



US 20140107810A1

(19) **United States**(12) **Patent Application Publication**
Johnson et al.(10) **Pub. No.: US 2014/0107810 A1**(43) **Pub. Date: Apr. 17, 2014**(54) **METHODS AND APPARATUS FOR CONTROL USING CONTROL DEVICES THAT PROVIDE A VIRTUAL MACHINE ENVIRONMENT AND THAT COMMUNICATE VIA AN IP NETWORK**(71) Applicant: **Invensys Systems, Inc.**, Foxboro, MA (US)(72) Inventors: **Alexander Johnson**, Houston, TX (US); **Paul C. Badavas**, Southboro, MA (US); **T. Eric Christiansen**, Raynham, MA (US); **Peter D. Hansen**, Wellesley, MA (US); **Thomas B. Kinney**, Franklin, MA (US); **Seyamak Keyghobad**, Marion, MA (US); **Bo Ling**, Norfolk, MA (US); **Richard L. Thibault**, Virginia Beach, VA (US)(21) Appl. No.: **14/054,037**(22) Filed: **Oct. 15, 2013****Related U.S. Application Data**


(60) Continuation of application No. 12/334,191, filed on Dec. 12, 2008, now abandoned, which is a continuation of application No. 11/781,219, filed on Jul. 20, 2007, now Pat. No. 8,090,452, which is a division of application No. 11/354,586, filed on Feb. 14, 2006, now abandoned, which is a continuation of application No. 10/756,187, filed on Jan. 13, 2004, now Pat. No.

7,020,532, which is a continuation of application No. 09/591,604, filed on Jun. 9, 2000, now Pat. No. 6,788, 980.

(60) Provisional application No. 60/149,276, filed on Aug. 17, 1999, provisional application No. 60/144,693, filed on Jul. 20, 1999, provisional application No. 60/139,071, filed on Jun. 11, 1999.

Publication Classification(51) **Int. Cl.**
G05B 15/02 (2006.01)(52) **U.S. Cl.**
CPC **G05B 15/02** (2013.01)
USPC **700/19**(57) **ABSTRACT**

The invention provides improved methods and apparatus for control using field and control devices that provide a virtual machine environment and that communicate via an IP network. By way of non-limiting example, such field device can be an "intelligent" transmitter or actuator that includes a low power processor, along with a random access memory, a read-only memory, FlashRAM, and a sensor interface. The processor can execute a real-time operating system, as well as a Java virtual machine (JVM). Java byte code executes in the JVM to configure the field device to perform typical process control functions, e.g., for proportional integral derivative (PID) control and signal conditioning. Control networks can include a plurality of such field and control devices interconnected by an IP network, such as an Ethernet.






