when there is a misalignment of their respective goals. This is a likely situation in any typical supply chain whereby companies are driven by profits and self-interest. Even though the different organizations sought to work with one another, there will be instances when disagreements over certain issues will occur.

In addition, the entities of the supply chain can change over time due to free entry and exit from the industry. As such, there are no stipulated links or associations between the agents in such a dynamic MAS. At any one time, they can choose to participate in a particular activity in the system, stay but remain dormant or leave the system completely.

Agents use ACL to interact with one another and exchange information on suppliers, goods, logistical arrangements, business logic and rules contained in the ontologies. It operates based on negotiations conducted between agents. Task agents obtain information about prospective negotiation partners from information agents. The negotiation process is characterized by allocating quantities and weights to a set of parameters. Negotiations between agents representing the manufacturers, vendors, logistics providers, distributors, wholesalers and customers will eventually set up a virtual supply chain.

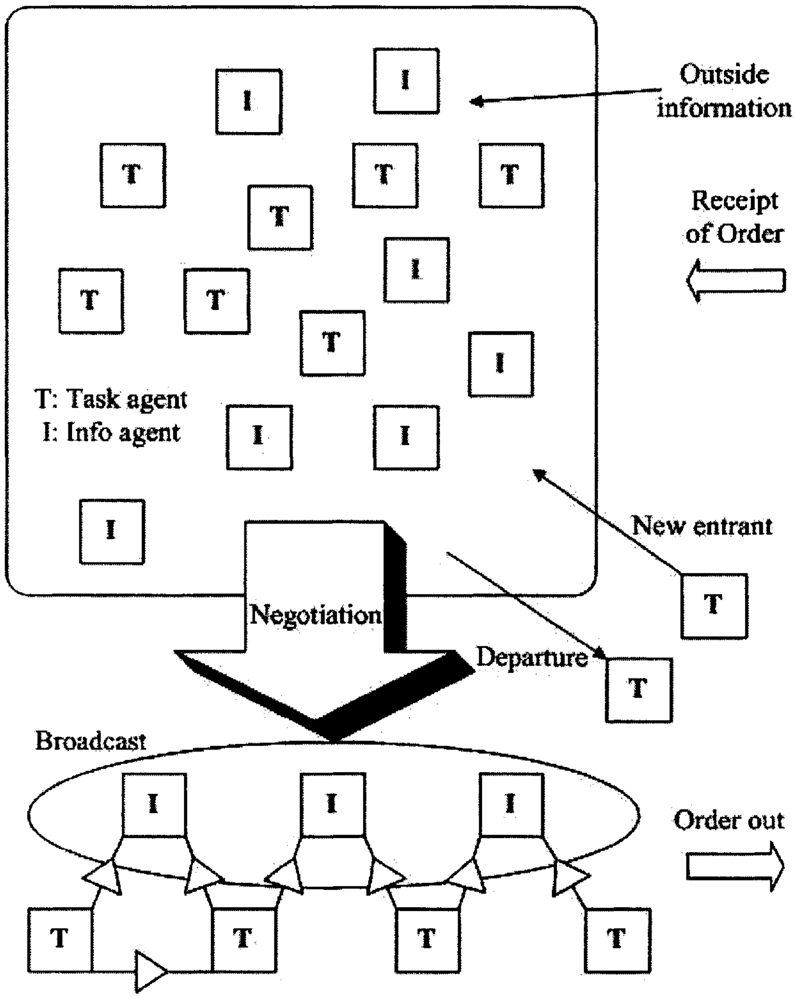

Figure 16: MAS architecture for MRO supply chain

In Figure 16, we find that the MAS for MRO supply chain has an open structure whereby existing agents can leave the system while new ones can join according to orders or decisions made by their respective owners. For example, a new supplier for engine bearings wants to start offering its product to the market. To join the MAS, its task agent will have to advertise information about the new product to one of the information agents who will subsequently broadcast it across the entire MAS community. Conversely, a supplier who wishes to exit the market has to notify an information agent about its departure.

When a new part order arrives, the suppliers' task agents will need to evaluate the requirements of the part and reconcile it with their capability to manufacture it according to the exact specifications. Other considerations will include the purchase price offered by the MRO shop, terms of delivery and logistics issues. After which, they will submit their respective proposals and attempt to win the contract.

A MRO shop that wishes to look for a suitable part supplier to provide engine bearings will then assess the proposals put forth and set conditions such as the availability of the part, selling price, quality and tolerances, distance between supplier and MRO facility, as key factors for consideration. The agent representing the MRO shop would then have to weigh different combinations of parameters offered by supplier agents, in order to select the supplier that can best meet its needs.

