Published February 1, 2003. Distribution restricted to Sponsors until May 1, 2003.

AUTO-ID CENTRE

WHITE PAPER

Auto-ID Based Control Demonstration
Phase 2: Pick and Place Packing with Holonic Control

James Brusey, Martyn Fletcher, Mark Harrison,
Alan Thorne, Steve Hodges, Duncan McFarlane

AUTO-ID CENTRE INSTITUTE FOR MANUFACTURING, UNIVERSITY OF CAMBRIDGE, MILL LANE, CAMBRIDGE, CB2 1RX, UNITED KINGDOM

ABSTRACT

A demonstration environment is being developed at the Auto-ID Centre in Cambridge, in order to show the benefits of Auto-ID technology when applied to manufacturing control systems. In this paper, the second phase of development is examined. The first phase combined conventional control with Auto-ID information. This phase specifically looks at how Auto-ID facilitates a distributed Holonic Manufacturing System (HMS).

Published February 1, 2003. Distribution restricted to Sponsors until May 1, 2003.

WHITE PAPER

Auto-ID Based Control Demonstration
Phase 2: Pick and Place Packing with Holonic Control

Biography

James Brusey
Senior Research Associate

James Brusey previously worked (for about 13 years) in computer system administration, specialising in IBM mainframe assembler. He received a B.Ap.Sci in Computer Science from RMIT University (Melbourne, Australia) in 1996. He began studying autonomous robot control in 1998 and was a team member and the main software developer for RMIT University's Formula 2000 RoboCup team, which made the finals in the 2000 games.

James' Ph.D. (submitted) is entitled "Reinforcement Learning for Robot Soccer". It developed a novel approach to bootstrapping reinforcement learning and also examined simulation-based reinforcement learning for a real robot.

Mark Harrison
Senior Research Associate

Mark Harrison is a Senior Research Associate at the Auto-ID Centre lab in Cambridge working on the development of a PML server, web-based graphical control interfaces and manufacturing recipe transformation ideas. In 1995, after completing his PhD research at the Cavendish Laboratory, University of Cambridge on the spectroscopy of semiconducting polymers, Mark continued to study these materials further while a Research Fellow at St. John's College, Cambridge and during 18 months at the Philipps University, Marburg, Germany. In April 1999, he returned to Cambridge, where he has worked for three years as a software engineer for Cambridge Advanced Electronics/ Internet-Extra, developing internet applications for collaborative working, infrastructure for a data synchronisation service and various automated web navigation/capture tools. He has also developed intranet applications for his former research group in the Physics department and for an EU R&D network on flat panel displays.

Martyn Fletcher
Applications Specialist

Martyn Fletcher works as an Applications Specialist at the Cambridge Laboratory of Agent Oriented Software Ltd. He graduated from Liverpool University in 1991 and obtained his Ph.D. from Keele University in 1997 with a thesis on the application of software agents in telecommunications for increased flexibility. He has four years post-doctoral experience in the research and design of intelligent agent systems for holonic manufacturing. Dr Fletcher is responsible for both implementing JACK in systems at Cambridge University's IfM, and consulting on agent/holonic technology to manufacturing and defence projects. His research focuses on problems related to multi-agent and holonic systems, production planning/ scheduling, and inventory management.

WHITE PAPER

Auto-ID Based Control Demonstration
Phase 2: Pick and Place Packing with Holonic Control

Biography

Alan Thorne
Laboratory Manager

Alan Thorne is the Program Manager for the Auto-ID Centre Lab at Cambridge University. Mr. Thorne graduated from Anglia Polytechnic University in Electronics and Control Systems and has a varied background in the field of Automation and Control. He has been involved in British Aerospace/IBM research projects as a systems engineer investigating flexible manufacturing systems on civil and military aircraft production. He has most recently been involved in projects relating to the development of novel AI-based control strategies.

Steve Hodges
Associate Director, Auto-ID Centre Europe

Steve Hodges is an Associate Director at the Auto-ID Centre Lab at Cambridge University. Steve received his first degree in Computer Science with Electronic Engineering, from University College London, and received his Ph.D. from Cambridge University Engineering Department in the area of Robotics and Computer Vision. His interests include embedded sensor systems, intelligent devices, computer augmented environments, mobile robotics, image processing, low-power radio communication, mass customization of consumer products and RF tagging technologies.

Duncan McFarlane
Research Director Europe

Duncan McFarlane is a Senior Lecturer in Manufacturing Engineering in the Cambridge University Engineering Department. He has been involved in the design and operation of manufacturing and control systems for over fifteen years. He completed a Bachelor of Engineering degree at Melbourne University in 1984, a PhD in the control system design at Cambridge in 1988, and worked industrially with BHP Australia in engineering and research positions between 1980 and 1994. Dr McFarlane joined the Department of Engineering at Cambridge in 1995 where his work is focused in the areas of response and agility strategies for manufacturing businesses, distributed (holonic) factory automation and control, and integration of manufacturing information systems. He is particularly interested in the interface between production automation systems and manufacturing business processes.

WHITE PAPER

Auto-ID Based Control Demonstration
Phase 2: Pick and Place Packing with Holonic Control

Contents

1. AUTO-ID IN MANUFACTURING SYSTEMS

1.1. Trends in Manufacturing

Future manufacturing systems may be quite different from those of today. Although mechanisation and automation have done much to increase efficiency, improve reliability, and reduce the cost of producing goods, traditional manufacturing systems tend to be inflexible and lack robust handling of disturbances. It can be said that they lack **agility**, in the sense that they are incapable of responding to change gracefully. An agile manufacturing system should be able to handle significant disturbances and reconfiguration changes as 'business-as-usual'.

Furthermore, in today's markets people and companies want goods that are customised to their specific needs, and they are not prepared to allow long lead times for delivery. For example, people want to specify how to mix and match constituent elements of a product via a web page and have it presented to their door the next day. Similarly, companies require smaller batch sizes and increased product variety. In other words, manufacturing businesses are being pushed, by market forces, to provide both **mass customisation** of their product families, and to react more quickly to consumer demands [8].