Fig. 1

FIG. 2

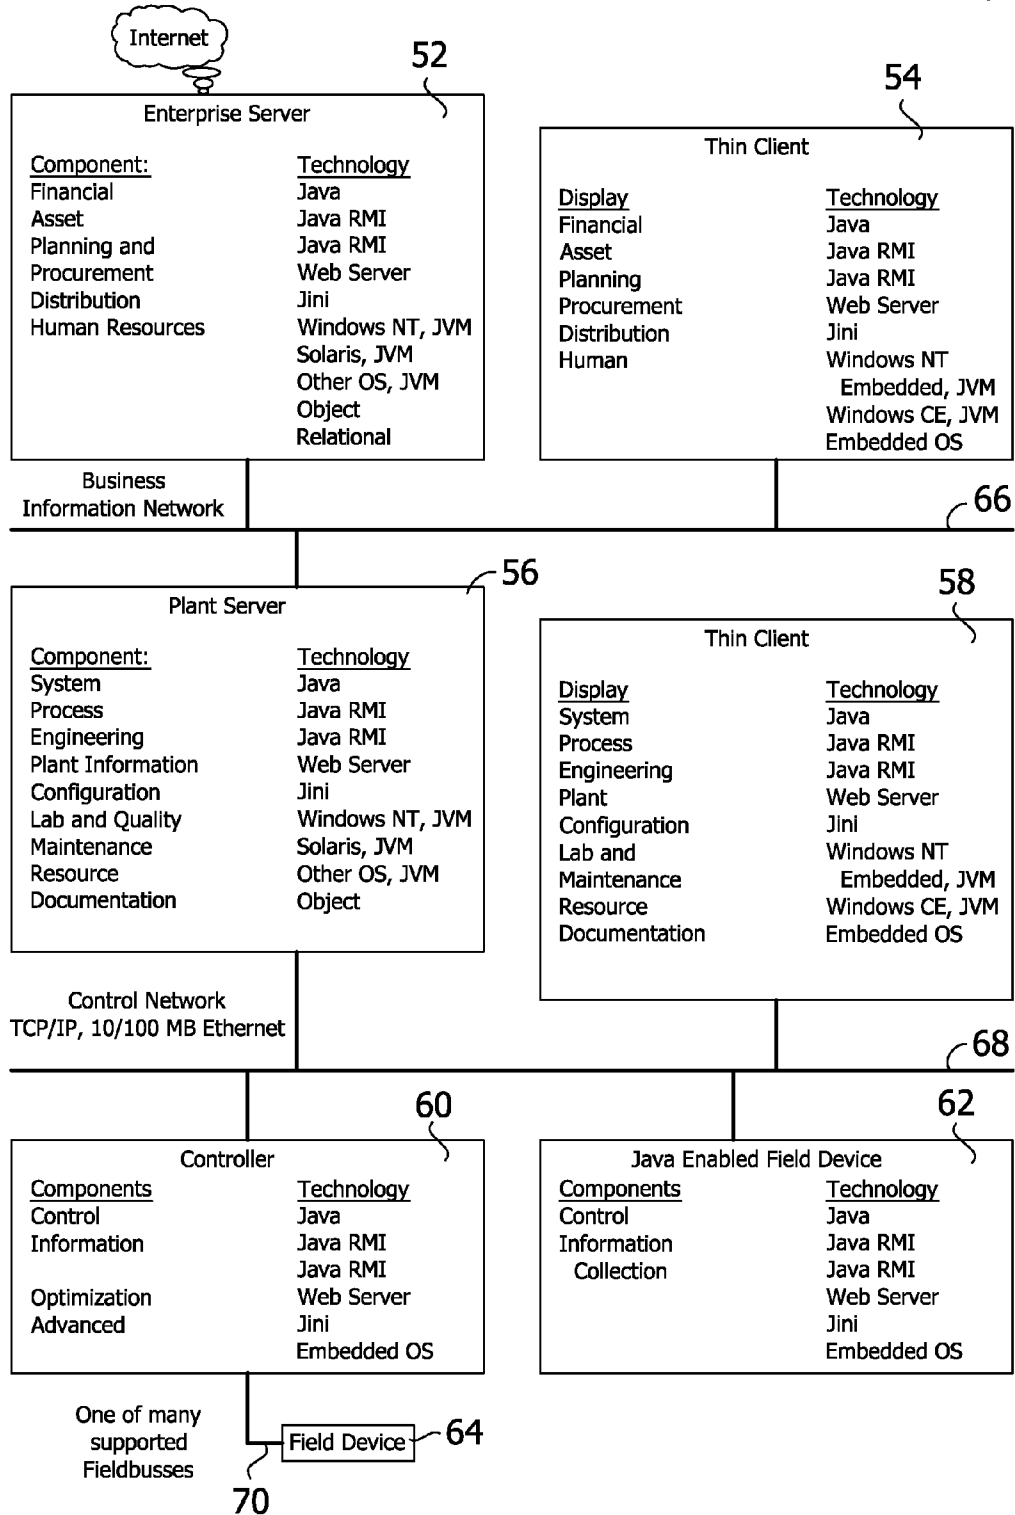


FIG. 3


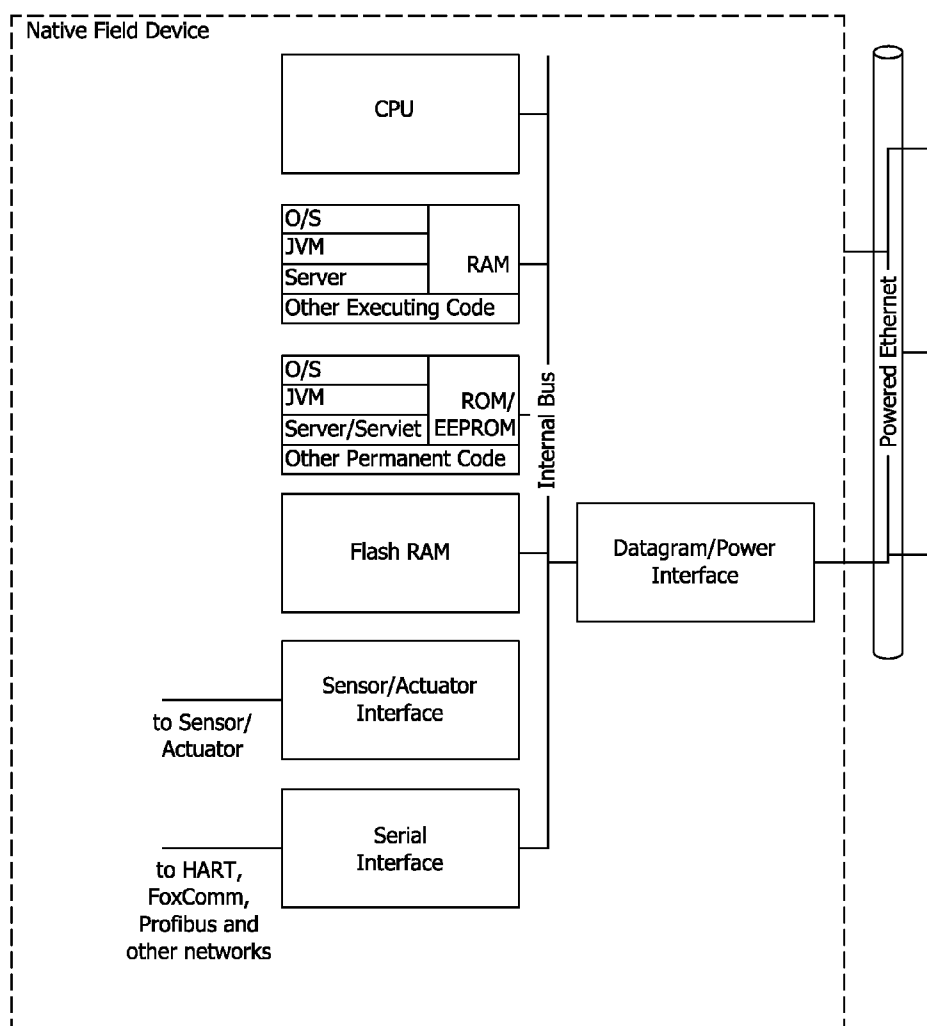


FIG. 4






Fig. 5

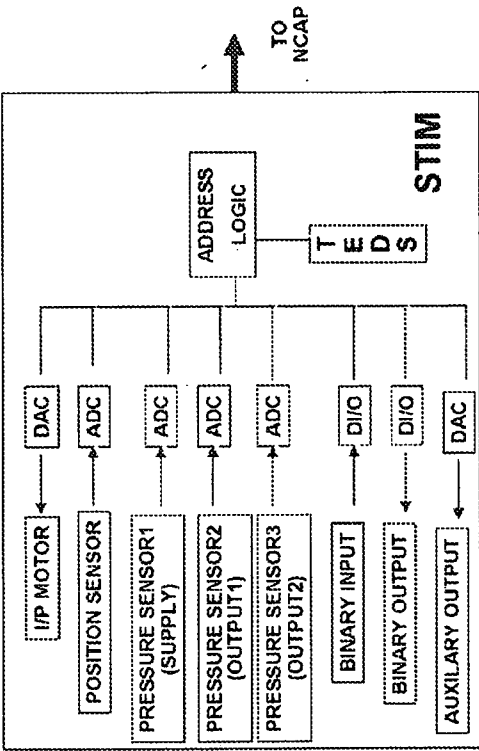


Fig. 6

**METHODS AND APPARATUS FOR CONTROL
USING CONTROL DEVICES THAT PROVIDE
A VIRTUAL MACHINE ENVIRONMENT AND
THAT COMMUNICATE VIA AN IP
NETWORK**

[0001] This application claims the priority of the following United States patent applications: U.S. Patent Application Ser. No. 60/139,071, entitled OMNIBUS AND WEB CONTROL, filed Jun. 11, 1999; U.S. Patent Application Ser. No. 60/144,693, entitled OMNIBUS AND WEB CONTROL, filed Jul. 20, 1999; U.S. Patent Application Ser. No. 60/149,276, entitled METHODS AND APPARATUS FOR PROCESS CONTROL ("AUTOARCHITECTURE"), filed Aug. 17, 1999; U.S. patent application Ser. No. 09/345,215, entitled PROCESS CONTROL SYSTEM AND METHOD WITH IMPROVED DISTRIBUTION, INSTALLATION, AND VALIDATION OF COMPONENTS, filed Jun. 30, 1999.

BACKGROUND OF THE INVENTION

[0002] The invention pertains to control systems and, more particularly, to methods and apparatus for networking, configuring and operating field devices, controllers, consoles and other control devices.

[0003] The terms "control" and "control systems" refer to the control of a device or system by monitoring one or more of its characteristics. This is used to insure that output, processing, quality and/or efficiency remain within desired parameters over the course of time. In many control systems, digital data processing or other automated apparatus monitor a device, process or system and automatically adjust its operational parameters. In other control systems, such apparatus monitor the device, process or system and display alarms or other indicia of its characteristics, leaving responsibility for adjustment to the operator.

[0004] Control is used in a number of fields. Process control, for example, is employed in the manufacturing sector for process, repetitive and discrete manufactures, though, it also has wide application in utility and other service industries. Environmental control finds application in residential, commercial, institutional and industrial settings, where temperature and other environmental factors must be properly maintained. Control is also used in articles of manufacture, from toasters to aircraft, to monitor and control device operation.

[0005] Modern day control systems typically include a combination of field devices, controllers, workstations and other more powerful digital data processing apparatus, the functions of which may overlap or be combined. Field devices include temperature, flow and other sensors that measure characteristics of the subject device, process or system. They also include valves and other actuators that mechanically, electrically, magnetically, or otherwise effect the desired control.

[0006] Controllers generate settings for the control devices based on measurements from sensor type field devices. Controller operation is typically based on a "control algorithm" that maintains a controlled system at a desired level, or drives it to that level, by minimizing differences between the values measured by the sensors and, for example, a setpoint defined by the operator.

[0007] Workstations, control stations and the like are typically used to configure and monitor the process as a whole. They are often also used to execute higher-levels of process

control, e.g., coordinating groups of controllers and responding to alarm conditions occurring within them.

[0008] In a food processing plant, for example, a workstation coordinates controllers that actuate conveyors, valves, and the like, to transport soup stock and other ingredients to a processing vessel. The workstation also configures and monitors the controllers that maintain the contents of that vessel at a simmer or low boil. The latter operate, for example, by comparing measurements of vapor pressure in the processing vessel with a desired setpoint. If the vessel pressure is too low, the control algorithm may call for incrementally opening the heating gas valves, thereby, driving the pressure and boiling activity upwards. As the pressure approaches the desired setpoint, the algorithm requires incrementally leveling the valves to maintain the roil of the boil.

[0009] The field devices, controllers, workstations and other components that make up a process control system typically communicate over heterogeneous media. Field devices connect with controllers, for example, over dedicated "fieldbuses" operating under proprietary or industry-specific protocols. Examples of these are FoxCom™, Profibus, ControlNet, ModBus, DeviceNet, among others. The controllers themselves may be connected to one another, as well as to workstations, via backplane or other proprietary high-speed dedicated buses, such as Nodebus™. Communications among workstations and plant or enterprise-level processors may be via Ethernet networks or other Internet Protocol (IP) networks.

[0010] Control device manufacturers, individually, and the control industry, as a whole, have pushed for some uniformity among otherwise competing communication standards. The Foundation Fieldbus, for example, is the result of an industry-wide effort to define a uniform protocol for communications among processor-equipped (or "intelligent") field devices. Efforts such as this have been limited to specific segments of the control hierarchy (e.g., bus communications among field devices) and are typically hampered by technological changes that all to soon render the standards obsolete.

[0011] Still less uniform are the command and operation of control devices. Though field devices may function at the direction of controllers and controllers, in turn, at the direction of workstations (or other plant-level processors), proprietary mechanisms within the individual components determine how they perform their respective functions. Even the commands for invoking those functions may be manufacturer- or product-specific. Thus, the commands necessary to drive actuators of one manufacturer will differ from those of another. How the corresponding commands are processed internally within the actuators differ still more (though, hopefully, the results achieved are the same). The specific programming codes used to effect a given control algorithm likewise differs among competing makes, as do those of the higher-level control processors.

[0012] Industry efforts toward harmonization of software for command and operation of control devices have focused on editing languages that define process control algorithms. This is distinct from the codes used to effect those algorithms within control devices and, rather, concerns software "tools" available to users to specify the algorithms, e.g., editors including IEC-1131 standard languages such as Field Blocks, Sequential Function Charts (SFC), Ladder Logic and Structured Text.

[0013] Less concerted are industry moves to extend monitoring and limited configuration capabilities beyond in-plant

consoles, e.g., to remote workstations. An example of this was the abortive Java for Distributed Control (JDC) effort, which proposed enabling in-plant workstations to serve web pages to remote Java bytecode-enabled client computers. The latter used the to web pages to monitor and set control parameters which the workstations, in turn, incorporated into their own control schemes.