A lengthy process of negotiation will take place between MRO and suppliers' agents before the MRO shop eventually makes its choice on the most compatible contractor it deems to be able to meet its demand for engine bearings. The final order leaves the MAS with the successful bidder having to arrange a delivery schedule and restock the part according to the terms agreed upon.

Chapter: 10 Implications of Agent Technology in MRO

A mechanism for maintenance scheduling and supply chain management of aircraft parts relating to the commercial airline industry is complicated because it has to be quick to respond to changes in passenger loads, fleet status, flight routes, crew levels and weather conditions. It usually involves a dynamic and iterative process of re-evaluation of both short and long-term plans, re-prioritization and re-allocation of scarce resources.

Here we discuss the implications of our proposed multi-agent system for MRO. The consequences can be broadly classified into two categories – aircraft maintenance (Section 10.1) and supply chain management of aircraft parts (Section 10.2).

10.1 Aircraft Maintenance

The MAS proposed in Chapter: 9 for aircraft repairs and inspections may provide a means to perform real-time maintenance scheduling and decision-making in a more productive and adaptive manner. The benefits of such a MAS include congregating information under one common system, integrated decision-making, improved deployment of resources, more efficient and cost-effective aircraft MRO. If done correctly, it is possible to increase resource utilization (aircraft, replacement parts and maintenance staff), increase throughput of repairs and reduce the number of late jobs. It also creates a flexible and dynamic architecture that responds to continuously changing market conditions.

10.1.1 Amalgamation of Information Systems

Presently, we find that some airlines operate a number of detached but similar information systems for access, storage and operation of different sets of data. Although the systems run parallel to one another, they can be incompatible when it comes to sharing, swapping and exchanging data between two or more platforms. They practically have zero interoperability and common information cannot be reused without first duplicating it from one system to the next.

For instance, maintenance personnel at a MRO depot wishing to carry out a routine repair might have to retrieve information about the aircraft, its maintenance history, repair procedures and parts availability from a number of data sources. This can be tedious and time consuming, leading to delays in job completion and increased downtime for the aircraft.

In the proposed MAS for aircraft inspection and repairs detailed in Figure 12, these information are kept in compatible information pods and accessible by information agents via an answer and query mechanism. This information can be retrieved from within any part of the MAS almost instantly and this helps to cut down on repair time and improve overall MRO efficiency.

The electronic retrieval, fusion and transmission of information are highly automated and are less prone to human errors and other discrepancies. Based on information from aircraft maintenance schedules, it is possible for an agent to provide an alert if a maintenance activity is not completed by a particular date as originally planned.

Each aircraft has its own separate maintenance plan, repair history and information on previous downtime, lubrication requirements and equipment records. The systematic record of aircraft faults and repairs can be used for future reference and similar recurring problem can be treated the same way for other aircraft. Such data can also be used for failure prediction and condition monitoring.

10.1.2 Integrated decision-making

Human decision-making tends to be a protracted and time-consuming process especially if there are many conflicting factors for consideration. Even when a decision is eventually reached, it may not be the most appropriate to others because of the fact that a full picture or appreciation of the problem is not realistic because of imperfect knowledge. There is a certain element of human bias and personal judgment involved when humans make decisions.

On the other hand, if we can accurately model the value (or utility) of various scenarios and situations, then MAS offers an alternative that allows decisions to be concluded collectively and arrive at the most optimum solution much faster for the given set of factors and constraints. This is made possible through the process of decomposing the larger problem to smaller tasks and finding the most capable task agents to perform them, at the same time balancing the opportunity cost of not having these agents available to carry out other tasks.

Each task agent has a set of objectives, preferences and utility functions. The idea behind integrated decision-making is to maximize the joint utility of the agents while solving the overall problem. Similar to information retrieval in the MAS, this negotiation process for task allocation can also be automated to avoid unnecessary time wastage and achieve greater efficiency. The negotiation agents will evaluate the pros and cons for the different coalition of task agents and choose the grouping that best satisfies the job requirement as well as having the highest joint utility. They must also ensure that there are no conflicts of interest among the task agents and their respective goals are met. The average aircraft inspection and repair time will decrease as a result of the automated allocation of tasks within the MRO shop.

10.1.3 Better Management of Resources

As a result of the union of all available information sources coupled with the capability for integrated decision-making, it is possible for the MAS to 'look into the future', predict the issues that may occur and contrive a feasible course of action to deal with them. This can be particularly useful when it comes to real-time scheduling of maintenance activities in a MRO shop.