1.2. New Strategies for Manufacturing Control Systems

We believe that the introduction of Auto-ID technology in a manufacturing environment can provide the basis for implementing much more agile systems by providing continuous, real-time information relating to the movement of physical items. As unexpected events relating to the movement of goods occur – such as late deliveries, mis-placed components, inaccurate shipments – this can be detected and the appropriate information passed to the manufacturing control system. Auto-ID technology also provides a mechanism for implementing systems that can readily handle mass customisation, due to the unique numbering scheme used. Since individual items can be uniquely identified and tracked, small batch sizes and even one-off production runs are feasible.

The **generation** of data relating to the environment provides a starting point for building the manufacturing systems discussed above. However, they need to **make decisions** and **automatically act upon** the Auto-ID data that is generated – commonly referred to as "closing the loop" – to fully make use of Auto-ID for control [7].

Unfortunately, the traditional **hierarchical** manufacturing control systems are geared to mass production of low-variety goods to fixed schedules. Often a schedule of the manufacturing operations, on a given day, is developed well in advance. Any unexpected events that invalidate it are hard to deal with. Also, as batch sizes decrease, the complexity of the scheduling process increases accordingly, meaning that the customisation of products is difficult to deal with.

An alternative strategy is to build a **distributed** (or **holonic**) manufacturing control system [1, 2, 3, 5]. Here, rather than building a monolithic control system which is pre-programmed to act in certain ways in response to certain conditions, the system is split into a number of smaller systems that are responsible for the control of sub-parts of the domain. These smaller units (**holons**) communicate with each other, and react appropriately according to the information they pass to each other.

It is interesting to note that it is possible to build Auto-ID based control systems using a traditional, hierarchical approach, and it is also possible to build a distributed, holonic control system that makes no use of Auto-ID data. However, we believe that the combination of these two approaches provides a very powerful architecture for creating sophisticated manufacturing systems that will meet future production and consumer demands.

1.3. Building a Demonstration Environment

A demonstration environment is being developed in three phases at the Cambridge Auto-ID Centre to show the benefits of using Auto-ID for control. The first phase of the demonstration [6] integrated the basic item identification and location information generated by a minimal Auto-ID system, with a traditional robot cell controller. The Auto-ID data that was generated by this system provided a number of benefits, which are reviewed in Section 2.1. The second phase of the demonstration, which is the subject of this paper, has two main aims, namely:

1) to consolidate and build on the advantages of applying Auto-ID to control shown previously, and
2) to integrate a distributed control architecture for the manufacturing system.

By demonstrating these features, we hope to better understand the issues related to implementing both distributed control and Auto-ID based control systems in real industrial applications. We also hope to promote both of these approaches by demonstrating the working manufacturing cell to visitors and collaborators.

1.4. Overview of Paper

The following section overviews the demonstration environment, and includes a comparison of the environment that was used in phase 1 with that used currently. In Section 3, the base architecture which provides a foundation for the holonic system, is examined. Section 4 examines the design and implementation of the holonic system is examine in detail. The paper concludes by running through some of the benefits of using a system such as the one described in this paper, and mentions some avenues for future research.

2. AN OVERVIEW OF THE CAMBRIDGE DEMONSTRATION ENVIRONMENT

The demonstrator environment at Cambridge is a simplified automated robotic packing cell that is used for the packing of men's personal care gift boxes. A variety of Gillette products are used in the demonstration, namely razors, deodorant, shower gel and shaving cream. These items may be packed into empty gift boxes – there are two different styles of gift box, each holding any mix of three products. For simplicity, the different products to be included in the gift boxes are pre-mounted into individual, standard sized, plastic item carriers to enable simple robotic handling, stacking and packing. These item carriers and the gift boxes are tagged using RFID labels pre-programmed with EPCs™, and the packing cell is populated with a number of RFID readers. Figures 1 and 2 show two varieties of packed gift boxes.

Figure 1 & 2: Mixed products packed by Cambridge packing cell

The packing cell is based around an anthropomorphic robot arm, with conveyor loops on either side that are used to carry both items and empty boxes into the cell, and to allow packed boxes to leave the cell. A storage system (or mini warehouse) acts as a buffer that allows the cell to operate when the rate and type of arrival of items does not match the box packing requirements.

2.1. Review of Phase 1

The Phase 1 demonstration [6] successfully showed some of the advantages that result when Auto-ID data is passed into the manufacturing control system. Figure 3 gives a plan view of the cell hardware that was used in Phase 1. The generic benefits of Auto-ID based control, which were demonstrated in a number of scenarios, are listed below:

- Ability to deal with mixed products.
- Capabilities for late stage customisation of products.
- Coping with random/unconstrained arrival of raw materials.
- Dealing with product disturbances in the warehouse area.

