Process Solutions

Honeywell

White Paper

The Undiscovered Country: The Future of Industrial Automation

Executive Summary

The Internet of Things (IoT), and specifically the Industrial Internet of Things (IIoT), has the promise and potential to be the most influential and disruptive influence on automation systems since the advent of microprocessor based Distributed Control Systems. Early architectural styles are emerging for IoT in which ubiquitous sensing is coupled to cloud-borne Data Analytics and Storage systems. While these are certainly viable for a broad-class of IoT solutions, e.g. Smart Grid and Consumer-grade Appliances, this article considers how these architectural styles should be adapted to create IIoT architectures for major classes of industrial automation systems: Continuous Process automation, Discrete Manufacturing Systems, SCADA systems, and combinations of these three. Additionally, the question of how the conventional Purdue model for automation systems, and today’s installed base, fits with new IIoT architectures is examined. Finally, the architectural mechanisms required to deliver key system attributes such as performance, availability, reliability, safety, and security are considered. In conclusion, we canvas some of the key issues in how to migrate from existing automation systems to the IIoT-based automation system of the future. The resulting vision is a new form of automation system architecture that balances the computational and lifecycle benefits of Cloud Computing with the requisite on premise Appliance-hosted capabilities necessary to provide safe, secure and long-lasting automation for complex manufacturing systems and processes.

Table of Contents

Table of Figures

Introduction

The Industrial Internet of Things is the intelligent application of digital technology to solving the automation needs of an analog world. IIoT could equally be called the *Intelligent* Internet of Things. There are two fundamental pillars to IIoT: digital automation systems, and the Internet itself. Although current DCS, PAS and Discrete manufacturing systems have been described as IIoT solutions, they are not true IIoT systems without the Internet, and Internet based cloud-borne technology. That is not to say that today's automation systems are outmoded; they do in fact play an important role joining with the Internet and Cloud Computing to form an IIoT architecture. And, given the widespread adoption of digital Industrial Automation systems to current manufacturing operations, their evolution to IIoT must be managed in an intelligent and beneficial manner.

This paper explores three key topics: the core nature of an IIoT, a classification of the main forms of an IIoT architecture, and the evolution of today's digital automation systems to IIoT. The goal of the whitepaper is to orient the reader to the key characteristics of IIoT architectures and how they apply to the industries served by Honeywell Process Solutions (HPS). The whitepaper provides the basis for an exploration of how the HPS strategy can be used to serve both existing and new markets.

Section 2 provides a brief introduction to the Industrial Internet of Things (IoT) and several of its main variants in terms of their general architecture and key elements. Section 3 describes the IIoT in particular and how it relates to both the IoT and distributed control systems as they currently exist. Section 4 provides an overview of key issues and enablers for migrating to IIoT-based architectures.

The Internet of Things

The development of the Internet over the past three decades has led to connectivity between people, organizations, and businesses on a scale that would have been difficult to imagine when it first emerged in the 1980s. This ubiquitous connectivity is rapidly extending beyond people to "things" as all manner of devices, sensors, controllers and actuators become connected in what is now referred to as the Internet of Things, or IoT. However, simply connecting vast numbers of objects from daily life into an Internet of Things is not sufficient to enable interesting and useful new ways of living and doing business unless there are platforms, tools, algorithms and applications to analyze, distribute, and act on the huge amounts of data that result from this connectivity. Consequently, the IoT, as it is currently understood, lies broadly at the intersection of ubiquitous device connectivity, cloud storage for the very large amounts of data produced, statistical and machine learning algorithms for analyzing and acting on that data, and new human computer interaction technologies provided by mobile and wearable computing devices.

The Emergence of IoT

IoT has its roots in the early 1970's and can be considered to have an epoch date of 1969, the year that the Internet itself (then ARPANET) was first deployed, when UNIX was released by Bell Labs, and when Honeywell first conceived of a micro-processor based Totally-Distributed Control system (TDC-2000). Given the premise that IoT is based on the harmonious alignment of the Internet to smart digital sensors and devices it is clear that ARPANET and TDC-2000 are foundational pillars of IIoT. UNIX is foundational as well, as it formed the underlying basis and structure for client-server computing, workstations, Personal Computing, server farms and virtualization.

Evolution of these core systems has yielded smart, Internet connected sensors and ubiquitous computing that weave computing into every aspect of life by allowing it to occur anywhere in formats that make sense in any particular situation. The realization of this vision is being driven by a wide range of networking technologies such as cellular wireless, corporate and municipal Wi-Fi connectivity, and short range connectivity via Bluetooth and Near Field Communication technologies. This pervasive networking environment has been coupled with a wide range of sensors built into household appliances, homes and buildings, personal devices measuring a variety of health parameters, and so on, together with user interfaces provided by web browsers and applications running on devices ranging from desktop computers to smartphones, tablet computers, and smart watches. The result is a range of new applications that reveal new insights into our own lives and the world around us.