For example, an aircraft might be called into a MRO service bay for unscheduled maintenance on a faulty landing gear. However, when enquiring about the existing status of the aircraft the task agents realize that the plane has a scheduled 'D' check (i.e. engine overhaul) coming up soon. Therefore, the task agents can be programmed to be intelligent enough to propose that the 'D' check be brought forward to an earlier date, in order to coincide it with the repair on the aircraft door. By doing so, it will minimize overall aircraft downtime and prevent the same aircraft from being called into the service bay twice within a short period of time.

Using data on the number of flight hours of each aircraft, it is possible for relevant usage-based and calendar-based maintenance to be scheduled accordingly. In addition, it can be beneficial to schedule complementary maintenance activities to be carried out on the same aircraft simultaneously and task agents can provide alerts when such a situation arises.

If multi-repair tasks can be performed during a scheduled maintenance stop, this will reduce the unnecessary duplication of routine jobs such as initial aircraft inspections that is associated with each visit to the MRO hangar. This eventually leads to better management of resources and ultimately cut aircraft downtime.

The creation of virtual teams of agents permits decentralized work coordination without the need for supervisory control and management. In other words, there is no additional requirement for any follow-up action or administrative chore once a work team is formed after the negotiation process. The work teams are left on their own to fulfill their respective tasks and places no further strain on resources. This is beneficial for resource management because of the predictability it offers.

10.1.4 Responsive and cost-effective scheduling mechanism

The automated MAS solution for airline MRO is highly responsive because information on any changes relating to passenger loads, fleet status, flight routes, crew levels and weather conditions can be directly fed into the data stream and made accessible through the information pods. These changes are reflected almost immediately when they occur.

For example, the pilot of an aircraft experiencing an in-flight malfunction of one of its engines will notify the ground crew of the problem. This information will be relayed and updated in the aircraft information pod and a virtual team of task agents will instantaneously be formed to tackle the problem of recovering the operational status of the failed engine.

The process of coalition formation of agents, tasks allocation, provision of maintenance resources and aircraft replacement parts can be concluded even before the aircraft reaches its next MRO depot. This means that it is possible for the actual engine replacement as well as the engineers and equipment required to perform the overhaul are available as soon as the aircraft arrives. Once the repair job is completed, the virtual team disbands and those task agents in the team are once again free to take on other assignments.

Besides being responsive, the MAS solution can also be cost-effective when implemented in the long run. The automated system means that the necessity for any human interference is removed once the objective of each task agent is carefully defined, together with their value propositions and mode of negotiations with fellow agents.

The abolition of the traditional practice for MRO foreman or supervisor to assign jobs among maintenance personnel eradicates the opportunity for favoritism, discriminatory work practices and standardizes the man-hours for each job across different aspects of maintenance. The task of job assignment is taken over by the task agents in the MAS solution and this does away with the layer of middle management and helps to lower the wage bill.

Task agents can perform flexible work scheduling to cater to the staff preference for working hours, length of works shifts, public holidays and weekends. Although the weekly hours of MRO personnel are comparable across the board, however they can choose to work different shifts or days depending on their own inclination.

The flexibility in work scheduling depends largely on the abilities of the task agents to dynamically ensure the MRO shop is operational at all times to meet the human resource requirements (i.e. mechanics and engineers) for inspection and repair jobs while trying to accommodate the special requests of each individual at the same time.

10.2 Supply Chain Management of Aircraft Parts

Other than streamlining the aircraft inspection and repair process, agent technology can also be applied to the supply chain management of aircraft parts. We first provide the definition of supply chain management as the examination of the entities and logistical processes within the supply chain with the aim to cut costs, improve quality and delivery timings.

It can be further described as "a collaborative-based strategy to link inter-organizational business operations to achieve a shared market opportunity" (Donald, 1999) and is usually divided into three areas - sourcing, manufacturing and delivery. As proposed in Chapter: 9, our agent architecture for MRO supply chain takes care of the inventory control, forecasting, sourcing and delivery of aircraft replacement parts for MRO.

10.2.1 Inventory control system

To evaluate the implications of agent technology in airline MRO supply chain management, we will need to look at a set of performance metrics involving responsiveness, flexibility and cost that is typically used to benchmark the effectiveness of any inventory control system.

10.2.1.1 Responsiveness

Since we measure responsiveness in terms of the performance in providing parts as and when needed, therefore the responsiveness of an inventory system for aircraft parts can be quantified by its issue effectiveness and the average waiting time for a part request. The issue effectiveness is the fraction of part requests that can be processed immediately and do not encounter any availability problem due to a shortage in inventory. The average waiting time for a part request is the period between receiving a request for a part to its actual delivery.