[0014] An academic system along these same lines was suggested by the Mercury Project of the University of Southern California, proposing the use of a web browser to enable remote users to control a robotic arm via a server that controlled the arm. A related company-specific effort included that announced by Tribe Computer Works that allegedly enabled users to manage routers and remote access servers over IP networks using web browser software. See, "Tribe Defines Net Management Role For Web Browser Software," *Network World*, May 22, 1995, at p. 14.

[0015] Thus sets the stage for the present invention, an object of which is to provide improved methods and apparatus for networking, configuring and operating field devices, controllers, consoles and other control devices. A related object is to provide such methods and apparatus for process control.

[0016] Further objects of the invention are to provide such methods and apparatus as reduce the confusion, complexity and costs attendant to prior art control systems.

[0017] Related objects of the invention are to provide such methods and apparatus as can be implemented with commercial off the shelf hardware and software.

[0018] Still further objects of the invention are to provide such methods and apparatus as achieve confusion-, complexity- and cost-reduction without hampering manufacturer creativity and without removing incentives to development of product differentiators.

SUMMARY OF THE INVENTION

[0019] The foregoing are among the objects attained by invention which provides, in one aspect, an improved field device for a process or other control system. The field device includes a virtual machine environment for executing Java byte code (or other such intermediate code) that, for example, configures the device to execute a control algorithm.

[0020] By way of non-limiting example, the field device can be an "intelligent" transmitter or actuator that includes a low power processor, along with a random access memory, a read-only memory, FlashRAM, and a sensor interface. The processor can execute a real-time operating system, as well as a Java virtual machine (JVM). Java byte code executes in the JVM to configure the field device to perform typical process control functions, e.g., for proportional integral derivative (PID) control and signal conditioning.

[0021] Further aspects of the invention provide a field device, such as a low-power intelligent actuator, that incorporates an embedded web server. This can be used to configure, monitor and/or maintain the device itself (as well as other elements of the control system) via a browser attached directly to the device or coupled to it over the network. To this end, the field device can incorporate a configuration editor, e.g., operating on a processor within the field device, that an end-user executes via the browser and web server.

[0022] Such a configuration editor can, in related aspects of the invention, be enabled or disabled depending on the environment in which it is used and more specifically, for example, the type of network in which it is incorporated.

Thus, for example, the editor can be disabled when the field device is incorporated in a process control network that includes, e.g., an applications development environment suitable for configuration of the field device. Conversely, it can be enabled when the field device is incorporated in a network that lacks such a capability.

[0023] Still further aspects of the invention provide a field device as described above that includes an interface to an IP network, through which the device communicates with other elements of the control system. The IP network can be, for example, an Ethernet network. Moreover, it can be "powered," carrying electrical power as well as packets, datagrams, and other control or data signals. The field device, in related aspects of the invention, draws operational power, e.g., for its processor and other components, from such a network.

[0024] Yet further aspects of the invention provide a field device as described above that obtains configuration information and/or its network address from such an IP network upon start-up. To this end, on power-up or coupling to the network, the field device can supply an identifier (e.g., attained from a letterbug, assigned by a hub, or otherwise) to a DHCP or other server on the network. Once provided with an IP address, the field device can formally enter into the control network, e.g., by posting its characteristics to a network bulletin board, e.g., using a network enabler such as a Jini and/or JavaSpace server, or the like. Other network devices monitoring or notified via such a bulletin board can send configuration information to the field device or otherwise.

[0025] Still further aspects of the invention provide control devices, such as servers, control stations, operator consoles, personal computers, handheld computers, and the like, having attributes as described above. Such control devices can have other attributes, according to further aspects of the invention. Thus, by way of non-limiting example, they can provide web servers that collect process data from one or more control devices, generate source for operator displays, provide access to the control system, and host an applications development environment.

[0026] Still other aspects of the invention provide process, environmental, industrial and other control systems that comprise field and control devices as described above that are coupled via an IP network and, particularly, for example, by a powered IP network.

[0027] Additional aspects of the invention are directed to DHCP servers and network enablers (optionally, including web servers) for use in control systems as described above. Related aspects provide such servers and enablers as are embodied in solid state technologies, e.g., with no moving parts.

[0028] These and other aspects of the invention are evident in the attached drawings, and in the description and claims that follow.

BRIEF DESCRIPTION OF THE DRAWINGS

[0029] A more complete understanding of the invention may be attained by reference to the drawings, in which:

[0030] FIG. 1 depicts a process control system 10 according to one practice of the invention;

[0031] FIGS. 2 and 3 depict more particular embodiments of a control system of the type shown in FIG. 1;

[0032] FIG. 4 depicts a native intelligent field device according to the invention;

[0033] FIG. 5 depicts a native control device according to the invention; and

[0034] FIG. 6 depicts a native intelligent positioner implementation according to the invention.

DETAILED DESCRIPTION OF THE ILLUSTRATED EMBODIMENT

[0035] FIG. 1 depicts a process control system 10 according to the invention. The system includes networked control devices that monitor and control a hypothetical mixing process that utilizes mixing chamber 22, fluid inlets 24, 26, fluid outlet 28, paddle 30, cooler 32, and cooler inlet 34. Though illustrated and described below for use in connection with process control, those skilled in the art will appreciate that apparatus and methods according to the invention can be used in connection any industrial, manufacturing, service, environmental or other process, device or system amenable to monitoring or control (hereinafter, collectively, "control").

[0036] The networked control devices include actuators, such as the valves depicted as controlling inlets and outlets 24-28 and 34. A further actuator is shown controlling paddle 30. These and other actuators utilized by the control system are constructed and operated in the conventional manner, as modified in accord with the teachings hereof. The actuators operate under control of respective field device controllers, labeled CTL, that are also constructed and operated in the conventional manner to provide initialization, signal conditioning and communications functions.

[0037] Rather than using separate controllers CTL, the actuators can be of the intelligent variety and can include integral microprocessors or other digital data processing apparatus for control, initialization, signal conditioning, communications and other control-related functions. For sake of convenience, the label CTL is used regardless of whether the control-related functionality is integral to the actuators (e.g., as in the case of intelligent actuators) or otherwise.

[0038] Illustrated sensor 29 monitors a temperature, level or other characteristic of fluid in chamber 22. The sensor 29, as well as other sensing apparatus utilized by the system, are constructed and operated in the conventional manner known in the art, as modified in accord with the teachings hereof. They can be coupled to the control network via a transmitter or other interface device INT that, too, is constructed and operated in the conventional manner, as modified by the teachings hereof. The interface devices facilitate initialization, signal conditioning and communications between the sensors and the control system. As above, one or more sensors can be of the intelligent variety, incorporating integral microprocessors or other digital data processing capabilities for initialization, signal conditioning, communications and other control-related functions. Here, too, the label INT is used in reference to the control-related functionality, regardless of whether embodied in an intelligent transmitter or otherwise.

[0039] The networked control devices include one or more controllers 36 that monitor and control respective aspects of the hypothetical mixing process in the conventional manner, as modified in accord with the teachings hereof. The controllers can comprise mainframe computers, workstations, personal computers, special-purpose hardware or other digital data processing apparatus capable of performing conventional monitoring and control functions. Preferred controllers are constructed and operated in the manner of the CP control processors commercially available from the assignee hereof, as modified in accord with the teachings herein.

[0040] The control system 10 includes a variety of devices that serve as user interfaces and that provide configuration and/or control functions, all in the conventional manner as modified in accord with the teachings hereof. Illustrated for these purposes are workstation 40, laptop computer 42 and handheld computer 44. These devices can provide configuration and control functions directly, as in the case of workstation 40, or in cooperation with server devices, e.g., as in the case of handheld computer 44 and server 46. Apparatus 40-44 can couple with the control network directly, e.g., via bus or network connection, or indirectly, e.g., via satellite, wireless connection or modem connection.

[0041] The control devices 36-46, CTL and INT, collectively, referred to as "native" devices, are coupled for communications via a medium that permits at least selected ones of the devices to communicate with one another. To this end, in the illustrated embodiment those devices are coupled via one or more networks 48 that are, preferably, IP-based such as, by way non-limiting example, Ethernets. The network(s) can include, as indicated by the multiple segments shown in the drawing, multiple segments such as various wide and local area networks. They may also include high and/or low bandwidth components, such as phone lines, and low and/or high latency components, such as geosynchronous satellites networks.

[0042] FIG. 2 depicts a more particular embodiment of a control system 50 of the type shown in FIG. 1. The system includes an enterprise server 52, a first thin client 54, plant server 56, a second thin client 58, a controller 60, a Java-enabled field device 62, and one or more field devices 64, coupled to one another, e.g., in the manner illustrated, by one or more networks 66, 68.