The current IoT landscape is characterized by a large number of emerging application areas and supporting technologies, many of which are still in the early stages of development. The result is that much of the promise of IoT is still to be delivered and there will inevitably be a certain amount of churn as competing approaches attempt to establish themselves and the hype surrounding IoT meets the reality of deploying commercially viable solutions. Nonetheless, the confluence of technology coming together in IoT approaches do enable new sorts of applications and business models that will undoubtedly create new markets and disrupt existing markets.

The Internet of Things has, to a large extent, been enabled by the rapid emergence of a series of technology inflections. These technologies are virtualization, cloud computing, pervasive networking, big data analytics and machine learning, smart devices, mobility, and cyber security. These technologies enable the new types of systems typical of what is recognized as the IoT.

Virtualization

Virtualization technology that allows software workloads (entire operating systems, individual applications, etc) to be decoupled from the hardware on which they run allows a range of deployment scenarios that can significantly reduce cost, simplify management, improve availability, and avoid problems associated with churn in the underlying hardware platforms. Virtual deployment of computing resources can be either on premise or off premise.

Cloud computing

Cloud computing provides virtualized platforms with elastic compute and storage capabilities. Cloud platforms are usually run as a service based on large data centers that make it possible to easily acquire additional computational (i.e. CPUs) and storage capacity (i.e. disk space) as it is needed. This can entirely remove the need for a software-based enterprise to acquire and manage their own computing infrastructure.

Pervasive networking

The definition of the “Internet of Things” is sometimes abbreviated to simply “everything connected to everything else”. While this is inadequate as a definition, it does point to an essential, enabling element of IoT and IIoT – ubiquitous connectivity through pervasive networking. More and more devices, both in consumer, commercial, and industrial markets come equipped with some form of connectivity. This connectivity can take the form of direct Internet connectivity via 3G and 4G cellular networks, indirect Internet connectivity via Wi-Fi, or local connectivity via Bluetooth or Near Field Communication. This allows devices to participate directly in cloud-based services or indirectly through gateway devices such as a smart phone that is connected to a cloud-based service. Pervasive networking provides opportunities for establishing relationships between elements of the physical world that do not exist otherwise, enabling a range of new applications and services.

Big data analytics and machine learning

The elastic compute and storage provided by cloud computing together with pervasive networking makes it possible to collect very large amounts of data from an increasingly wide range of sources. Collecting lots of data about a lot of things (big data) provides opportunities for analysis of phenomena not possible otherwise. For example, analysis of energy usage across entire service jurisdictions based on information from individual metering devices located with the final consumer allows analysis and management of energy provision in new and more effective ways. Big data is characterized by a large volume of data (in the order of terabytes and petabytes) that requires new techniques to store and analyze, such as massive parallel data stores and statistical techniques or identifying correlations and patterns in data. The availability of data for analysis has also spurred the application of machine learning techniques including artificial intelligence algorithms to the big data stores as well.

Smart devices

Not only are devices becoming more connected, they are becoming smarter. The availability of small, low power processors that can be embedded in many devices allows them to act as more than mere sensors or actuators. Local computational resources allows devices to act on their local environment becoming more interactive and autonomous. Examples include smart key fobs that keep track of personal items and smart home monitors that learn daily routines to automatically control lighting, climate, etc. The trend toward increasing connectivity and capabilities in a wider range of finer grained smart devices is also known as ubiquitous computing. The aim of ubiquitous computing is for computers to blend into the surroundings so that they become an integral part of the physical and virtual world that are available when needed without having to explicitly use any form of conventional computer interface.

Mobility

One area where smart devices are making a very large impact is in the area of mobile computing. Smart phones and tablets enable a wide range of highly interactive context sensitive (location, time, task, etc) applications. However, the current trend is toward disassembling the smart phone and distributing its capabilities across a series of smart wearable devices. The current crop of smart watches and head up displays are a good example of this. Smart wearables are able to collect more data from a wider variety of sensors (cameras, voice, gesture, biophysical, etc) and render information across a wider range of modalities (vision, speech, haptics, etc). This enables richer applications that are usable in a wider variety of circumstances (hands-free, eyes-free, etc).