The MAS solution proposed in Figure 15 is more responsive than conventional methods of managing inventory due to better forecasting for the demand of parts and faster restocking procedures. The depletion of parts can be more precisely estimated with the aid of an extensive collection of information sources ranging from aircraft specifications, maintenance history, flight schedule, fleet size and age.

All these information are obtainable through information agents spanning the entire multi-agent architecture that would otherwise not be available to an inventory supervisor with only a narrow view of things that happen in the warehouse he manages. He would be unable to make accurate projections of part demand if he is unaware of changes in flight frequencies, routing or additions to the airline fleet. In contrast, task agents handling forecasting of parts have a better understanding and broader picture of the dynamics in the airline industry as well as factors that might affect the MRO shop itself. They are hence in a stronger position to foretell any shortage (or surplus) and decide when to restock (or hold off purchasing) parts, based on the additional information they possess.

As detailed in the MAS solution for restocking parts (see Chapter: 9), the restocking procedures are straightforward and can be implemented autonomously and expeditiously. Each part is assigned an information and task agent who will oversee the restocking of the part they represent when stocks of it run low. A broker agent will subsequently advertise the demand and a suitable supplier will be found in turn.

Besides the customary responsibility of placing purchase orders for parts, task agents can also fulfill two other roles. Firstly, they can be employed to monitor items with short shelf life and sound an alarm when the expiry date of the items is approaching. Secondly, when delivery of parts are not received from the manufacturer or supplier by the delivery date, they can either issue an alert notice to the warehouse manager or liaise directly with the task agents of the supplier to clarify the delay and possibly send a reminder.

10.2.1.2 Flexibility

The demand for aircraft spare parts is motivated by part failure and can fluctuate during a business cycle or due to changing economic conditions. A flexible inventory system should be able to handle such non-linear behavior of demand and adjust accordingly. Flexibility of the system is characterized by three factors:

- Length of time average demand can be met with current levels of inventory
- Proportion of parts in inventory running a surplus at the end of each month
- Time needed to source a part that is unavailable

According to the first two factors mentioned above, the goals of achieving greater inventory flexibility and keeping costs down are often conflicting in nature because the former would necessarily entail having a larger inventory and thus incurring higher storage and inventory costs. We will discuss more on cost in the next section (10.2.1.3). For now, we will explain how agent technology can give rise to an increased pool of suppliers, otherwise known as multiple sourcing, that reduces the time needed to source a part that is unavailable, increases part availability and improves flexibility.

Agent technology allows one to reach out across the Semantic Web and invite all suppliers to tender a bid, thus adding to the number of possible suppliers. The fact that there can be a greater pool of suppliers will make suppliers curb any monopolistic and anti-trust practices. It is also beneficial to have the flexibility in switching between suppliers and the lack of any long-term contracts allows constant reviews of the acquisition process to ensure that it conforms to current best practices.

Another implication of agent technology in supply chain management is the way in which other MAS solutions can be employed to adopt a strategic sourcing policy. Strategic sourcing differs from conventional sourcing methods used in the acquisition of parts for MRO operations. Traditional means typically involve a tender-bid process for all distinct parts whereby the supplier with the lowest price usually wins.

However, under strategic sourcing parts that are closely related in their features, have common sources of supply and procurement schemes, are grouped together as one common commodity. For example, all electric cables and wires used in aircraft cabin are considered to be one commodity because they all serve a similar purpose and are bought from suppliers belonging to the same industrial sector.

With strategic sourcing, the selected supplier(s) is chosen based on its ability to meet the overall supply needs of the entire commodity rather than just the individual constituent parts (Favre, 1998). This means the supplier can offer the best value on the whole in terms of fulfilling the orders for a particular commodity, despite not always being the lowest cost provider for all the parts. There is greater central planning because cost-benefit analysis and selection process are done across the entire MRO set-up.

Agent technology can be employed as a strategic sourcing tool once the parts in a particular commodity are defined. Based on some selection criteria, potential groupings of suppliers can be weighed in terms of their pros and cons to determine which set of suppliers can best meet the business objectives of the MRO operation.

10.2.1.3 Cost of MRO activities

This is perhaps the most explicit measure of the success of agent technology in the airline MRO industry. Airlines strive to minimize the cost of maintaining and servicing their aircraft so as to improve profit margins and eventually their balance sheets. They would want to limit the cost of their MRO activities while sustaining an operational and airworthy fleet of aircraft.