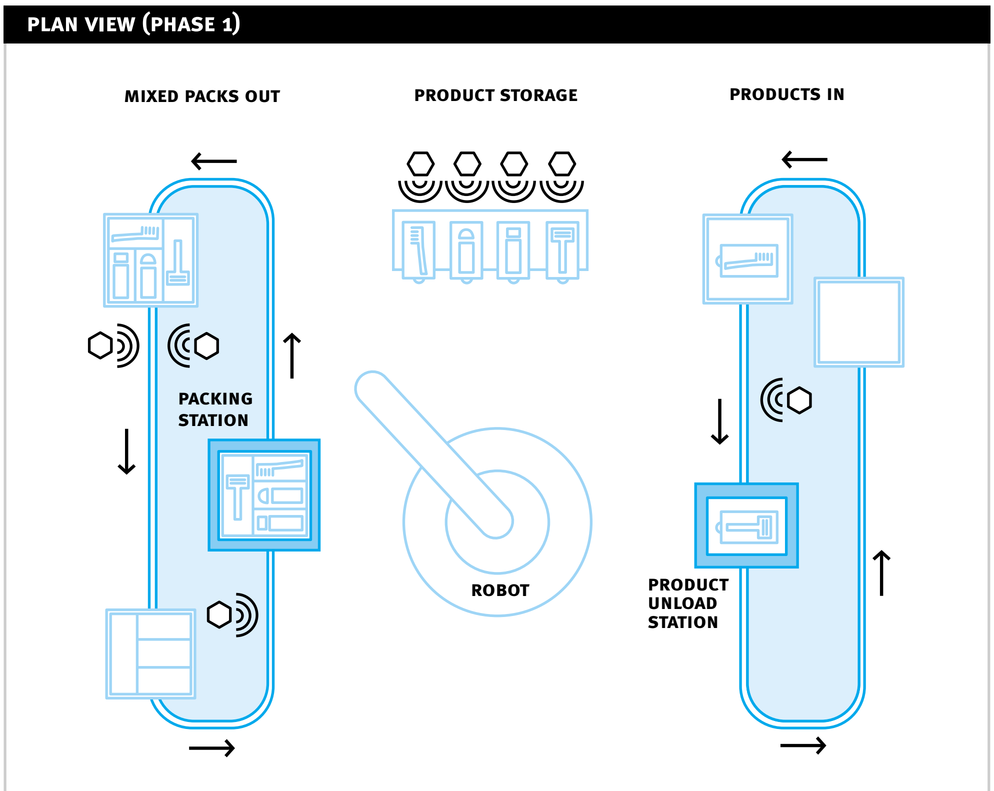

Figure 3: Schematic of the packing and storage system for phase 1.

2.2. Changes developed for Phase 2

Various enhancements have been made to the packing cell for the second phase. This enables the extra benefits that are possible when integrating Auto-ID data into a distributed control system to be demonstrated. The hardware changes are listed below, whilst the Phase 2 packing cell can be seen in Figure 4:

- A main conveyor feeder loop had to be installed. This feeder loop allows dynamic routing of items, boxes and packed boxes between the packing stations and the quality control area.
- The quality control station was moved onto the feeder loop.
- Two gate sets provide for routing of the items, boxes and packed boxes between the conveyor loops. The gate sets can be thought of as points on a railway line.
- The operation of the docking stations has been made generic allowing each docking station to perform a locating task for either incoming items needing to be stored, or boxes needing to be packed or reworked.
- The material handling operation of the robot was expanded, by teaching extra pick and place locations, thus ensuring the generic functionality of the docking stations could be utilised by the robot.
- The flexibility of the conveyor shuttles was enhanced, allowing them to be used for transporting either boxes or parts as the situation demands.
- Extra tag readers were added at decision and operation points throughout the packing facility, enabling the control system to exploit timely, item level data provided by the Auto-ID infrastructure.

This last point is helpful to the holonic control infrastructure, as will be discussed in Section 4. By arranging a reader at the point of entry to each gate and docking station, the holonic control is simplified as the event of a shuttle tag being read corresponds to a decision that must be made.

Figure 4: Schematic of the packing and storage system for phase 2.

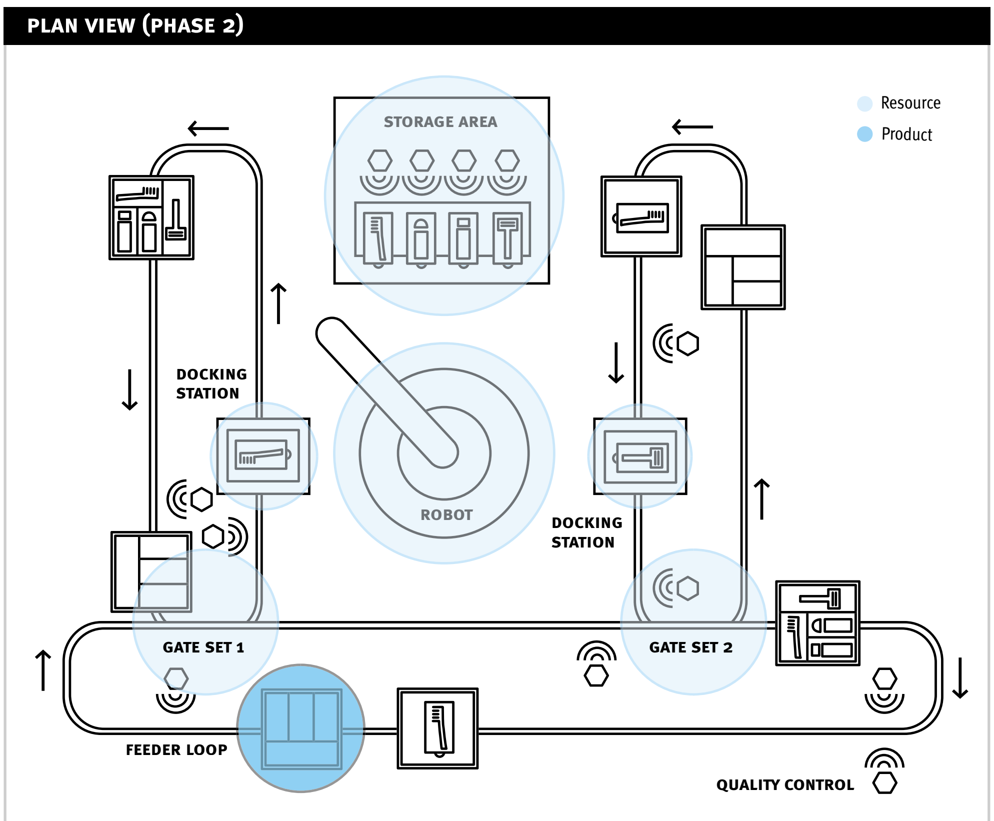

3. THE BASE ARCHITECTURE

In the following section, the base architecture of the cell is described. This is divided into the low-level architecture, which includes the PLC control and the Robot controller, and the virtual warehouse, which provides information about the state of the items and products in the cell.

3.1. The Low-Level Architecture

A critical aspect in the design of the low-level architecture was to ensure that different resources were controlled independently. This enabled the use of holonic control at the high-level. In addition, it meant that control software could be replicated along with the physical resources that it controlled. This allowed, for example, the gate control software to be developed in a modular fashion. It may seem surprising that this is not normally done, but it is more common for control of a cell to be monolithic in structure as this allows the designer, when designing control for one part of the system, to make use of information available from another.

Another important aspect of the low-level architecture is that it provides a robust control interface to higher level systems, so that these might be developed without concern for damaging equipment. In addition, the internal operation of the resource is fail-safe in nature, for similar reasons.