Cyber security

Cyber security has become a predominant factor in any computing system and especially in systems that leverage the technology inflections outlined above. This is particularly true given the increasing pervasiveness of networking and connectivity, which creates more complex webs of communication that need to be secured to ensure that the operation of systems and the integrity of intellectual property are protected.

IoT architecture

The technologies outlined above come together in a general architecture that consists of three main domains – the cloud, the network, and the edge – as illustrated in Figure 1.

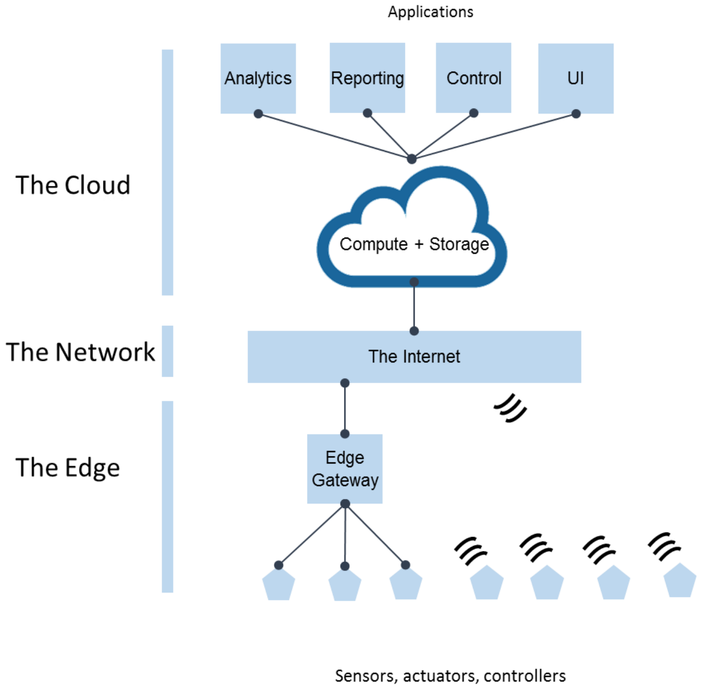

Figure 1 Internet of Things Architecture

The cloud includes compute and storage mechanisms together with applications including analytics, reporting, control and user interfaces. The user interfaces may actually live at the edge and are often combined with sensors as in the case of smart phones.

The network is the means by which the components of the architecture are connected together. This connectivity is built on IP based protocols, some of which are conventional protocols such as HTTP with others being more specialized protocols designed specifically to enable IoT-based applications involving large amounts of data collection and distribution.

The edge consists of the “Things” in the Internet of Things such as sensors, actuators, and controllers. In some cases these devices are connected directly to the network via 3G/4G cellular or Wi-Fi. In other cases an intermediary, an edge gateway, provides connectivity to one or more devices that support only local connectivity.

The next section explores the specialization of IoT for Industrial Automation – the IIoT.

The Industrial Internet of Things

As noted in the previous section, the core ideas of the Internet of Things have broad applicability. The Industrial Internet of Things is, in broad terms, the application of these ideas to the planning, running, analysis, and optimization of industrial enterprises. The application of these ideas to industry are being developed in a number of initiatives such as Industry 4.0 and the Industrial Internet Consortium and aim to bring together the means of production (the physical plant) with the advanced Internet-based computational and analytic capabilities to create “cyber-physical” systems that transcend the capabilities and scope of current automation systems. In simplistic terms the IIoT connects the world of industrial Things (sensors, actuators, controllers, robots, etc) to computational capabilities residing in Internet-based storage and analytics.

In attempting to understand the particular characteristics of the IIoT and how Honeywell Process Solutions intends to leverage IIoT approaches in future solutions the next two sections compare IIoT to more generic IoT concepts and to existing automation systems such as Honeywell’s Experion PKS. The following sections then describe the key elements of the IIoT approach and the ways in which it enables an industrial enterprise that is safer, more secure, more reliable, and more efficient.

IIoT vs IoT

IIoT differs from the more generic concept of the IoT with respect to several key quality requirements that result in architectures that expand on IoT approaches. A fundamental difference is that IIoT aims to enhance the operation and management of industrial production processes, many of which involve exothermic reactions for which safety is a primary concern. Security of IIoT-based systems is also of paramount importance not just from a safety perspective, but also in cases of the production of essential and strategically important goods and services. This results in more stringent security, reliability, availability requirements and the ability to continue operation with intermittent access to Internet resources. When failures do occur, the system must continue operation where possible or degrade gracefully, deterministically, and safely.