There are a number of constituents affecting the overall cost of MRO activities, and not all of them can be improved by employing agent technology. Largest among the cost components is the money paid out to purchase aircraft replacement parts and other technical support provided by the manufacturers or suppliers.

Apart from the actual amount of money paid to the suppliers, there is also significant cost from the transaction itself. This administrative cost results from marketing, research, offering and awarding the contract, maintaining the contract, monitoring suppliers' performance and the logistics required to pick, pack and ship the parts (Bickel, 2003). There is also the cost of human resources (engineers, foreman, inspectors, workers) to manage the MRO set-up. The wage bill of these personnel can be quite substantial.

We will however focus our attention on the cost of maintaining parts in the inventory system because this is where agent technology can make a large impact. The inventory holding costs include opportunity cost of capital to purchase the parts, storage and warehousing fees, depreciation of parts if they become damaged or obsolete while in storage. Damage could be due to moving the part within the storage facility or higher rates of failure that comes with age, especially for parts with electronic components.

Agent technology can help to lower the cost of individual parts through multiple sourcing or reduce overall cost by having a leaner inventory system and embracing just-in-time (JIT) principles. The next section (10.2.2) discusses JIT in greater detail. Due to the lack of perfect knowledge and information in a marketplace, the total number of potential suppliers for a given part far outstrips those actually bidding to offer its product and services.

Classical economic theory suggests that a lack of competition will drive up prices and bring down quality. This implies that the chosen supplier might not be the one offering the best deal. However, the price of the part will be least expensive if we have the maximum number of suppliers to compete against each other and allow market forces to operate freely. Agents can also reap possible economies of scale through bulk purchasing arrangements with suppliers.

Although the advantages of having a bigger pool of suppliers are evident, nevertheless there are also certain drawbacks. For the overhaul of a given aircraft, it is likely that there will be more parts suppliers as a result of multiple sourcing. More suppliers obviously mean more coordination work to ensure timely delivery and make it more challenging for the MRO shop to review and provide feedback on the quality, tolerances and specifications of parts.

While each individual part is of the lowest price and highest quality available in the market, the increase in the number of delivery orders processed might contribute to the total cost of overhauling the aircraft. There will be higher contract administration costs in order to monitor the numerous contracts for the purchase of parts. There is also a higher chance that various parts from different suppliers do not mesh well with one another, leading to longer repair time and possibly higher failure rates in future.

10.2.2 Just In Time (JIT)

Agent-based technology can be adopted as a means to implement some of the JIT strategies, so as to cope with the problems and issues surrounding supply chain management and manufacturing. The term, JIT, is coined to represent a process that is competent enough to react immediately to a change in demand thus making it pointless for the build-up of inventory that may otherwise arise due to the anticipation of demand changes or inefficiencies in the process (Hutchins, 1999). The lack or minimization of stock will help keep inventory costs low.

The JIT philosophy comprises of an entire spectrum of issues that involve radical changes in organizational planning, management thinking, workforce standards, industrial cultural values and factory layouts (Majima, 1992). The benefits of JIT are not restricted to a single organization but apply throughout the whole supply chain.

Admittedly, agent software is powerless to implement all of these principles and is therefore not sufficient to achieve the full potential of JIT on its own. However, it does enable the computerization of the stock control system, ensure the smooth flow of materials along the production line and efficient distribution of the end products.

To achieve this, agents can be made accountable to provide alerts to the factory or warehouse manager before stocks of a particular raw material or spare part runs low. As was mentioned earlier in our proposed agent architecture for restocking aircraft parts, more sophisticated agents can even be assigned to be directly responsible for restocking parts at regular time intervals or when a stipulated event occurs.

To counter the unpredictability of deliveries by suppliers, agents can employ multiple sourcing methods on the Semantic Web in search of quality-controlled inputs for their factory or warehouse at the lowest price while tracking the conveyance in collaboration with the suppliers' agents. Likewise, agents can also be utilized to manage the distribution of finished goods to customers. They can interact with customers' agents to determine where and when the goods are needed, and in what quantities and specifications.

Essentially, software agents can aid in production planning by ensuring that adequate resources that are necessary for production are in the right place at the right time. It can also help achieve a good channel and network for distribution of the finished products. This will eliminate inventory, minimize work-in-progress and working capital, and ultimately get us closer to achieving the JIT beliefs.

Chapter: 11 Conclusions and Recommendations

11.1 Summary of Findings

There are many inherent features of the airline MRO industry that are inhibiting greater productivity and competency. Poor management of aircraft spare parts, human-related maintenance errors and the lack of coordinated decision-making in maintenance scheduling are just some of the traits inhibiting greater productivity.