3.2. The Virtual Warehouse

For Phase 2, we have developed a more generic PML server for storing information about the products and batch orders. The PML server is now implemented as a SOAP (simple object access protocol) service with a set of generic methods to retrieve the whole PML or read/write to any node within the PML hierarchy, in order to access any particular property of the object. A more detailed description of the recent developments with the PML server is given elsewhere [4].

We have also developed additional SOAP services and an underlying PostgreSQL relational database, to implement a "virtual warehouse" data representation of the current state of the system, such as the location of shuttles, or the identity, location and count of items in the storage units. The virtual warehouse receives software events from Savant™ upon the addition or removal of tags from the field of RFID readers, and uses these to maintain a representation of the current **state** of the system. For many situations it is just as important to know the states of the system, while for other situations it is important to be able to trigger events on detection of changes between states. The early version of Savant has a strong emphasis on reliable filtering of data from RFID tag reads and generating software events. It can be configured to selectively filter into logical queues, using pattern matches on the reader EPC™ or tag EPC™. The virtual warehouse is intended to complement the event-driven Savant™ by storing the current "state" of all the properties of the system. The control software is also able to update the virtual warehouse, and can, for example, mark a docking station as unavailable if it is aware of a blockage or other problem which prevents normal operation.

A major motivation for the development of the virtual warehouse was the need to be able to demonstrate and control the Phase 2 manufacturing environment over the Internet. We have developed a web interface to the manufacturing facility that displays a plan view of the facility and shows such things as the inferred locations of shuttles (see Figure 5). The contents of shuttles (unpacked items/empty boxes/ packed boxes) are also depicted, drawing on updates to the PML of the batch order upon packing and 3-level aggregation inferences performed as a shuttle passes through the quality control station (for example, inferring that an item is within a box that is upon a shuttle). The web interface also shows the number of items in each of the storage chutes and the EPC™ of the item available at the base of each chute.

Figure 5: Web-bases control interface to phase 2 packing cell.

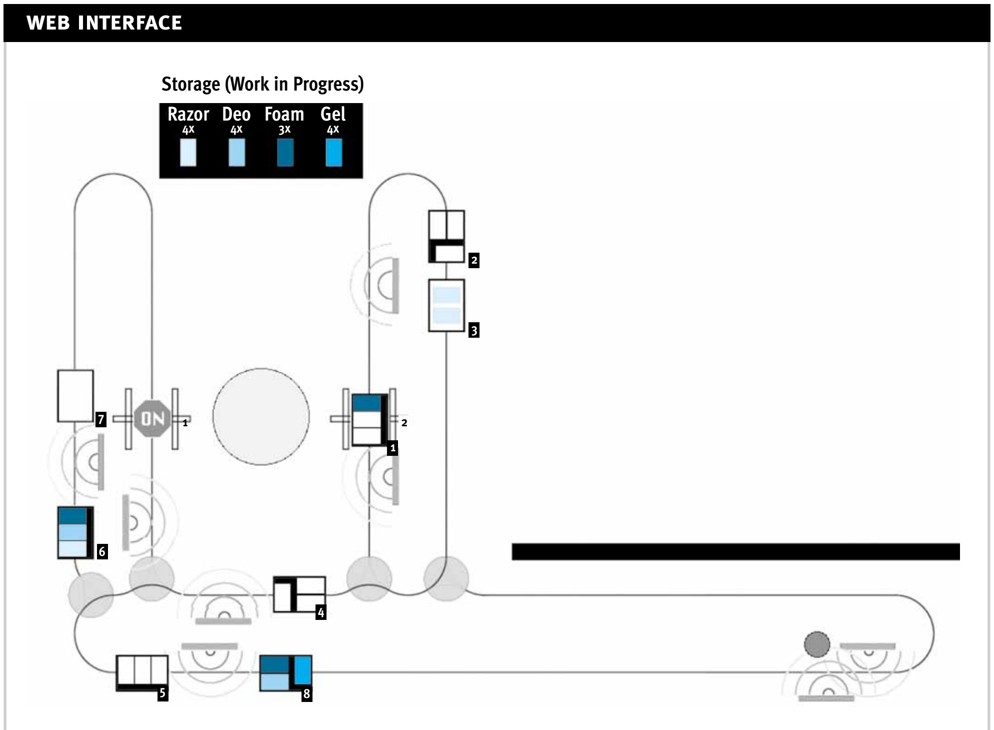

In addition to the web interface, we have also deployed three web cameras that send live video images via the Internet for remote monitoring of the cell. One camera is positioned to give an overview of the cell, while the other two show activity at each of the docking stations where packing and unpacking are performed.

4. THE NEW HOLONIC ARCHITECTURE

Rather than developing a holonic version of the system developed during Phase 1 from scratch, the holonic architecture was built on top of the existing control infrastructure that was described in the previous section. In the following section, we provide a review of the background for this architecture, and explain the original ideas and framework behind the design of the system, and then describe the system as it was implemented.

4.1. Background

A **holon** (from the Greek **holos**, meaning whole, and the suffix -**on**, meaning particle or part) is a well-defined component of a larger system, that has interactions with other holons in the system. A holon may be comprised of sub-components, which might also be considered holons. Apart from where they interact with other holons, they are autonomous entities. Holons are important conceptually because they help us divide a large system that has complex internal interactions into a number of smaller entities that each have both internal and external interactions.

Holonic Manufacturing Systems (HMS) use the concept of holons as the basis for the design of a manufacturing system. HMSs gain robustness and flexibility by allowing individual holons to be autonomous but also allowing for interaction and negotiation between holons. Such systems have increased robustness by avoiding any central point of control, and by reducing the amount of coupling between parts of a system, and thus reducing the impact of a failure in any single component. They are more flexible since holons can potentially be added or removed at runtime to allow for changes in manufacturing workloads and product characteristics. The holonic approach departs from traditional manufacturing control, which tends to be hierarchical, centralised and inflexible.

4.2. System Design

The theoretical or high-level design of a system specifies the objectives of the manufacturing system without defining how those objectives are met. The framework for this specification is illustrated in Figure 6.