[0043] Native enterprise server 52 (corresponding, by way of non-limiting example, to server 46) comprises a main-frame computer or engineering workstation that executes enterprise-level applications such as, by non-limiting example, those for financial, asset planning and procurement, distribution and/or human resources. In addition, it supports web serving, as well as optional object, relational and other database management systems. Server 52 executes Windows NT, Solaris or another conventional commercial or proprietary operating system. It is also equipped with a Java Virtual Machine (JVM), e.g., capable of executing Java virtual machine instructions (bytecodes), of performing remote method invocation (RMI), and of supporting Jini networking. The enterprise server 12 can be coupled to further networks, e.g., to the Internet, as shown, in any manner known in the art.

[0044] Native thin client 54 (corresponding, for example, to handheld computer 44) provides similar functionality as server 52, though its actual processing activity is limited to user input and output. Application processing, such as financial, asset planning and procurement, distribution and/or human resources, are performed on behalf of client 54 by a server, such as server 52. The operating system and JVM functions may be embedded in the conventional manner of a thin client. The thin client 54 is coupled to the enterprise server 52 over a business information network 66 (corresponding, for example, to network 48), typically, an Ethernet or other IP network, configured in a LAN, WAN or the like.

[0045] Native plant server 56 (corresponding, by way of non-limiting example to workstation 40) comprises a plant control console or engineering workstation modified in accord with the teachings hereof, executing plant-level control applications including system, process, engineering,

plant information, configuration, lab quality, maintenance, resource and documentation applications of the type known in the art. Like enterprise server **52**, plant server **56** can execute Windows NT, Solaris or other conventional operating systems. It is also preferably equipped to execute a Java Virtual Machine as described above. Plant server **56** is coupled to the enterprise server **52** and thin client **54** over the business information network **66**.

[0046] Native thin client **58** provides similar functionality as server **56** though, again, relies on that (or another) server to perform most processing activity. As above, the operating system and JVM functions may be embedded in the conventional manner of a thin client. The thin client **58** is coupled to the plant server **56** over a control network **68** (corresponding, for example, to network **48**), e.g., an Ethernet.

[0047] Native controller **60** (corresponding, for example, to controller **36**) executes control algorithms to control associated non-native field devices **64**, e.g., via any variety of commercial and/or proprietary field bus **70** hardware and protocols. Where processing resources are limited, the controller **60** utilizes an embedded operating system that supports web serving and the JVM. The controller **60** is coupled to the plant sever **66** and to the thin client **58** via control network **68**.

[0048] Native field device **62** is a sensor, actuator or other field device. The illustrated device is of the intelligent variety, including processor (not shown) and operating system. It can be of the type commercially available in the marketplace, as modified in accord with the teachings hereof. The illustrated device supports web serving and JVM, as described above. The device **62** provides information collection and control functions, as illustrated.

[0049] FIG. **3** depicts another a more particular embodiment of a control system **50** of the type shown in FIG. **1**.

[0050] Referring to FIG. **1**, the illustrated control system **10**, **50** uses web browsers, Java, Jini, JavaSpaces, TCP/IP and other technologies normally associated with the Internet and the world wide web to create a self-defining control network that minimizes complexity while emphasizing the portability and re-usability of the user's application. These technologies are advantageously employed to eliminate proprietary hardware and software to preserve the user's investments; eliminate network configuration and system management through self-configuring networks; increase the user's choice of algorithm suppliers by implementing control in Java; preserve the user's applications through hardware and system software changes through the use of Java and Web Browsers for process displays; minimize the wiring costs; reduce maintenance by making intelligent field devices a practical reality.

[0051] In addition to web technologies, the illustrated control system **10** uses an object location service, a commercial messaging service, and standard networking equipment, such as hubs, routers, network interface cards, where applicable. It also relies on common standards, where applicable, such as 802.3, Internet RFCs, the IEEE 1451 sensor standards. The system **10** supports a heterogeneous computing environment and utilizes industry standards (e.g., Microsoft standards) to provide communications between the native control components to the business and desktop environments.

1 Device Hardware

1.1 Platform-Defining Control Devices

[0052] In the illustrated embodiment, native control devices such as controllers **36**, **60**, workstation **40**, servers **46**,

52, **56** providing a platform for the control system typically include a central processing unit (CPU), memory (RAM), network support hardware, access to permanent storage, an operating system, and a Java Virtual Machine (JVM) including the TCP/IP suite. In addition, the devices include web server software, e.g., software of the type that serves graphical "web" pages in response to requests by other devices. Configurator software can be provided, as well, permitting each device to configure the control system or selected portions of it. Those skilled in the art will appreciate that not all of these components need be included in all native control devices, e.g., some commercially available JVMs can serve as an OS themselves.

[0053] Process control object (PCO) software provided on the platform-defining control devices comprise a collection of data and methods implemented in Java and executed on the native control devices' JVMs to perform typical process control functions, by way of non-limiting example, signal conditioning and PID control. Likewise, station management object (SMO) software on the devices comprise data and methods that allow the device to report its health, performance, and other status information in a uniform manner. The SMO software can implement SNMP or the equivalent Java functionality.

[0054] Software is also provided on the platform-defining controls devices for messaging services for data transfer and alarm/event notification, as well as software comprising system management pages. Each control device may include additional software, of course, depending on its functionality.

1.2 Field Devices

[0055] Native intelligent field devices typically include a low power CPU, e.g., a NetARM or Java Chip; a real-time OS like VxWorks, QNX, PSOS; RAM; FlashRAM to serve as permanent storage; ROM/EEPROM to serve as the home for the OS and other permanent software; an Ethernet interface; power from the Ethernet interface (in the event a powered Ethernet network or hub is used) or otherwise; a sensor interface, e.g., IEEE 1451.1 and 1451.2; JVM; a web server, and a device specific configurator servlet. A field device so constructed can be configured and monitored via a lightweight web browser, e.g., handheld computer **44**, coupled to the device over the network **48**.

[0056] The field devices can include serial interfaces to allow the attachment of these devices to HART, FoxComm, Profibus and other networks operating under a protocol different from that of network **48**. Combined with appropriate software, these devices provide the user with a single transmitter suitable for use on any field network.

[0057] One configuration of a native intelligent field device including the elements described above these elements is shown in FIG. **5**. Those skilled in the art will appreciate that other configurations can be realized, as well, in accord with the teachings hereof.

1.2.1 Stoic Sensors

[0058] The control system **10** can include one or more stoic sensors, not illustrated, that constitute minimal sensing elements. Typically, these are simply silicon chips that are packaged to allow them to sense the process environment and that have only enough electronics to relay measurements to a pair of wires that carry the raw signal to another device for processing. Some stoic sensors generate a signal without external

power; others require some excitation. In a preferred embodiment, native devices **36-46**, CTL and INT according to the invention support the attachment of many stoic sensors.

1.2.2 I/O

1.2.2.1 Overview

[0059] Some I/O devices, e.g., thermocouples, RTDs and analog transmitters, may not use a transmitter compatible with the network **48** but, rather, send raw sensor output to a multiplexor for processing. To accommodate these devices, the system **10** includes two types of intelligent I/O cards: network I/O cards supporting IP and an API, and native I/O cards that support the protocol of network **48** (i.e., the “native” protocol) directly.

[0060] Network I/O cards are coupled directly to network **48**. Native control devices are IP enabled and support the proprietary API of the cards, thereby, permitting reading and writing of their I/O registers. The cards’ respective APIs support retrieval of data in several formats (raw counts, linearized counts, engineering units, etc.), as well as the assignment of simple configuration information to those registers, if the device supports such functions. To this end, the native devices utilize I/O classes that corresponding to the native I/O protocols. Redundant I/O devices can be physically interfaced transparently through a single I/O card or through multiple independent network I/O cards. Logically, PCOs support at least the use of multiple independent cards. Alternative I/O features may be used in addition. Native I/O cards are substantially similar to native-enabled transmitters, except for packaging and scale. These I/O cards may be considered as small multi-loop controllers or multiplexors.

1.2.2.2 Foreign Device Integration

[0061] Networked I/O cards are utilized with devices that cannot otherwise directly interface with the network **48**. The cards, in this case, provide an interface that permits at least reading and writing of relevant device values, which can be stored internally to the card and accessed via its API. Foreign devices that are control systems in its own right typically maintain objects or other data structures representing process values and associated data. A preferred interface to these “devices” are through objects that wrap the foreign function blocks with the services expected of native devices, e.g., detail displays and configurators, which are accessed in exactly the same manner as native ones of those same services. For example, a name service is able to locate foreign blocks and the native API of system **10** permits querying those blocks for values.