To streamline some of the MRO activities so that they are harmonized and complement the future air transportation system envisioned by the NAS, it has been proposed that a MRO agent-architecture and its relevant ontologies be developed. This will integrate some of the airlines functionalities and their corresponding information systems, and possibly offer a solution to the inefficiencies currently persisting in the MRO industry.

In this thesis, we proposed multi-agent architectures for aircraft inspection and repairs, restocking aircraft spare parts as well as managing the overall MRO supply chain. It is hoped that these solutions will lead to faster decision-making as a result of more rapid information access, gathering, filtering and fusion. This should eventually translate to more responsive aircraft maintenance and lower the cost of MRO as a proportion of airlines' total operating costs.

The two main entities in MAS are agents and ontologies. An agent is a computational entity that performs a stipulated task on behalf of the party it represents. It behaves autonomously because its objective can be predetermined. Equally important are ontologies that are organized, computer-coded hierarchy of concepts, ideas and controlled vocabularies with related meaning covering a single particular domain. They form the backbone and building blocks of Web services that constitute the Semantic Web.

Agents exploit embedded information in ontologies to make sense of our world and understand the underlying logic, rationale and inter-relationships. The possibility of employing numerous agents to work together on one single task allows the formation of MAS. The prospect of greater processing speed, system reliability and the capability to handle multi-dimensional optimization problems formed the rationale behind pursuing MAS solutions.

Another advantage of MAS stems from the openness of the system that allows agents to leave or enter at any time and making the system dynamic in nature. This means that new agents with enhanced capabilities can be introduced to the system as and when they become available without diminishing the overall system performance. These agents can be produced and initiated without altering the framework of the existing system.

The outcome of applying agent technology to airline MRO can only be positive if it provides a mechanism that manages maintenance scheduling and supply chain management of aircraft parts in a responsive and cost-effective way. Depending on the MRO agents' relationships and linkages (defined in the relevant ontologies) with the airline's flight operations, the agents can potentially put forward an optimized maintenance schedule based on the airline's passenger loads, fleet status, flight routes, staff levels and aircraft spare parts inventory.

11.2 Future Research

We have only just embarked on the journey of exploring how agent-based technology can be applied to the airline industry. We have so far restricted our agent applications to the airline MRO industry by proposing multi-agent architectures for aircraft inspection and repairs, restocking aircraft spare parts as well as managing the overall MRO supply chain.

The next logical step for future research would be to propose an agent-architecture to handle the complexity in airline flight scheduling. This is somewhat similar to maintenance scheduling in MRO except for different parameters and constraints. It is foreseeable that the MAS for flight scheduling will be capable of forecasting the number of flight-capable hours in each aircraft, based on its specifications, age and frequency of use. It is then possible to improve overall fleet utilization by allocating flight routes in such a way that it balances the flight-capable hours with the actual flight-hours flown of those operationally available aircraft.

In fact, the airline can schedule its flights in order to meet MRO objectives and vice versa. This requires a higher level of sophistication because it involves the coupling of two MAS for flight scheduling and maintenance scheduling.

Other areas for future research include the pilot implementation and validation of agent enhanced MRO IT infrastructure. This signifies a shift from the proposal phase to the actual development of the agents and ontologies themselves. It would require the use of translators, bridges and gateways to transform existing aircraft catalog, taxonomies and documents to universal ontologies so that all aircraft documents and information exist in universally understood language.

As discussed in this paper in several places, agents and automated processes are ideal choices for making sense of ontologies and are key components in the next generation Semantic Web. However, one of the main issues is how we can publish metadata (of which ontologies are a subset) and make them available to agents.

Nokia™ has presented a model for how ontologies can be made available with minimal infrastructure changes. Nokia™ has proposed a few simple extensions to the HTTP protocol. In their model, the HTTP protocol is the Web while the HTML sits on top[22]. By hooking up with an existing core infrastructure Nokia makes the process of getting to the Semantic Web much easier. This appears to be a good solution although its acceptance at this point is still unknown.

---

[22] http://sw.nokia.com/uriqa/URIQA.html

Our industry partner for this thesis, Agent Software, is currently researching the design and implementation of an agent system for airline MRO. Agent Software is looking to past agent design patterns like reduced network bandwidth through local processing, reliance and scalability through distributed components in order to build their frameworks and applications.

In the Agent Software framework, a server runs in the background and brokers agent requests. Agents are carried in SOAP packets while the server provides access to services and the bundling up of the return agent and data as a SOAP response.