Figure 6: Design framework for generic holonic manufacturing systems

1. Influences Scope
2. Constraint
3. Function & Performance
4. Function & Performance

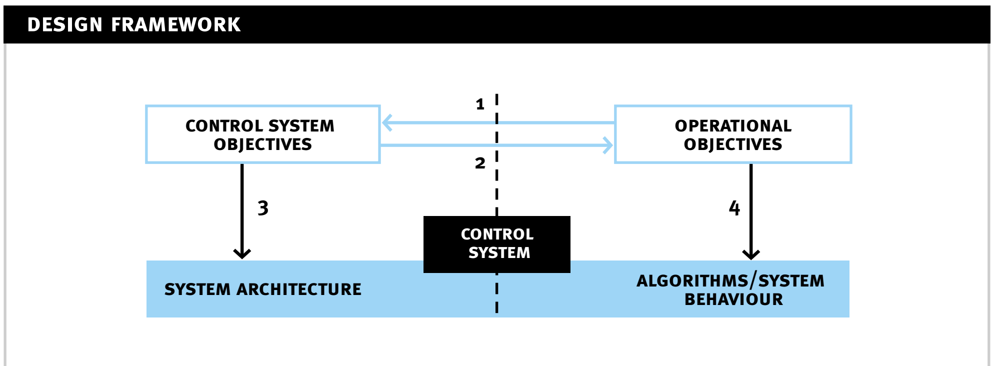

As the diagram shows, the overall objectives for a system can be split into control system and operational objectives. Control system objectives relate to the functionality of individual holons, whilst the operational objectives relate to the behaviour of the whole system. There is an interplay between the two, as control system objectives may constrain what can be achieved operationally. Conversely, an expansion of operational objectives may introduce additional requirements for the control system.

Both sets of objectives can be further divided into function and performance objectives. For example, as part of the control system objectives, there may be a functional requirement for a holon to manipulate a product in a particular way. In addition, there tend to be performance objectives that are specified in terms of specific measures, such as how many products can be processed per minute.

In a typical manufacturing system, control system objectives include such things as the ability to reliably deliver raw materials into the system (materials delivery), to manage parts during production (product management), to allow movement of materials and parts around the system (material handling), to store or buffer materials or parts (material storage), and to extract out finished products (material removal).

In contrast, operational objectives define how the system appears from the outside. This begins with a specification of what the inputs are and what is produced but also includes aspects to do with the system's agility, such as how flexible it is in terms of handling different types of raw material types, whether delivery and removal facilities can be interchanged, the ability to reconfigure to meet different demands, or the ability to handle growth. The control system and operational objectives lead to a specification of the system architecture and overall system behaviour, respectively.

In the following section, we demonstrate how this framework applies to the Cambridge packing cell.

4.2.1. Control System Objectives

The first issue in the design of the control system is the delivery of materials. Materials enter the system by being placed on shuttles and transported around the system. This has an associated shuttle loading holon. Similarly, finished products leave the system by being removed from shuttles.

Product management is assigned to a product holon. This holon tracks the product through the system and ensures that it is manufactured correctly. Material handling is provided by robot holons, loading/unloading station holons, and gate holons. These holons manage the physical resources to manipulate and move raw materials and products. Material storage is provided by stack holons that manage the storage stacks. Since raw materials can also be stored on shuttles, shuttles correspond to "virtual stack holons" that can be called upon when stack levels become low.

4.2.2. Operational Objectives

From the operational perspective, the main objective is to fulfil orders. Although there need not be a holon to manage each order, the system must be able to associate an order with a set of product types and quantities. Furthermore, each product type has an associated recipe. This describes the parts that constitute the product and how they are assembled.

When an order arrives, product holons are generated and they seek and allocate their required material resources, such as empty boxes and parts to be packed into the boxes. They also need to allocate appropriate hardware resources by negotiating with the associated resource holons. The product holons then apply their recipe to create the product. A quality control phase ensures that the recipe has been completed and that the product has been put together correctly.

To allow the system to be flexible and responsive, product holons have an associated priority so that it is possible to give preferential passage to parts associated with high priority products. In the extreme case, partly-created low priority products can be disassembled to allow higher priority orders to be filled. To increase the agility of the system, an additional design objective is to allow late changes to orders. This might be simply to change the product type or to manipulate some other aspect of the recipe, or it may be to increase or decrease the quantity.

Figure 7: Original (generic) design of holonic packing system

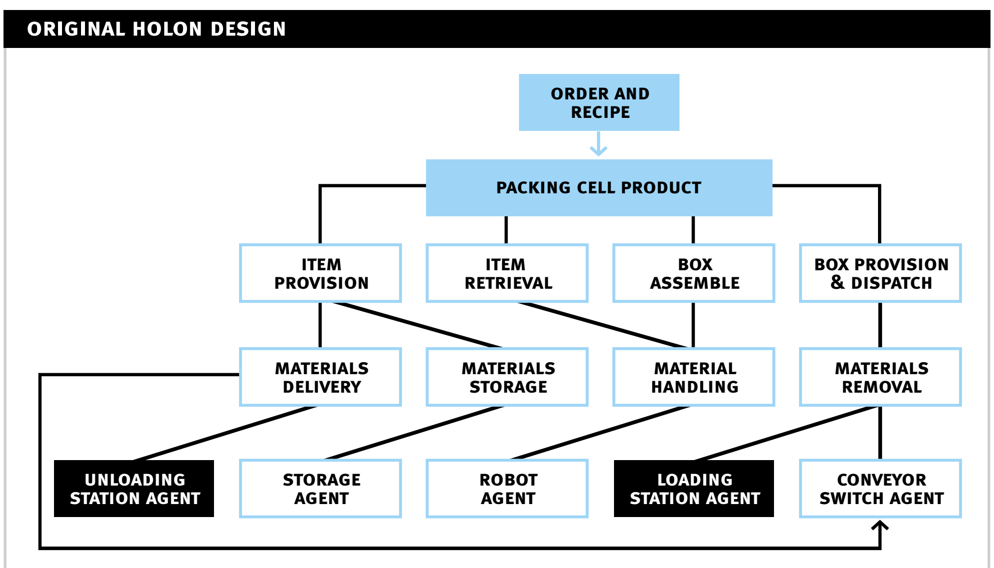

Figure 7 summarises the original design of the packing system. Note that unloading and loading holons provide similar functionality and could potentially be merged to provide a single holon to manage both operations. This seems a sensible simplification since the same resource is typically used for both operations.