1.3 Native Controllers

[0062] In a preferred embodiment, all native devices may be used as controllers, though, dedicated I/O-less native controllers, such as controllers **36** and **60**, are typically required for unit-wide operations. Thus, for example, control can also be provided by personal computers, workstations and servers of the type shown as elements **40**, **42**, **46**, in FIG. 1. Native controllers, e.g., **36**, **60** are of the same general design as native boards used in the transmitters and actuators. However, they support a faster CPU clock cycle and more memory (volatile and non-volatile) than their smaller siblings. High performance native controllers have a still more powerful CPU, more RAM, and maybe a removable FlashRAM card

for bulk storage and backup. Like the native transmitters and actuators, these native controllers can receive power from a powered Ethernet connection or otherwise. Native controllers can be used in redundant and non-redundant configurations. Since the I/O is independent of the native controller, redundancy can be implemented in a number of manners, e.g., hardware fault-tolerance or transaction based synchronization between stations with clustering.

1.4 Native Compliant Workstations

[0063] Native workstations, e.g., such as element **40**, can be constitute web browser-enabled devices that communicate with web servers, such as element **46**, that actually collect the process data from other native devices. In addition, the workstations can consist of a flat panel display, OS and web browser in ROM (or on a memory card), web page (process graphic) database/cache, connections for optional annunciator keyboards, option for wireless Ethernet, and, optional, battery operation. Data supporting the operator interface comes from native devices directly.

1.5 Native Web Server

[0064] If a generic web browser-enabled device is the operator’s console, a native web server, such as servers **46**, **52**, **56** sources the operator displays. This can be implemented with redundancy using technologies NT’s clustering capabilities, or the like. Additionally, a native web server provides access to a illustrated control system **10** to users not physically attached to the illustrated control system **10**, e.g., as illustrated with respect to elements **44** and **46** of FIG. 1. In this mode, it centralizes data requests to maximize efficient use of communication resources.

1.6 Native Enablers

[0065] 1.6.1 The illustrated control system **10** includes native enablers **49**, such as Jini and JavaSpaces devices and Solid-State DHCP servers. These enablers, which are optionally redundant, are preferably fully solid state with persistent memories. They do not use batteries for their persistent memory, though they may well plug into a wall outlet for power in non-industrial environments. The Jini and JavaSpaces serve as community “bulletin boards.” These are used to notify system monitors that native devices have been added to or removed from the system. The system monitors respond to such notices by triggering appropriate actions.

[0066] The illustrated control system **10** assumes that each native device is able to acquire an IP address when it comes on-line. To minimize configuration, a DHCP server, illustrated for example by element **49**, is required to furnish these addresses. To ensure maximum availability of this critical server, it is preferably a solid-state device and, optionally, includes a redundant partner DHCP Server. A particular native DHCP server provides addresses for a portion of the network **48**. It obtains its configuration by notifying the system **10** that it has come on line.

[0067] Those skilled in the art will, of course, appreciate that IP addresses can be selected by mechanisms other than DHCP servers.

1.6.2 Native Internetwork Server

[0068] An internetwork server **47** is used to link separate control systems networks **48**. It provides a platform for internetwork communications, controlled access to data, name

conflict resolution, and other services need to efficiently, securely, and transparently accomplish the connection of one or more illustrated control systems to another.

2 SubSystems

2.1 Software Downloads

[0069] The illustrated system permits the electronic downloading of software and/or licenses for execution on the native devices. This includes native software objects, updates and/or licenses for same. To ensure proper operation of the process facility even in the event of insufficient licenses, downloaded software registers with the system monitor and declares whether or not it is licensed. In the case of insufficient licensing, the system monitor's notifies the customer appropriately.

[0070] In instances where appropriate native software modules and authorizations are in place, the system **10** can access a download site, e.g., an Internet e-commerce site, and determines if updates are available. If so, it notifies the user and invites him/her to request downloading of the upgrade. Such an e-commerce or other web site can also provide configuration tools that allow the user to design, implement, and test a new PCO block. This allows the user to access complex and/or rapidly improving software tools without having to maintain them locally. For this purpose, the use is charged, e.g., a modest access fee.

2.2 Security

[0071] The system **10** includes a native security system to control access to managed objects based on the type of access, e.g., (i) extra-network, i.e., users who are accessing the illustrated control system **10** from a non-native station; (ii) inter-network, i.e., users who are operating from a native workstation on a different physical native network; and (iii) intra-network, i.e., users who are operating from a native workstation within the particular native network.

[0072] Secured extra-system access is provided through a native secure web server, e.g., server **46**, that permits dial-up, network, and other remote access and that supplies and defines the permitted extra-network access, including the API available to applications hosted outside the network or on the server. See, for example, elements **44** and **46** of FIG. **1**. Access is controlled by user name, by user location (IP address or workstation name depending on the network), and by the type of the targeted object.

[0073] Inter-system access is provided by a gateway device, such as server **47**, that permits the secure transfer of data. This device negotiates secure access, deals with name conflicts between systems, and provides support for various physical media. A pair of such devices are provided to account for situations where the source is local to the sink. In a preferred system, the server **47** or other gateway encrypts the data so that others cannot read it. Likewise, it authenticates message sources to verify that they are coming from a matching device. A preferred gateway minimizes the number of packet transfers so as to minimize delays over slow or high latency links.

[0074] Secured intra-system access is controlled based on the user and the workstation, leveraging the Java security model and other security models to the extent possible. The native security system authenticates users, e.g., regardless of whether they are attempting access from operator's console or via an application program.

[0075] The native security system manages access to objects, including PCO attributes (i.e., variables and/or parameters of PCOs), system monitors, and native messages. It allows applications to add objects to the list of managed objects. Read and write access is preferably controlled separately, i.e., as if a single PCO attribute was two separate objects.

[0076] The security model is based on a lock and key approach. Each PCO attribute, for example, is assigned to one of a large number of locks. The user is given different keys for read and write access to object attributes. This key may be modified according to the type of native workstation used for access. The key is passed to the object attribute when the object attribute is accessed. If the access is a connection, it need be passed only once. The object attribute compares the key to its lock and allow or deny access as appropriate.

[0077] A software tool is supplied with the applications development environment (ADE) or configurator that allows identification of users; access points allowed by the user; object attribute groups accessible to the user; and access type (read or write) to the object types. The object attribute groups define collections of object attributes via similar security access guidelines, e.g., alarm limits might be in one group and tuning parameters might be in another. The security model is embedded in the APIs providing access to the secured objects. This assures that that access is granted only after the native security system has performed the necessary verifications.

2.3 Maintenance and Support

[0078] The illustrated control system **10** supports the on-line upgrade of all application software from remote and local workstations. Application software, for purposes of this discussion, includes the system monitor, PCOs, and other Java-based code. Native diagnostic and maintenance tools work over user-supplied or other IP-based networks, providing that they allow services such as the world wide web to operate over them to external sites. For networks that do not permit this, the illustrated control system **10** provides modules that support direct connections to a native support center via dialup analog phone lines, and dialup high speed lines, e.g., ISDN and ADSL. Native maintenance, including software updates, operate over these connections; hence, it is not necessary for a person to be physically present to update system or application software.

[0079] The illustrated control system **10** includes diagnostics to track problems from the hardware upwards. Maintenance staff can run in parallel diagnostic versions of the software. Such software running in open loop or switched into operation as necessary is believed particularly advantageous for support purposes.

2.4 Configuration

[0080] An application development environment (ADE) coordinates configuration in the illustrated control system **10**, e.g., except for devices (such as certain transmitters and positioners) where use of an ADE is not advantageous and for which internal configuration is preferred. Characteristics of the ADE are its use in configuring/building control applications including control schemes, displays, historians, etc.; support of multiple simultaneous users; support of remote, concurrent development; support of change tracking and change management including task based activity traces, summary reports, change annotation, approval cycles, and

user identification; a human interface component that runs in any web browser coupled to the system; a server component that runs on any user-supplied system; permanent storage comprising a commercial database for which a JDBC implementation is available; allowing the definition of configuration object templates; allowing the instantiation of configuration object templates into configuration objects; allowing the user to add and remove editors for configuration object components in real-time; limiting the user's access to configuration capabilities in the various pieces of equipment; and supporting application distribution along with verification of download permissions.

[0081] With respect to the support of multiple simultaneous users, system editors provide object-based "check-out" and "check-in." No other user is allowed to edit a checked-out object or any of the objects that it contains. As an alternative, editors support the concept of task-based configuration activities and the related "check-out" and "check-in" facilities required to make it work. When an object is checked out, it is assigned to a task. Objects are not checked back in, tasks are. A build manager has the responsibility of integrating the tasks.

[0082] The native ADE delivers lower engineering and maintenance costs through the unification of application development, preservation of application expertise, reduction of application development effort, and the deployment of developed applications. Is it based on industry standards (e.g., IEC 1131 and IEC 1499) to preserve the user's investment in training and applications. It also produces appropriately formatted Java class for execution in native devices. Since the native ADE produces output that can be read and loaded into a JVM and since the native control devices include JVM, one control configurator can configure all native devices.