An issue that comes up with agents is how does the system that launched the agents keep in contact with the agents, for status updates or possibly to alter their plans. In the case of synchronous agents this is not a problem. On the other hand, asynchronous agents' requirements are a bit more complicated and would require further research. Another concern is agent-to-agent messaging and communications. Expanding the framework to provide a shared blackboard could possibly provide richer agent-to-agent communications without a direct link.

Bibliography

Adams J.L., Abell J.B., Isaacson K.E., "Modeling and forecasting the demand for aircraft recoverable spare parts," prepared for the United States Air Force, Office of the Secretary of Defense. Rand Corporation, 1993.

Balabanovic M. and Shoham Y., "Learning Information Retrieval Agents: Experiments with Automated Web Browsing," Proc. AAAI Spring Symp. Information Gathering from Heterogeneous, Distributed Resources, AAAI Press, Menlo Park, Calif., 1995, pp. 13-18.

Belobaba P., MIT lecture on "Airline operating costs," 2002.

Berners-Lee T., Hendler J., and Lassila O., "The Semantic Web," Scientific Am., vol. 284, no. 5, May 2001, pp. 34-43.

Bickel R.W., "Improving Air Force Purchasing and Supply Management of Spare Parts," RAND Graduate School Dissertation, 2003.

Box D. et al., "Simple Object Access Protocol (SOAP) 1.1," 2000, www.w3.org/TR/SOAP.

Box D., "Inside SOAP," 2000, www.xml.com/lpt/a/2000/02/09/feature/index.html.

Bray T., "What is RDF?" 2001, www.xml.com/lpt/a/2001/01/24/rdf.html.

Brazier F., Steen M. van and Wijngaards N., "On MAS scalability," Agents 01 Workshop on Infrastructure and Scalability for Agents, ACM Press, New York, 2001, pp. 121-126.

Dada, M., "Inventory systems for spare parts," MIT Thesis Mgt 1984 Ph.D.

DAML-S Coalition, "DAML-S: Web Service Description for the Semantic Web," Proc. Int'l Semantic Web Conf. (ISWC), LNCS 2342, Springer Verlag, 2002, pp. 348-363.

Davis J., Huhns M., and Bonnell R., "Using Objects and Patterns for Building Compliance Agents in Healthcare," Proc. OOPSLA 98 Midyear Workshop OOT for Insurance and Health Care, 1998, pp. 1-12.

Decker S. et al., "The Semantic Web: The roles of XML and RDF," IEEE Internet Computing, vol. 4, no. 5, Sept/Oct 2000, pp.63-74.

Donald J., Closs D.J., Stank T.P., 21st Century Logistics: Making Supply Chain Integration a Reality, Oak Brook, IL: Council of Logistics Management, 1999, p.6.

FAA resource on National Airspace System (NAS)
http://www.faa.gov/education/resource/National Airspace System (NAS).htm

Farkas C., and Huhns M.N., "Making Agents Secure on the Semantic Web," IEEE Internet Computing, Nov/Dec 2002, pp.76-79.

Favre D., Toole R., Vidal O., "Strategic Sourcing: Aligning procurement needs with your business goals," in Gattorna, John, ed., Strategic Supply Chain Alignment: Best Practice in Supply Chain Management, Brookfield, VT Gower, 1998, p.285.

Filman R.E., "Semantic services," IEEE Internet Computing, July/Aug 2003, pp. 4-6.

Foner L., "Yenta: A Multi-Agent, Referral Based Matchmaking System," Proc. First Int'l Conf. Autonomous Agents (Agents 97), ACM Press, New York, 1997, pp. 301-307.

Graeber R.C. and Marx D.A. "Human error in aircraft maintenance", in N. Johnston, N. McDonald & R. Fuller (eds), Aviation Psychology in Practice, Aldershot, Avebury, 1994.

Goh S.S., "Sustainment of commercial aircraft gas turbine engines: An organizational and cognitive engineering approach," MIT M.S. thesis, 2003.

Gray J. and Reuter A. "Transaction Processing: Concepts and Techniques," Morgan Kaufmann, 1993.

Guarino N., "Formal ontology and information systems," Proc. Formal ontology and information systems, IOS Press, pp. 3-15, 1998.

Guilfoyle, C., "Vendors of Agent Technology," UNICOM Seminar on Intelligent Agents and their Business Applications, 8-9 November 1995, London, pp. 135-142.

Hansman R.J., MIT lecture on "Safety, reliability, certification, maintenance," 2003.

Heimlich J., "Economic Trends and Challenges for the U.S. Airline Industry," Airline Transportation Association.