Another fundamental difference between IIoT and human and consumer applications of IoT is that an industrial plant is a very long-lived, capital intensive asset requiring long term support in the face of rapid technological advances. This requires support for existing, often ageing equipment and infrastructure and a means of protecting investments in intellectual property concerning the planning, execution, and optimization of production activities. In contrast, other applications of IoT involve short product lifecycles that are often driven by whims of fashion and budget. Consumers are willing to rip and replace to get improved functionality (e.g., IoT enabled lightbulbs, thermostats, refrigerators, etc). Similarly, manufacturers generally don’t want to retrofit existing appliances as it doesn’t move new product – it is part of a planned obsolescence that leads to increased / new sales. On the other hand, it is very expensive to shut down an industrial process to replace / upgrade equipment. Instead industrial enterprises favor keeping things running as long as possible – as exemplified by the huge spare parts business for obsoleted systems. One consequence of this is that many devices that will form part of the IIoT will continue to communicate via existing, often old protocols and will need special mechanisms to integrate them into the wider IIoT environment.

Many applications of IoT are human-centered in which information delivery and interaction is aimed primarily at human users. The IIoT, on the other hand, focuses on automation of industrial processes with a trend toward less manual human involvement in production. Human participation is still a key element, but increasing levels of automation continue to move the human participants in an industrial enterprise away from direct control of the process toward higher level planning and supervisory roles. The ultimate goal of IIoT might be considered completely autonomous “lights out” operation over an increasingly large scope of production from the autonomous operation of individual units today to the autonomous operation of an entire site or collection of sites in the future.

IIoT Architecture

Bringing IoT ideas to the industrial enterprise means reconciling and integrating them with existing automation systems. As previously mentioned, the IIoT is an extension of concepts that Honeywell pioneered in the 1970s with the introduction of the Totally Distributed Control system, a precursor to the IoT concept of Edge Computing. A large DCS is a complex network of sensors, actuators, controllers, and computational capabilities. The lower layers of DCS tend to be autonomous with responsibility for the real time control of the process and can operate with a high degree of safety and reliability. The layers above provide various supervisory capabilities including advanced and supervisory control, and Human Machine Interfaces (HMIs) for management of the process by human operators. Above this are facilities for capturing and analyzing a continuous historical record of the process and tools for planning and scheduling production activities that are passed down to the lower layers for execution.

It is tempting to draw a direct comparison between the DCS of today and the IIoT-based automation system of the future and claim that we are already doing IIoT, but to do so ignores the significant changes to the DCS as we understand it that will occur with the introduction of the IIoT. The IIoT arises from the combination of core DCS concepts such as local, high availability real-time control of industrial processes together with the technologies and architectures that enable the IoT.

Figure 2 illustrates the key elements of the IIoT.