4.3. Implemented System

A key characteristic of holons is that they are typically made up of a number of resources that have functional cohesion. These resources can include people, software components, machinery and parts. Two example holons identified in this design are the robot and gate holons. As Figure 8 shows, the robot holon is comprised of a specialised software agent, the blackboard system (denoted "BB" in the diagram), a bridging interface, and the robot itself. The software agent interacts by writing messages to the blackboard. The robot bridge periodically checks the blackboard for commands and when it finds them sends them to the robot. The commands are always similar in form: pick up an object from one position and place it in another. Logic in the low-level robot controller performs each operation.

Figure 8: Holons can be made up of a combination of software agents and physical resources or other elements.

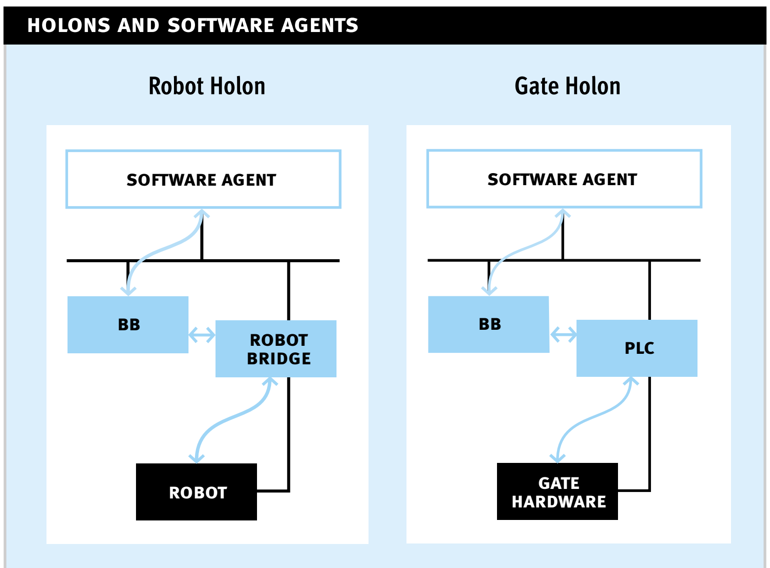

In this way, the robot holon consists of a number of software and hardware components, and importantly the software to drive the robot is split over a number of platforms.

The gate holon operates in a similar way. As Figure 8 shows, the commands to the gate are sent via the blackboard. This time, a PLC (programmable logic controller) interfaces directly with the blackboard system and switches the gates as required. Note that each gate holon actually controls a gate pair.

4.3.1. Holon Types

The holon types that have been implemented in the control system (with the number of instances) are: Order Holon (an instance is spawned for each box ordered), Production Manager Holon (1), Gate Holon (2), Reader Holon (7), Docking Station Holon (2), Robot Holon (1), Storage Holon (1), Box Manager Holon (1), and Track Holon (1). Broadly speaking, these holons can be divided into resource holons that represent physical resources, and order holons that represent the set of requirements associated with each order.

The software agent part of each of these holon types has been encoded using the JACK Intelligent Agents™ environment from Agent Oriented Software Limited.

We focus now on the features of the robot holon software agent because it is one of the more complex resource holons, and its processing typifies the interactions an autonomous holon has within itself and with the external world. It can be observed that the robot holon supports the material handling functional objective via:

- Scheduling jobs based on reward.
- Performing various pick and place operations in order to pack boxes, unpacking boxes, sorting the storage area and unloading items into storage.
- Interacting with the physical robot.

Each of these is modelled as a capability (containing appropriate events, plans and belief structures) within the robot. Using JACK's design tool, we can represent these agent, capability, plan, event, and belief relations (see Figure 9).

Figure 9: Example of agent design diagram showing beliefs, capabilities and plans.

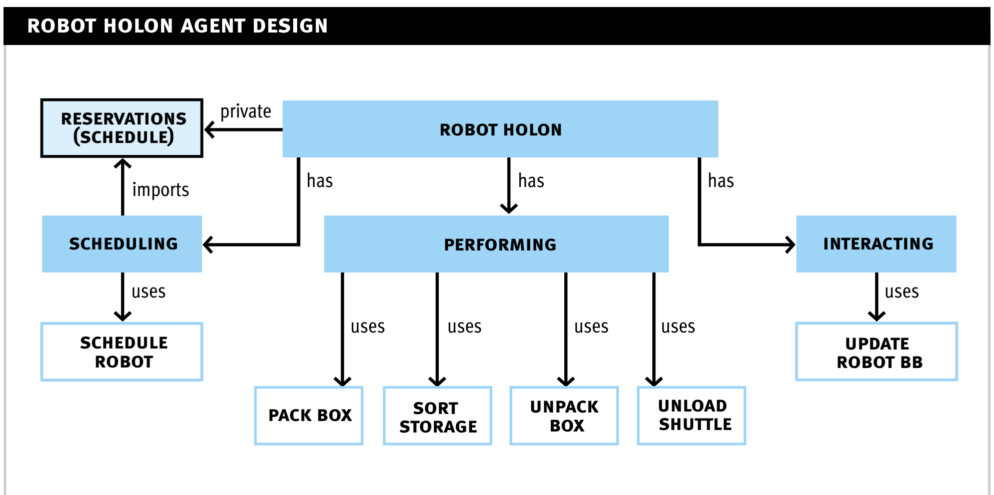

The robot holon takes action when its software agent receives a `ROBOTJOBARRIVED` event. Note that a `SCHEDULEROBOT` plan takes care of sorting jobs in order of priority. Four plans in the `PERFORMING` capability can handle this event. The choice as to which one will handle the event initially is determined by the number of boxes and items associated with the shuttle that has docked. For example, the `PACKBOX` plan is used if,

$$N_b = 1 \wedge N_i = 0,$$

where N_b is the number of boxes and N_i is the number of items associated with the `ROBOTJOBARRIVED` event. In other words, when a robot job arrives requiring one box and no items, this triggers the box to be packed. The `UNLOADSHUTTLE` plan, which takes raw materials and puts them into the stacks, is used when,

$$N_b = 0 \wedge N_i = 2.$$

Figure 10 illustrates this event/plan relationship. The remaining events handled and sent by the ROBOTHOLON have been removed from the figure for clarity.