Heflin J. and Hendler J.A., "Dynamic Ontologies on the Web," Proc. Am. Assoc. for Artificial Intelligence Conf. (AAAI), AAAI Press, Menlo Park, Calif., 2000, pp. 443-449.

Hessburg J., "Air carrier MRO handbook," McGraw-Hill, 2001.

Huhns M.N., "Agent Societies – Magnitude and Duration," IEEE Internet Computing, Jan/Feb 2002, pp.79-81.

Huhns M.N., "Agents as Web Services," IEEE Internet Computing, July/Aug 2002, pp.93-95.

Hutchins D., "Just in time – 2nd edition," 1999. ISBN 0-566-07798-1. Gower Publishing Limited.

Jackman S., "MRO market value down more than 10%," Overhaul & Maintenance. Vol. 8(4). Aviation Week Group, 2002.

Kalakota R., Stallaert J., Whinston A.B., "Implementing Real time Supply Chain Optimization System," in Proceedings of the Conference on Supply Chain Management, Hong Kong, 1995.

Kautz H., Selman B. and Shah M., "ReferralWeb: Combining Social Networks and Collaborative Filtering," Comm. ACM, Vol. 40, no.3, 1997, pp.63-65.

King F.H., "Aviation maintenance management," Southern Illinois University Press, 1986.

Lea D., "Middleware for Web Services," IEEE Internet Computing, Jan/Feb 2003, pp. 28-29.

Lee H.L., Padmanabhan V., and Whang S., "The Bullwhip Effect in Supply Chains," Sloan Management Review, Spring 1997, pp.93-102.

Majima I., "The shift to JIT: how people make the difference," 1992. ISBN 0-915299-93-3. Productivity Press, Inc.

Marx D., "Maintenance human factors at the turn of the century," Engineering, Maintenance & Material (EM&M) forum. Memphis, TN, 1999.

Mitra P., Melnik S., and Decker S., "Framework for the Semantic Web: an RDF Tutorial," IEEE Internet Computing, vol. 4, no. 6, Nov/Dec 2000, pp. 68-73.

Noy N. and Hafner C., "The state of the art in ontology design," AI Magazine, vol. 18, no. 3, 1997, pp. 53-74.

Noy N. et al., "Creating Semantic Web contents with Protégé-2000," IEEE Intelligent Systems, vol. 16, no. 2, 2001, pp. 60-71.

Petrie C., "Service Agents and Virtual Enterprises: A Survey," IEEE Internet Computing, July/Aug 2003, pp. 68-78.

Phillips E.H., "MRO Operators Want to Help Airlines Manage Assets," Aviation Week & Space Technology, April 13, 2003.

Rankin W., "Boeing applications of human factors related to maintenance," Engineering, Maintenance & Material (EM&M) forum. Phoenix, AZ, 2000.

Rose J.R. and Huhns M.N., "Philosophical Agents," IEEE Internet Computing, vol. 5, May/June 2001, pp. 104-106.

Schofield A., "FAA Unveils Latest Version of OEP; Few Timetables Changed," Aviation Daily, January 9, 2003.

Staats R., Glicoes L., and Huhns M.N., "A Multiagent Environment for Department of Defense Distribution," Adaptation and Learning in Multiagent Systems, G. Weiss and S. Sen, eds., Lecture Notes in Artificial Intelligence 1042, Springer-Verlag, Berlin, 1996, pp. 53-84.

Stifflemire P.F., "Study of Aircraft Maintenance Demand and Supply," Allegheny Institute Report, June 2002.

Sycara K., Shehory O., Sukthankar G., Mukherjee V., "Agent aided aircraft maintenance," The Robotics Institute, School of Computer Science, Carnegie-Mellon University, May 27, 1999.

Tan J.J., Titkov L. and Neophytou C., "Securing Multi-Agent Platform Communication," Working Notes Second Int'l Workshop on Security of Mobile Multiagent Systems, 2002, pp. 66-72; www.dfki.de/~kuf/semas/semas-2002/WorkingNotes/semas-2002.zip.

Turkett W.H., Rose J.R. and Huhns M.N., "Massive Deliberation," IEEE Internet Computing, Jan/Feb 2003, pp. 72-75.

Vinoski S., "Web Services Interaction Models. Part I: Current Practice," IEEE Internet Computing, May/June 2002, pp. 89-91.

Walsh N., "What is XML?" www.xml.com/pub/a/98/10/guide1.html.

Wilde E., "XML Technologies Dissected," IEEE Internet Computing, Sept/Oct 2003, pp. 74-78.