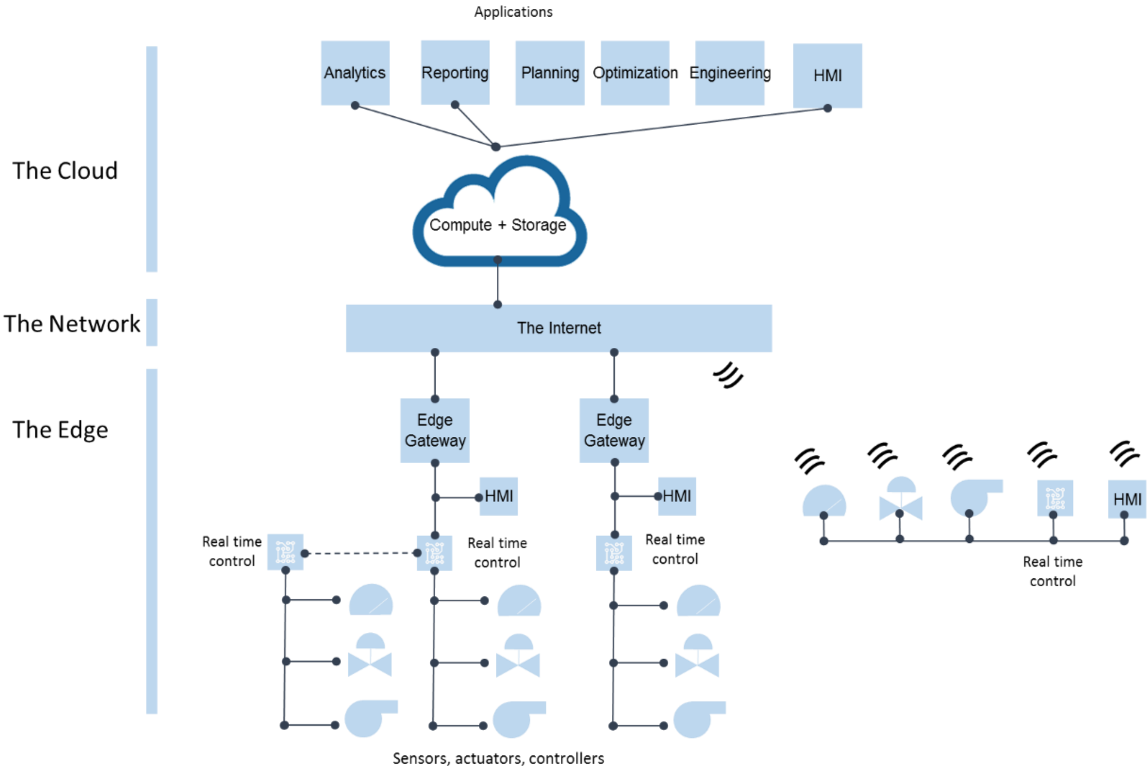

Figure 2 Industrial Internet of Things Architecture

In the case of the IIoT, the applications running against cloud-based storage include applications geared toward industrial enterprises such as planning/scheduling, optimization, and engineering.

This model provides an overview of the general structure of IIoT-based systems. This is not a one-size-fits-all model. There will be variations on this architecture for specific types of systems and customer sites. In some cases the cloud environment may be deployed on premise (either at the customer's site or in the customer's data center).

One of the main differences between IoT and IIoT architectures concerns the nature of the edge computing environment. In the IIoT the edge computing environment provides the opportunity to address key requirements in the areas of performance and robustness needed in industrial process control. Another significant characteristic of the edge computing environment in the IIoT that sets it apart from the IoT is a high degree of heterogeneity in the devices used and the protocols with which they communicate.

The IIoT edge computing environment consists of a wide range of devices including sensors, actuators, controllers, and human machine interfaces. These devices are located in close proximity to the production process and may communicate directly with cloud-based services or via an edge gateway that acts as a data concentrator and/or filter and protocol converter. Edge devices may act collectively in a federation of devices to provide an autonomous coordinated set of capabilities at the edge. For example a federation of sensors, actuators, controllers, and HMIs may provide real-time control and management for a process unit or area. Such a federation would utilize peer-to-peer communication amongst devices using a variety of protocols. While there is a trend toward open IP-based protocols in the IIoT such as OPC UA, there will continue to be a role for existing protocols such as HART, FieldBus, Modbus and so on, particularly for existing installed devices. Edge gateways are used to interface heterogeneous devices and protocols with cloud-based services.

IIoT vs DCS

Some of the key differences between an IIoT architecture and a conventional DCS architecture can be illustrated by comparing the architectures at their highest levels. The structure of a DCS and associated applications typically confirm to the well understood Purdue Enterprise Reference Architecture developed in the 1990s. The Purdue model structures an industrial enterprise into a series of layers ranging from the physical process (Level 0), basic control (Level 1), area control (Level 2), site manufacturing operations and control (Level 3), and business planning and logistics (Level 4). Enterprise wide business systems, such as ERP systems are often considered as Level 5 of the Purdue model.

This abstract model typically has a corresponding realization in the topology of the system in which boundaries between levels are often expressed as network boundaries across which security can be enforced.

Figure 3 illustrates the basic organization of the Purdue model, including a Level 3.5 DMZ that helps segregate the system in terms of access control and cyber security. The IIoT architecture illustrated previously in Figure 2 is, at the highest level, separated into two major subdivisions – the edge and the cloud. This structure can be further broken down into a seven level model, also shown in

Figure 3.