Figure 10: Agent design diagram showing the RobotJobArrived event

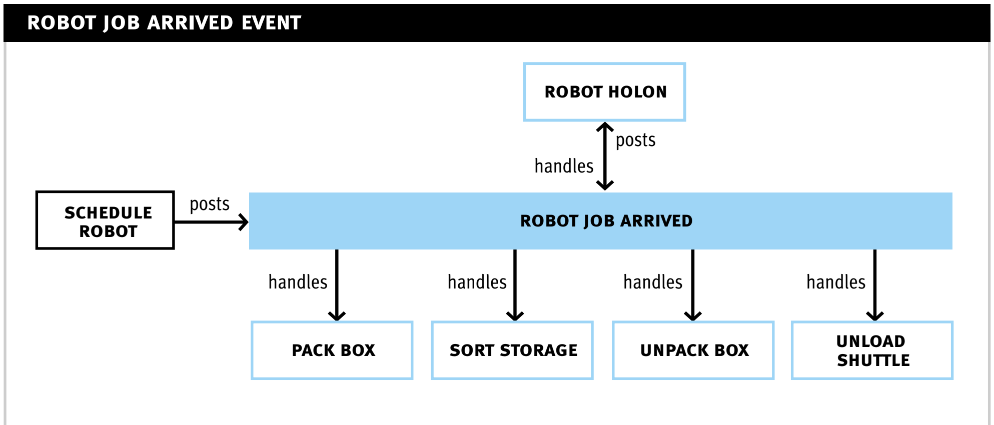

4.3.2. Holon Interactions

Interaction between holons is largely performed via software agent messages. These messages can be either synchronous, where the sender waits for a reply, or asynchronous. An example of such interaction is given in Figure 11. The figure shows a synchronous message where the robot holon asks the order holon for the order parameters. This is done using a QUERYORDERPARAMETERS event. On the other hand, the PACKBOX event is an asynchronous message and corresponds to a notification from the robot holon to the order holon that box packing has been completed. In this case, no reply is necessary.

Figure 11: Agent design diagram showing the message interaction between the robot and order holons.

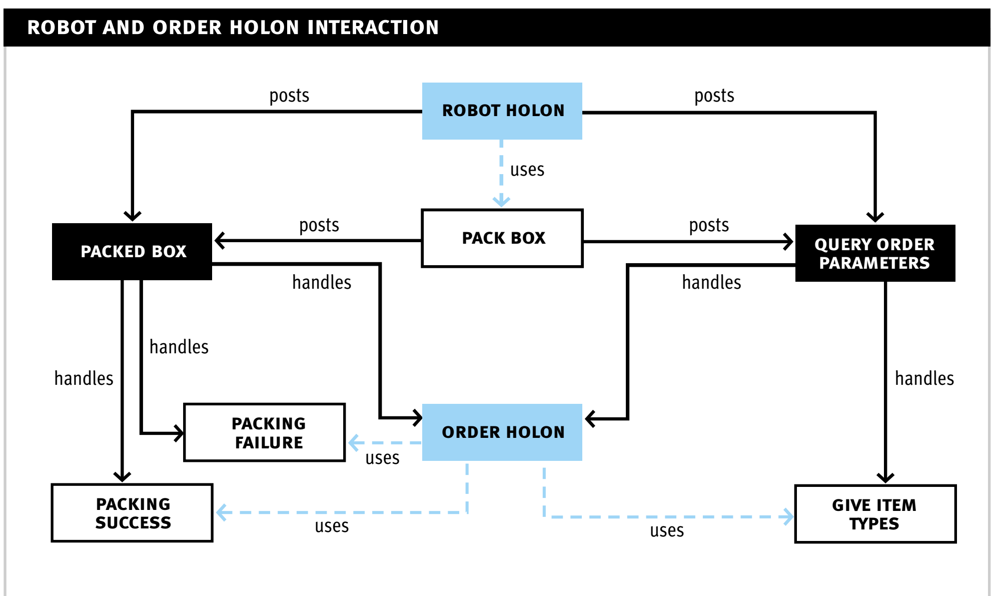

This messaging architecture is coupled with Auto-ID by triggering message events as shuttles pass readers. Due to the layout of the cell, these message events correspond to shuttles arriving at gates or docking stations. This provides the gate or docking station holon with exact knowledge about which shuttle is arriving. This reliability of information leads to a correspondingly reliable control for the whole cell.

5. SUMMARY

There are significant advantages in combining Auto-ID technology with the holonic manufacturing system philosophy. Auto-ID enables holonic control by providing reliable information that identifies each item and its location uniquely. This combination leads to a number of benefits.

5.1 Benefits Demonstrated in the Packing Cell

The product's location and status can be determined at key decision and operation points around the packing cell through the use of strategically placed RFID readers. Coupled with the distributed control architecture used in the the cell, this means that the production route of a product can be varied to make best use of the available physical resources. If either packing or unloading station fails due to a hardware fault, the remaining station can be used to carry out both tasks. In this way, the production system is agile in the face of unexpected conditions. Similarly, the system can handle returned goods seamlessly, because those items can be uniquely identified within the cell. Small batch sizes are also facilitated through the ability to identify items at the product level – this is demonstrated in the packing cell with the ability to suspend current work in order to do rush orders.

5.2. Generic Benefits

Of course, the specific benefits gained by integrating Auto-ID data into a distributed control system and demonstrated in the cell demonstration environment are applicable much more generically. In particular, distributed Auto-ID based control:

- Enhances the customisation capabilities of the production system by supporting smaller batch sizes and even mass customisation.
- Enables novel scheduling techniques, allowing dynamic response to unexpected conditions.
- Makes better use of elements in the production facility.
- Provides graceful degradation paths when dealing with equipment failure.

These benefits are in addtion to those demonstrated in Phase 1, namely:

- Ability to deal with mixed products.
- Capabilities for late stage customisation of products.
- Coping with random/unconstrained arrival of raw materials.
- Dealing with product disturbances in the warehouse area.