[0083] In a preferred embodiment, the ADE imports configurations from legacy systems; supports the use of third-party control algorithms; and supports both bulk building of control configurations and on-line changes with validation.

[0084] Configuration is not limited to implementation on a web browser, though this can be advantageous since it allows configuration from many types of device without installation of special software, i.e., it provides a thin-client configuration mechanism. For device configuration, there are two cases: a device used in illustrated control system 10 and a device used outside illustrated control system 10. For devices used outside illustrated control system 10, the configurator is placed in the device so that it can be configured even without the native ADE. This on-board configurator allows the device to be configured for use with Profibus, Foundation Fieldbus, on 4-20 ma wires, and other industry-standard or proprietary networks. For devices used in the native environment, (i) a copy of the configurator is available from a native web server for off-line configuration and download, and (ii) the on-board configurator is disabled to prevent changes in the configuration except through the ADE.

2.5 Human Interface

[0085] The illustrated system's primary human interface (HI) device is a web browser. The current range of devices executing these spans cellular phones to mainframe computers. The native HI is multilingual, i.e., it supports the presentation of information in a default language and one or more secondary languages simultaneously. All standard native

applications support text substitution based on the currently selected language. Error messages are likewise in the currently selected language.

2.6 Process Control

[0086] Process Control is implemented using process control objects (PCOs) running in an execution environment consisting of a JVM and any associated applications required to load and execute the specified control strategies. Control strategies are specified in terms of PCO composites. PCOs are Java classes designed to be modular and easily upgraded even during operation. The native PCO configurator, from a high level view, creates instances of these classes, connects them as necessary, and installs them in particular devices.

[0087] Process control objects consist of two user-available types: blocks and composites. PCO blocks are, from the user's point of view, similar to conventional function blocks. Likewise, composites are similar to conventional collections of function blocks. Distinguishing external features of PCOs include the fact that changes involve the addition/deletion of PCOs and cause new versions to be loaded into the targeted native device. The new version runs in parallel with the old version until the engineer decides that it is operating properly and "swaps" the control from old to new.

[0088] Composite PCOs may span stations transparently. PCOs are bound to stations very late in the configuration process and may be migrated from one station to another at any time. Further, PCOs from different sources may operate in the same device if the configurator supports the use of multiple PCO libraries.

[0089] A native PCO configurator supports multiple libraries of PCOs from multiple vendors; permits the creation of composite PCOs from other PCOs; permits the use of composite PCOs as templates, with a composite PCO definition specifying which fields is altered in a template; permits the assignment of PCOs to physical devices; provides IEC 1131/1499 influenced view of the configuration process, i.e., support for Logic Diagrams, Structured Text, Sequential Function Charts, Function Blocks, and PCOs, while producing Java byte code as its output.

[0090] The creation of composites includes the raising of internal (deeply buried) names to the top level of the composite PCO. For example, a cascaded Temperature/Flow loop might have the temperature and flow measurements referenced as TC_CASCADE:PRIMARYFLOW and TC_CASCADE:SECONDARYFLOW or as the longer fully qualified name. The configurator works both off-line in a bulk creation mode and on-line in an individual correction mode.

2.7 Communications: Object Location, Message Transfer, and Data Transfer

[0091] The concepts of object location and data/message transfer are closely bound. A process control system imposes significant quality of service requirements on any services used to locate objects of interest and to acquire values from or make changes to those objects. The illustrated system includes an object location and data transfer service that provides a communication model definition, security, object location services, network types requiring support, quality of service, APIs (Java and non-Java), maintenance and upgrade strategy, and interoperability with existing control systems. In addition, the communications system addresses the particular needs of process data transfer.

2.8 Critical Applications

2.8.1 Time Synchronization

[0092] The illustrated control system **10** supports time synchronization to the millisecond in each station on the network **46**. Where equipment configuration renders this impossible, time synchronization to 50 ms is provided.

2.8.2 Alarm and Message Management

[0093] The illustrated control system **10** provides a facility that centralizes viewing, recording, organizing, and categorizing the messages generated by the system monitor, the operator action journals, the PCOs, and other systems that record textual information. This application is the basis for alarm management strategies including inferential alarming and alarm prediction. However, the message management facility is not in and of itself an alarm historian. Rather, it relies on a native historian to provide the long-term storage of this information.

2.8.3 Historian

[0094] The native historian provides the fundamental data collection, reduction, and archival services. The data collected includes process values and messages from various applications within the illustrated control system **10**. In addition, it provides historical data in support of several other common applications—typically known as plant information management system (DIMS) programs and trend window support. The historian is capable of exporting its data to user-selected databases.

2.8.4 Plant Information Management System

[0095] These applications include a simple and easy to report writing facility, a calculation facility that can use historical and real-time data to compute new values, and a desktop visualization system that uses the historians data instead of real-time data. The desktop visualization system uses the native graphics capabilities to connect the graphics to the historian and to allow the operator to “move” the entire graphic backwards and forwards through the historical data. The values calculated by the calculation facility are stored in PCOs that represent the data and generate appropriate alarms. The report writing facility supports shift reports with calendar adjustments as a simple to use feature. In addition, it facilitates the definition of ad-hoc reports using a page layout metaphor.

3 Interoperability

[0096] Interoperability of the illustrated control system **IO** with legacy process control systems can be facilitated by adding a networked I/O module to the I/O chain of existing I/O cards, adding a native device integrator to talk to other networks of devices, or utilizing a native API for Microsoft and Solaris-based applications.

4 Operating the Control System

4.1 System Startup and Management

4.1.1 Device Identification

[0097] Each native device is assigned a name that is used to identify the device and obtain any configuration information

that it needs. Embodiments of the invention use alternate approaches to identifying a device, including using a hub that is configured by the user to assign a name to each of its ports, e.g., using letterbug or software configuration; installing a letterbug, e.g., on the workshop bench, which tells the device its name; using a PC or handheld computer, e.g., on the workshop bench, to give the device its name; or using soft letterbugs.

4.1.2 Address Acquisition

[0098] After a native device has its name, it acquires its IP address from its environment. The illustrated embodiment uses DHCP for this purpose. Preferably, the DHCP server is a native device, such as enabler **49**, though other DHCP servers can be used, as well.

4.1.3 Registration of the Device

[0099] Once the device knows its address, it registers its characteristics on a native bulletin board, which can be colocated with the DHCP server, e.g., in an enabler **49**. Typically, many software services in the system **10** register with the bulletin board for notification of additions of native devices. Upon registration of a device, these services are notified and enter into relationships, if any, with the new device. A typical situation would involve notifying a system monitor process that a native device is active. That process then updates its web page with information about this device. In a preferred embodiment that utilizes a hierarchical network, each bulletin board covers a given “territory” or region of the network.

4.1.4 Configuration Acquisition

[0100] If a device is unconfigured, it marks the “Configured” attribute on its bulletin board entry as “False”. Any system monitor receiving device notifications to this effect generates an alert. Once a device has its name, it is able to communicate with any networked supply of non-volatile storage. This is particularly useful in the case of devices with limited in device non-volatile storage, whom a system configuration utility notifies of the location of configuration information. Devices that can retain all of their configuration information can start up using that configuration. Actions taken during start-up, e.g., taking PID algorithms out of hold, are configurable.

4.1.5 System Monitors

[0101] A system monitor monitors the health of system components (hardware and software), monitors the health and operation of user applications, provides configuration information to devices that request it, and monitors the licensing of software and upgrades for purposes of alerting the appropriate groups. The system monitor supports standard SNMP agents and any native specific system management protocols.

[0102] The predominant source of information used by the system monitors are the bulletin boards which, as noted above, are used by the native devices to record their current state of operation. As devices are added and removed from the bulletin boards, the system monitor displays that information. The information posted on the bulletin boards includes diagnostic information. A system monitor not only displays this information but also allows the user to access and operate any embedded diagnostic displays and operations.

4.2 Device Configuration and Management

4.2.1 Device Startup

- [0103] Initial start-up of a native device involves the following steps:
- [0104] 1) Install the device on the network **48** and turn on the power as required by the device.
- [0105] 2) The device obtains an IP address from a DHCP server **49** on the network **48**.
- [0106] 3) The device registers with the native name service.
- [0107] 4) The device waits for a configuration.
- [0108] 5) The device is now ready for configuration.
- [0109] Restart of a native device is effected by the steps of
- [0110] 1) Install the device on the network **48** and turn on the power.
- [0111] 2) The device obtains an IP address from a DHCP server on the network.
- [0112] 3) The device registers with the bulletin board.
- [0113] 4) The device begins operation according to the configuration stored in its non-volatile RAM. This RAM may be a network resource that it must recover. If that resource is unavailable, the device is essentially unconfigured and the above procedure is followed.
- [0114] 5) The device is now ready for configuration and operation.