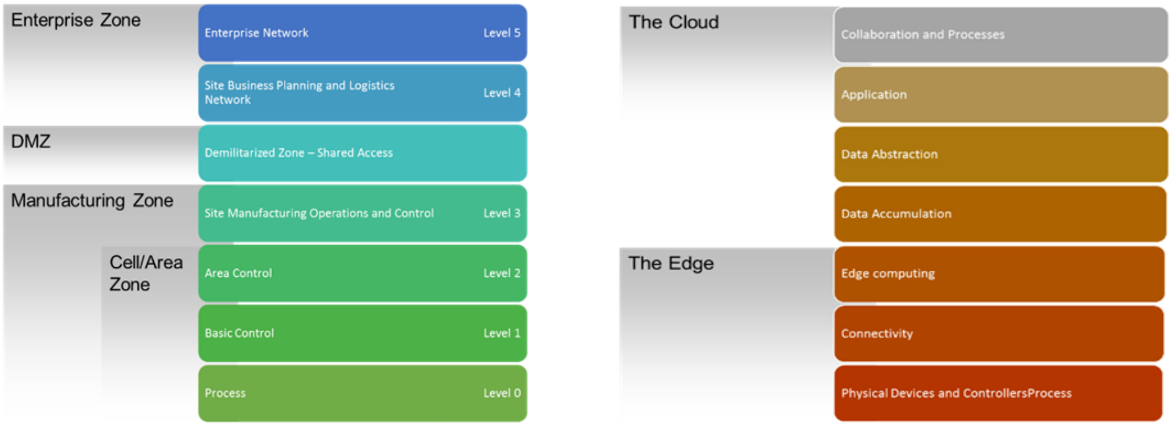

Figure 3 Purdue Enterprise Reference Architecture model on the left and IoT Reference Model on the right

Applying an IIoT architecture to an industrial enterprise requires reconciling these two different organizational structures so that the key architectural qualities provided by the Purdue model (safety, security, reliability, efficiency) are maintained and enhanced within an IIoT-based structure.

An initial reconciliation of the Purdue model with the IoT model considers the partitioning of the functionality represented by the four main layers in the Purdue model within the two main layers in the IoT model – the edge and the cloud. Level 1 of the Purdue model, basic control, moves to the edge in the IoT model, while Level 4, business planning and logistics, moves to the cloud. There is also a strong argument for moving much of Level 2, area control, to the edge to keep it close to the process being controlled for performance, security, and reliability reasons. The functionality represented by Level 3, site manufacturing operations will get pulled up into the cloud and pushed down to the edge depending on the balance of key system quality attributes. History, Advanced Process Control, S88 Batch, and Alarm Management are all examples of functions that can be deployed either in the cloud, or on premise in embedded devices, or both.

An approximate allocation of Purdue model levels to the basic IoT partitioning is illustrated in

Figure 4. In general, moving functionality to either the cloud or the edge represents a tradeoff amongst a number of system qualities. For example, moving functionality to the edge can improve performance and reliability at the expense of having to provision and manage functionality distributed across a large number of devices. On the other hand, moving functionality to the cloud makes it easier to install, scale up, upgrade, and retire at the expense of the functionality being remote from the devices and controllers on which the functionality may depend.

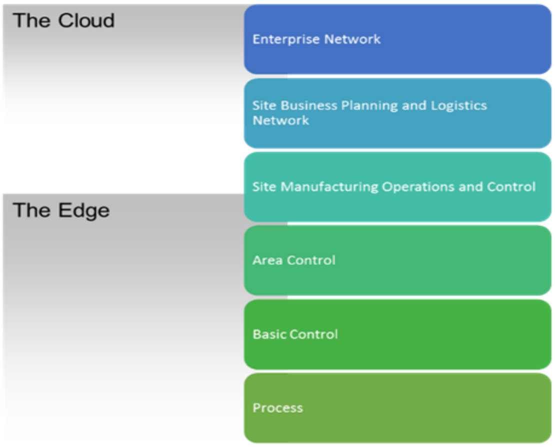
The Cloud
The Edge
Enterprise Network
Site Business Planning and Logistics Network
Site Manufacturing Operations and Control
Area Control
Basic Control
Process

Figure 4 Approximate correspondence between levels in the Purdue model and the basic structure of the IoT

In general the move to an IIoT-based architecture will result in a system unconstrained by the hierarchical structure of a DCS.

IIoT Benefits

This section will describe how IIoT can deliver better support for key requirements in the areas of safety, security, reliability, and efficiency.

Safety

The overriding concern in any industrial enterprise is safety. The operation of industrial processes must be done in a way that protects human workers from harm and the surrounding environment from accidental release of material into the atmosphere, water ways, etc. Safety issues are most critical in processes that involve hazardous materials and high temperatures and pressures for which the nature and magnitude of accidents are the most severe.

There are well-understood and well-developed sets of practices and standards concerning the basic safety of an industrial process. For example, the Safety Integrity Level model provides a quantitative measure of the risk reduction provided by Safety Instrumented Systems that are responsible for the basic safety of a process and formalized in IEC 61511. These practices are embodied in products such has Honeywell's Safety Manager that provides a SIL-certified safety platform based on a Quadruple Modular Redundant architecture and the Universal Safety Logic Solver that provide very high levels of reliability even in the face of component failures. There will continue to be a key role for Safety Instrumented Systems in any IIoT based automation system.