5.3. Future Work

There are a number of interesting avenues in terms of future research at the Cambridge Auto-ID Centre, using our unique laboratory demonstration environment. These areas include:

- Abstraction and extensions of distributed Auto-ID based control. It would, of course, be useful to abstract more generic learnings about distributed Auto-ID based control systems from our specific laboratory implementation. It may, in time, also be possible to integrate these with experiences gained in real industrial applications as the technology is trialled more extensively. Similarly, it would be useful to develop more quantitative metrics for designing and benchmarking Auto-ID based control deployments. The use of holonic manufacturing systems is also a new topic, and there will undoubtedly be on-going developments in this area.

- Creating a more dynamic production environment. The current environment is capable of reacting to the failure of one of the two docking stations by routing shuttles to the remaining station. It would be better still if such failures could be handled by making use of **different** types of equipment – for example a Fanuc SCARA robot arm might be used in place of an ABB anthropomorphic arm. The ability to do this would rely on the abstraction of the various tasks as generic 'recipes', rather than specific machine instructions – and of course, the recipes need to be automatically translated into machine instructions. With such a mechanism in place, it is much easier to explore the possibility of multiple productions paths, or routes through a factory, for similar orders. This would allow the production system to handle errors and failures much more readily. It may also be possible for materials handling systems to use descriptions of items, in the form of PML data, to adapt to those items – in this case turning parameterized item data into machine instructions – thereby handling materials intelligently.

- Considering the data flows required between different companies in the supply chain. As well as looking at the movement of items through a single factory environment, it is useful to look at the way in which Auto-ID data may enhance the movement of those items throughout the whole supply chain. One approach would be to build a 'mini supply chain' within the lab in Cambridge, emulating a multi-organisation setup. An ambitious goal is the integration of our laboratory equipment with traditional higher-level software systems, such as MES and ERP packages.

- Testing of the latest technological developments from the Centre and its sponsors. New and enhanced technologies are being developed in a number of areas by the Auto-ID Centre and various sponsors, and the Cambridge demonstration environment is a valuable test-bed for the integration of these. Examples include cheap tag and agile reader hardware, which need to be tested in real, industrial environments, Savant™ software developments, including more sophisticated plug-in filters and combinations of these filters, and ONS and PML server infrastructure.

- The integration of additional sensor data, such as environmental information, with Auto-ID data. Methods for performing such integration need to be designed and evaluated, and demonstration applications need to be built. At the low-level, the sensors may be fixed in the environment, with wired power and communications. However, in some applications it may be appropriate to use mobile, wireless sensors, which may well take the form of more sophisticated RFID devices. The specification and design of such higher functionality tags – Class 2 and beyond – is in itself a big topic of research. In addition to the use of the sensory data for real-time control (in combination with the Auto-ID data), it may be beneficial to generate 'sensor data trails' – histories of sensor data, relating to machine operating characteristics as well as to the environment. These may be useful for subsequent analysis to diagnose production quality issues (that only become apparent some time later).

- Investigation of alternative approaches to representation of item location based on Auto-ID data. The nature of the RFID technology that underlies Auto-ID is in some senses probabalistic rather than absolute. A tag that is located towards the boundary of the field generated by an RFID reader will, for example, sometimes be detected by that reader and sometimes not. Similarly, a case containing many items may straddle the field boundary of a reader, so that the case itself and some of its contents may be detected, but some of the contents will not be. A probabalistic model of the environment is perhaps an attractive approach to modelling this behaviour. A related issue concerns the communication of Auto-ID information. Whilst an event-based protocol is suitable for propogating discrete location information (and is much more efficient than a state-based protocol), in a probabalistic setting this approach may need to be reviewed.

– Consideration of some deployment issues such as speed and cost of production systems. Many industrial applications are very high volume, low variety – such as the production and distribution of soft drinks or cigarettes. The design of an Auto-ID deployment in a high-volume application may need careful consideration in order to guarantee appropriate performance. A related issue concerns selecting the most suitable scale for an Auto-ID deployment in more controlled environments, where disturbances and failures are rare, to alleviate the need for fine-grained Auto-ID tracking.

We intend to build upon the Phase 2 demonstration environment discussed in this paper, by incorporating some of the ideas outlined above, for the third and final phase of our work.

6. REFERENCES

1. **F. Maturana, W. Shen & D.H. Norrie, (1999), "MetaMorph: An Adaptive Agent-Based Architecture for Intelligent Manufacturing".**
International Journal of Production Research, 37(10), 2159–2174.
http://citeseer.nj.nec.com/maturana99metamorph.html

2. **V. Mařik, M. Fletcher & M. Pěchouček, (2002), "Holons & Agents: Recent developments and mutual impacts, "Multi-Agent Systems and Applications II".**
LNAI 2322, Springer, Berlin.

3. **H. Van Brussel, J. Wyns, P. Valckenaers, L. Bongaerts & P. Peeters, (1998), "Reference architecture for holonic manufacturing systems: PROSA".**
Computers in Industry, 37, pp. 255-274. http://citeseer.nj.nec.com/483772.html

4. **M. Harrison & D. McFarlane, (2002) ,"Development of a Prototype PML Server for an Auto-ID Enabled Robotic Manufacturing Environment".**
http://www.autoidcenter.org/research/CAM-AUTOID-WH-010.pdf, CAM-AUTOID-WH-010.

5. **S. Bussmann, (1998), "Agent-Oriented Programming of Manufacturing Control Tasks".**
Proceedings of Third International Conference on Multi-Agent Systems (ICMAS'98), pages 57–63, IEEE Computer Society. http://citeseer.nj.nec.com/bussmann98agentoriented.html

6. **S. Hodges, A. Thorne, A. Garcia, J.L. Chirn, M. Harrison & D. McFarlane, (2002), "Auto-ID Based Control Demonstration Phase 1: Pick and Place Packing with Conventional Control".**
http://www.autoidcenter.org/research/CAM-AUTOID-WH-006.pdf, CAM-AUTOID-WH-006.

7. **D. McFarlane, (2002), "Auto-ID Based Control – An Overview".**
http://www.autoidcenter.co.uk/research/CAM-AUTOID-WH-004.pdf, CAM-AUTOID-WH-004.

8. **J.H. Gilmore & B.J. Pine, II (Editors), 2000, "Markets of One: Creating Customer-unique Value Through Mass Customization".**
Harvard Business School Press.

Designed and produced by Foxner. www.foxner.com