4.2.2 On-Line Configuration

[0115] The steps involved in on-line native device configuration are:

- [0116] 1) The engineer starts a web browser on workstation **40**, handheld computer **44** (wireless or otherwise), a PC, a native workstation, etc.
- [0117] 2) The default home page on the browser is cached and contains an applet that locates the native bulletin board and raises the initial web page that lists the services available from that workstation.
- [0118] 3) The engineer selects the ADE and a native device of interest.
- [0119] If the device is on a native network **48**, the native configurator (ADE) is used to make the changes. In this case, the web browser is pointed at the web server that hosts the ADE. This approach eliminates conflicts between field change and database changes. The configuration of non-native devices is provided by a service that encapsulates the native commands and passes them through a native device to the actual foreign device.

[0120] The steps in PCO configuration are

- [0121] 1) As in the off-line mode, the engineer accesses the ADE and makes necessary changes.
- [0122] 2) Once the changes are complete, the engineer tells the new objects to begin execution in "open loop" mode, i.e., their outputs are not allowed to go to the field or to any other display station other than the one used to configure them. (For examples, their names are not registered with the native name service.)
- [0123] 3) In the open loop mode, the engineer can "Verify and Validate" his new control scheme by viewing the detail displays for the new and the old objects "in parallel". The set of available detail displays may include special displays for evaluation of parallel composites.

- [0124] 4) Once the configuration is complete, the device records its new configuration in its local non-volatile RAM. Optionally, it updates a remote configuration database or just set a "changed configuration" status indicator much like conventional intelligent transmitters.

5 Native Software Subsystems

5.1 Security

5.1.1 Object Access

5.1.1.1 Locks and Keys

[0125] As noted above, the proposed security mechanism is a lock and key. In this model, the attributes of an object that are available to be manipulated are assigned a lock number. An application registers with the security system at start-up to obtain a key. The value of the key depends on the effective user id (name) and the station on which the application is running. The key is never visible to the user, so he/she cannot alter it. This key is available to any library that uses the security system to ensure proper access.

5.1.1.2 Security Database

[0126] A security database contains a list of users. For each user it allows the specification of a master key, i.e., a bit array with a bit set for each lock the user is allowed to unlock. There is one master key for each object type supported by the security system. By default, the security system supports the creation of master keys for the system monitor, process data (PCO attributes), and native messages (events). The database also allows modifications to the master key for a user based on the station used to perform an operation. These exceptions can be expressed as IP addresses, station names, or by the type of access being used (e.g., extra-system, inter-system, or intra-system).

5.1.1.3 Operation

[0127] The security system supports caching the security system database on specific servers to improve redundant operation and to speed key queries. This is generally not a significant issue since the query is done once per application. Each time an operation is attempted or, in some circumstances, once per connection, the key is passed to the target object along with any other information required by the object to achieve the user's request. It is the responsibility of the object to match a bit in the key with the lock on the attribute or method being altered. If there is a match, the operation proceeds. If the match fails, the operation is aborted and the calling application is notified of the failure. If the lock is zero, the operation is allowed to continue.

5.1.1.4 Use

[0128] According to a preferred practice, the illustrated control system **10** supports at least 256 individual locks. This allows the application engineer to define **255** groups of objects for each type of object being accessed.

5.1.1.5 Impact on PCOs and Other System Components

[0129] The value record of a PCO attribute includes two keys: one for read access and one for write access. A default

value for the key can be inherited according to rules concerning the use of the attribute, e.g., all alarm attributes have an assigned group by default. A system monitor interface implements the same mechanism.

5.2 Maintenance and Support

5.3 Configuration

[0130] A configuration object (CO) is a collection of objects plus the editor used to manipulate them. A Configuration Object Template (COT) is a CO that has not been instantiated. The application development environment (ADE) is the environment for configuration object editors. It shows the user the existing COs in the CO being edited and a list of available templates. It provides the mechanism for adding objects to a CO by instantiating a COT and for removing COs. A Configuration Object Component (COC) is any of the objects found in a Configuration Object. Typical components are: PCOs, Process Graphics, Reports, Historian Configuration Objects, etc. The ADE has two components: the human interface and the server. Several human interfaces can access the same server at any time.

[0131] When a Native device is placed on-line, its internal configuration services are disabled. All configuration instructions come from the application development environment. Where on-line configuration is permitted, the system monitor or other system component flags any configuration mismatches.

[0132] Invoking the ADE is tantamount to starting the editor of a particular top-level CO. The ADE presents all of the available COs within the selected CO and all of the available COTs. The ADE provides a mechanism for manipulating those objects and for adding new objects by instantiating COTs. Since the objects within the ADE are also Configuration Objects, i.e., they are themselves collections of objects with associated editors, configuration involves selecting an object, invoking its editor, and making changes. Each editor knows how to store its object's data in the persistent storage.

[0133] The ADE preferably includes a system editor, a control editor, a graphic editor, a message manager editor, a historian editor. All editors support the use and creation of templates by the user. The also support the assignment of their objects to particular stations in the native network. Not all objects require assignment, but the editors are responsible for it.

5.3.1 System Editor

[0134] The system editor allows the user to configure a logical arrangement of native device objects. The only devices it configures are those that are native-enabled. If a device, such as Networked I/O, does not support Native, this editor is not interested in it. The following table lists those devices, what they do, and what needs to be configured.

Device Type	Description	System Editor Configured Attributes
DHCP/Bulletin Board Server	These devices allocate IP addresses to stations and act as the bulletin board for the stations that ask for addresses. They are not configured to know what devices might be attached.	Name and association with a Web Server

-continued

Device Type	Description	System Editor Configured Attributes
Native Controllers	These devices act as PCO platforms.	Name and association with a DHCP/bulletin board Server (which may be a Web Server)
Native Workstations	The devices that supply the human interface to the process.	Name and association with a DHCP/bulletin board Server (which may be a Web Server)
Native Web Servers	This devices host Native Applications and graphics. Native Workstations may also be Web Servers.	Name

[0135] As with the other editors, the system editor is object oriented and template driven. The user is allowed to select an object type from a palette and attach it to the appropriate stations. Individual Device Editors are provided for each of the device types.

5.3.1.1 DHCP/Bulletin Board Editor

[0136] If the device is assigned to another DHCP/bulletin board device or a web server, the editor allows the user to specify the number of IP addresses that this device needs to be able to supply below it. In this case, when the device starts up, it gets its own IP address and IP address range from the Web Server or DHCP/bulletin board device to which it is assigned. If it is not assigned to such a device, the editor allows a range of IP addresses to be assigned to it. It is assumed that the device gets its own IP address from another, non-Native DHCP Server.

5.3.1.2 Native PCO Platforms

[0137] Each device type on the network is provided a specific editor to control its specific configuration. Typical devices are: transmitters, actuators, controllers, etc. These editors are available for off-line configuration. In this case, the system monitor downloads any approved configurations at the next start-up. The editors may also be available on-board the device (e.g., for fixed function devices such as single station controllers), which allows direct configuration through a web browser. The on-board facility is automatically disabled once the device is placed on a native network. This ensures database consistency while allowing the use of the device on non-native networks. In the illustrated embodiment, device specific configurators are provide in ADE add-in and on-board, preferably, with the same code would be available and used for both.

5.3.1.3 Native Workstations

[0138] The configuration of this device includes user name (s) allowed; read only operation; full fledged operator, and so forth.

5.3.1.4 Native Server

[0139] The configuration of this station is a range of addresses if it is acting as a DHCP server.

5.3.2 The Control Editor

[0140] The Control Editor implements PCOs plus four of the IEC 1131-3 languages (LD, ST, SFC, and FB). Any control schemes built in those languages are translated into a Java file, compiled, and downloaded. The IEC 1131-3 concept of configurations, resources, programs, and tasks are translated into the native environment that consists of set of native platforms. In addition, the Control Editor supports defining of control schemes using PCOs in the same manner that it supports the use of IEC Function Blocks; supports the assignment of control entities to specific PCO platforms; the importation of instruments lists to generate, and PCOs.

5.3.3 The Graphic Editor

[0141] The Graphic Editor support all of the human interface (HI) functionality and graphics typical of this type of editor.

5.3.4 Application Configuration

[0142] Applications can only be assigned to native web servers.

5.3.4.1 Historian Object Editor

[0143] A historian object when found in a CO is a representation of a portion of the full historian configuration, e.g., a slice of the historian. The historian object consists of the name of the historian to receive the data and a list of PCO attributes. Each attribute is associated with certain critical information related to the historian, e.g., change deltas, sample rates, reduction operations, archive characteristics, etc. Starting the historian editor resulting in showing the full configuration of the historian(s) to which this object is assigned.

5.3.4.1.1 Historian Data Definition Object Editor

[0144] One historian object exists for each historian in the system. These objects are created automatically if a historian data definition object is created and assigned to a historian. The data in this object controls the general function of the historian, e.g., where it runs in the system, what its archive policy is, what its data reduction policy is, etc.