Security

A concern closely related to safety is that of security. Unless an automation system is secure from unauthorized access and activity its safety cannot be guaranteed. Aside from preventing compromises to the safety of the plant, security also serves to protect the intellectual property inherent in an industrial process itself and the procedures for planning, scheduling, executing, maintaining, and optimizing production on the process. Increasing levels of computer-based automation have increased the risks associated with cyber security attacks and has led to the development of cyber security standards and practices such as ISA/EC-62443 (formerly ISA-99).

Many existing DCS components have no inherent security built in. For example, they may lack any explicit access control mechanism and may transmit data on the network in plain text. The legacy components do not disappear in an IIoT-based system, but are confined to the edge computing environment to which access is strictly controlled. Access to legacy DCS components via edge gateway devices involves both access control and secure communications.

Another vulnerability in current automation systems stems from the use of open systems platforms, particularly in Levels 2 and higher in the Purdue model. These platforms pose risks due to their widespread use across many domains making their vulnerabilities and associated exploits well understood. Patching of these platforms in order to address vulnerabilities as they are discovered is a major cost in the maintenance of such systems and even if performed promptly do not protect systems from so-called zero day attacks for which patch is not available. The open system platforms also provide numerous attack vectors including standards-based interfaces such as USB ports.

IIoT helps address these issues by pushing automation system functionality either down into the hardened edge computing environment or up into the cloud. The cloud computing environment has rich access control and communications security mechanisms built-in and the centralized nature of the infrastructure makes it much easier to maintain in order to address vulnerabilities that are discovered.

Reliability

Reliability refers to the ability of a system to remain operational over time. The probability that a highly reliable system will fail to perform its intended function is very small and is a key characteristic of industrial automation systems. A reliable system is able to continue to perform its intended function in the face of a range of changing circumstances that include, but are not limited to equipment failure, maintenance and upgrade of components, plant expansion, and changes in the circumstances of those working in the facility. An IIoT-based approach can contribute to the reliability of an industrial enterprise both in terms of the reliability of the automation system itself as well as the reliability of the production process more generally.

The reliability of the automation system can be enhanced both by pushing functions out to the edge and into the cloud. As with safety, pushing functions, especially control functions, out to the edge allows those functions to act more autonomously with fewer dependencies on other components, reducing the potential causes of failure. Moving functions into the cloud allows them to be more easily managed, maintained and upgraded reducing the impact of these operations on the system. The decoupling of edge and cloud-based functions also allows them to be managed much more independently, again allowing the system to remain operational through a much wider range of lifecycle events. Cloud-based deployment also makes it very easy to expand a system with additional storage and computational resources in order to facilitate incremental plant expansions without the engineering overhead associated with a current DCS.

Efficiency

With a production process that is running safely, securely, and reliably, attention can turn to making production as efficient as possible in order to maximize the profitability of the enterprise. This amounts to optimizing operations in a range of areas such as maximizing throughput or yield, minimizing energy and raw material usage, minimizing engineering, maintenance and labor costs, and so on.

Optimization is essentially a decision making process directed toward achieving a specified goal subject to constraints that limit the actions that can be taken to achieve that goal. The decision making process may occur second to second as in the case of online optimization of process control variables or day to day as in the case of production and maintenance planning. The ability to make the right decisions depends on having the relevant information available on which to base the decisions and a means of determining the course of action to take on the basis of that information. The IIoT approach provides important opportunities for improving decision making by delivering the right information to the places in forms it can be acted on.

The ability to collect more data from uncorrelated sources provides opportunities for applying data analytics, modeling, and machine learning techniques to gain better insight into the current and future state of the enterprise. This information can then be delivered to those in decision making roles in ways that allow the decision makers to act on that information. Decision making can be decentralized and put in the hands of those responsible for the enacting the decision.

Analytics, modeling, simulation, and machine learning techniques also provide additional opportunities for closing the loop and enacting decisions automatically. In these cases, the decision making process can be pushed out to the edge environment to enhance the capabilities of the autonomous elements of the system. Moving model-based control into the real-time control platform means that on premise elements of the system are able to continue to optimize the process rather than just regulate it without depending on non-local resources.

Standards

A significant difference between today's DCS and an IIoT system has to do with heterogeneity; whereas a DCS tends to be a combination of a vendor's proprietary technology, the IIoT must accommodate fine-grained use of technologies and functions from multiple vendors, and to do so over a long time horizon. To do this, new standards beyond those allowing for communications interoperability (HART, Foundation Fieldbus, ProfiBus, OPC, etc.) will need to emerge to allow for functional alignment from multiple sources. OPC UA (OPC Unified Architecture) is one such standard; it has protocol support as well as the semantic richness of a well-defined namespace that encourages collaboration across many different application classes (history, alarms & events, control, security, etc.). Although OPC UA won't be the only standard employed in IIoT, it has the potential to be the lingua franca of interoperable and well-bred IIoT solutions.

Getting There from Here

The IIoT represents a step change in the evolution of automation systems. The benefits that flow from new, highly scalable deployment patterns, smarter devices, more comprehensive data collection and analytics, and broader reach through mobile applications are large. However, achieving these benefits requires an orderly transition from the automation system of today to the automation system of the future. This transition will be a step-wise initiative that will need to consider the following key aspects:

Preservation of core customer IP – Customers typically have very large engineering and intellectual property investments in their automation systems. Control strategies, supervisory applications, and HMI graphics need to be preserved as the automation system evolves. Re-engineering these is expensive and adds little value. It is far better to preserve this investment either by providing ongoing support for these items in their current form or by providing high fidelity translation to new forms.

Preserving in-place equipment, especially large capital equipment, and DCS/PLC solutions – In addition to the engineering content of an automation system there is a lot of associated equipment. Ripping this equipment out and replacing it with new equipment is usually not feasible or cost effective, so it is imperative that evolution to the IIoT accommodates existing equipment. A key strategy here is to provide support for existing communications protocols that allows existing equipment to be integrated into and IIoT architecture in a secure way.

Maintaining SIL levels – The Safety Integrity Levels of equipment and systems in an automation system are central to it achieving its primary goal of safe operations. Any move to new deployment patterns and new devices needs to maintain existing SIL levels. Of course, the same applies to maintaining the security of a system. In both cases, the evolution of the system should be seen as an opportunity to not only maintain levels of safety and security, but to enhance them beyond their current levels.

On-process updates to control hardware and software solutions – A step-wise evolution to new forms of automation systems will occur over a period of time. As changes to a system are introduced, they need to be done in a way that does not interrupt or compromise plant production. The updating of hardware and software as well as the introduction of new system components needs to be done “on-process”.

Performance/capacity of existing systems and demand placed on them from cloud based applications – The IIoT encourages the collection of more data from more sources. While this provides a platform for improved analytics, more advanced planning and optimization applications, and more powerful mobile solutions, the impact of this additional demand for data on the existing components of an automation system needs to be managed. There is little point in enabling new forms of application if the needs of those applications compromise the core mission of the automation system.

The good news is that Honeywell has a long history of exactly this sort of system evolution. The evolution of TDC 2000 to TDC 3000 and on to Experion PKS demonstrates Honeywell’s ability to institute significant architectural change in automation systems while honoring the key principles outlined above. This evolution continues as Experion evolves towards the IIoT.

Conclusion

In many ways, the IIoT represents an “undiscovered country”, full of promise, but waiting to be explored and mapped out. This article has attempted to map out this undiscovered country and provide pointers to how the promise of future automation systems will be realized. The resulting vision is a new form of automation system architecture that balances the computational and lifecycle benefits of Cloud Computing with the requisite on premise, Appliance-hosted capabilities necessary to provide safe, secure and long-lasting automation for complex manufacturing systems and processes.

About the Authors

Paul McLaughlin is Chief Engineer in Honeywell Process Solutions. He has worked in the automation industry for 35 years and been with Honeywell for the past 30. During that time Paul was responsible for system and product architecture across the range of HPS products including Experion PKS, TotalPlant Solutions System, Intuition, Engineered Field Solutions, and System Migration. Paul is currently responsible for the development of the HPS Industrial Internet of Things strategy. Paul has degrees in mathematics and computer science from the University of Delaware and the University of Pennsylvania.

Rohan McAdam is Chief Architect for Honeywell Process Solutions. He has worked in the field of industrial process control since 1988. Prior to joining Honeywell in 1993, he worked in the alumina industry in Western Australia. Rohan has worked primarily on the development of Human Machine Interfaces for HPS industrial process control products. Rohan has a mathematics degree from Charles Sturt University, a master's degree in cognitive science from the University of New South Wales, and a PhD in computer science from Charles Sturt University.

For More Information
Learn more about how Honeywell is delivering performance, reliability, safety and security, visit our website www.honeywellprocess.com or contact your Honeywell account manager.

Honeywell Process Solutions
Honeywell
1250 West Sam Houston Parkway South
Houston, TX 77042

Honeywell House, Arlington Business Park
Bracknell, Berkshire, England RG12 1EB UK

Shanghai City Centre, 100 Junyi Road
Shanghai, China 20051

www.honeywellprocess.com

Honeywell

WP-15-19-ENG
January 2016