5.4 Process Control

5.4.1 Definition of the Process Control Domain

[0145] The process control domain is usually divided into three primary control styles: (i) analog control (continuous), (ii) logic control (ladder logic), and (iii) sequential control (supervisory operation of logic and analog control structures). In addition, every process control system must address batch (or recipe based) manufacturing with the related topics of lot tracking and validation.

[0146] The illustrated control system **10** views these styles as particular visualizations of the process control functionality. This means that these three types of control do not exist as objects. Rather, the control system receives a class definition of a process control object (PCO) which it instantiates and executes. It is the responsibility of the control editor to create the correct representation of the control structure while also representing that control structure in one of the industry standard formats, i.e., ladder-logic, function blocks, or structured text, or PCOs directly

[0147] Sequential Function Charts are a meta-language used to control the execution of the other three languages. In the illustrated embodiment, it serves as a wrapper class for the other classes. The source code for a control scheme and the critical values for its PCO attributes, i.e., its checkpoint file, are held in the non-volatile memory of the Native device.

5.4.2 Process Control Objects

5.4.2.1 PCO Attributes

[0148] The fundamental control object is a PCO attribute. As noted above, these are provided in two types: variables and parameters. An example of a variable attribute is a measurement in a PID block. In the illustrated system, a PCO attribute is an object in its own right and it contains certain vital information when it is retrieved. This information is "atomic", i.e., all information is obtained with an attribute, not just its "value". The minimal set of information from a PCO variable attribute is value, data type, quality information time tag in milliseconds, and alarm status. It is this object that is transferred from one PCO to another to implement the transfer of information needed in an actual control scheme.

[0149] Parameter attributes bundle less information than a variable attribute. The typical example of a parameter attribute is the proportional band of a PID block. In the illustrated control system **10**, this attribute has less supporting information than a variable attribute, e.g., it would not have an engineering unit range.

5.4.2.2 PCO Blocks

[0150] PCO attributes and methods to operate on those attributes are combined into PCO Blocks. These blocks are analogous to the control blocks in other systems, e.g., I/A Series FoxBlocks, IEC 1131-3 Function Blocks, and Foxboro Microspec Blocks. However, because they are implemented in Java, they can be executed on any microprocessor running any OS that supports a JVM. In addition, they are designed to occupy only the amount of RAM required for the selected options and to be on-line replaceable.

[0151] PCO blocks capture and encapsulate knowledge of regulatory control algorithms needed to by continuous processes. The PCO user block can be used to integrate particular user-defined control algorithms into a PCO control scheme. It is the PCO user block that is used to implement control strategies that are defined using one of the IEC 1131-3 control languages.

5.4.2.3 PCO Composites

[0152] PCO Composites are implementations of a control strategy. They consist of PCO blocks and other PCO composites linked in such a manner as to provide the control mechanisms needed for a particular process. It is the PCO composite that captures application level control expertise, e.g., complex control structures for major pieces of equipment or even the simpler idioms of cascade control.

5.4.3. PCO Execution Environment

[0153] The PCO execution environment is the software in a native control device that supports the execution of PCOs. This software includes the JVM and any other required code, such as a class loader that loads stand-alone applications for each composite, a class loader that loads an applet for each composite, a block processor that loads Java byte codes for a

composite and executes each composite in turn, and a block processor that loads Java byte codes for PCO Blocks and executes each blocks in turn.

5.5 Communications

5.5.1 Requirements

5.5.1.1 Communication Model Definition

[0154] Several types of communication are provided in the illustrated control system **10**. These types include data transfer, message transfer between application programs, object movement (downloads), class file (code) transfers, and events (unsolicited messages).

[0155] The communications mechanism is redundant, i.e., the failure of one station in the system does not prevent communication (new or established) between any two other stations. It also minimizes the impact on uninvolved stations when the service is used by its clients, e.g., a message is not sent to a network segment if it is not going to be used there. The communications mechanism, moreover, provides an encapsulated implementation that can be replaced “under the hood” as new technologies become available. Also, it allows the illustrated control system **10** to be configured for the appropriate communication resources, e.g., the number of outbound connections. Moreover, it provides performance statistics, e.g., high, low, average transfer times between two stations, bytes per seconds between two stations, etc., for viewing in the system monitor. Still further, it automatically reconnects to an object if it is relocated on or removed from/restored to the network.

[0156] Data transfer in the illustrated control system **10** is provided by a mechanism that permits clients to subscribe to updates in a object. It also provides a mechanism for one-shot get/set access to an object. A client of the data transfer facility supplies the control system with a (i) number of objects that it needs, (ii) conditions under which it wants updates: change delta, maximum time between updates, or both, (iii) be capable of retrieving the values all at once (a snapshot), and (iv) be capable of retrieving the values as series of changes with the option to block, for a user specified amount of time, while waiting for a change. Client publication (write lists) are not supported in the illustrated embodiment. Target devices subscribe to the data in the source device.

[0157] With respect to client one-shots (get), the client of the data transfer facility supplies the system with a number of objects that it needs and retrieve the values all at once (a snapshot). Conversely, with respect to client one-shots (set), the client supplies the system with a number of objects that it wishes to change, and set any part of the object if it has access. The parts referred to are the supporting information in a PCO attribute as well as the value part.

5.5.1.1.1 Message Transfer

5.5.1.1.1.1 General

[0158] While the message transfer mechanism is a general case of the data transfer mechanism, its requirements are given separately. The client of the message transfer facility must support private conversations between applications using a reliable communications mechanism and support by subscription publication of information from one server to many clients.

5.5.1.1.1.2 Private Conversations

[0159] The private conversation mode allows applications to find each other by their application name: the client application is not required to know the station name or IP address of the server application. It also allow applications to send records (messages) of arbitrary size: the private conversation mode guarantees that the message boundary is recognized and that order is preserved. Moreover, it allows both the sending and receiving process to determine if the connection is lost. Still further, it supports transaction based message transfer, i.e., the sender does not get an acknowledgement until the receiver says that the message is safely stored.

5.5.1.1.1.3 Subscriptions

[0160] The subscription facility supports senders and listeners. A sender is a process that has data that may be of interest to one or more listeners. The subscription mode supports this type of operation by allowing a “sending” process to register that it will produce messages of a particular type, allowing a “listening” process to register to receive all messages of a particular type, supporting transaction based subscription if the client so requests (this request would have no impact on the sender), and supporting transaction based publishing by a server if the server requests it (this request would have no impact on the receiver). If no listeners are registered, the subscription facility defines how the message is handled.

5.5.1.2 Security

[0161] The native communications system is linked to the native security system to ensure that unauthorized users are unable to access objects. The illustrated control system **10** does not arbitrate between users of a resource. The prevention of multiple access to an object is the responsibility of the object. The breaking of an application secured link by an operator is not a function of the illustrated control system **10**.

5.5.1.3 Object Location Services

[0162] Object are located by name. To this end, the object location service supports the location of named objects; supports hierarchical object names; allows rule based specification of the name delimiting character; locates an object based on a “longest fit” because (a) not all parts of an object name are globally known and (b) not all parts of an object are in the same physical location; supports the implementation of naming scopes, i.e., limiting the visibility of names; supports the use of a name search path so that relative names can be located; is redundant; supports locating many types of objects, e.g., server processes, PCO attributes, and application programs; and supports the addition of related information about the object in the name directory, e.g., object type, short-cut information to be used by the process that owns the object, object type/implemented interface, and information. With respect to the last item, the object location services permit a client to request the name of a service that supports a specific interface. This allows new interfaces to be added to a service without breaking an old one.

5.5.1.4 Networks

[0163] The illustrated control system operates over any IP based network **48**. By way of non-limiting example, the illustrated control system **10** operates over wide area networks supplied by the customer, installed plant LANs supplied by

the customer, dedicated LANs supplied either by the customer or the vendor, redundant LANs supplied by the vendor. This operation is transparent to all applications using the system **10**. While it is recognized that performance are degraded over low bandwidth (phone lines) and high latency (geosynchronous satellites) networks, the system **10** optimizes as best it can to minimize the impact of different network types. In short, the illustrated control system **10** does not assume that it “owns” the network.

5.5.1.5 Quality of Service

5.5.1.5.1 Performance

[0164] Assume that the illustrated control system **10** is operating in support of a PCO in a dedicated control station (not a transmitter) or in a PC, it supports at least 10,000 object value gets/sets per second if the object is in the local station; supports at least 100 get/set per second requests from a client if the object is in a remote station; supports at least 50 get/set requests per second made to a server from a remote client; is able to locate 3000 objects per second; is able to detect a communication failure in no more than 1000 milliseconds; ensures that the time to move data from a server station to a client station (API call to API call) is less than 100 milliseconds in a normally configured network; allows a server station to move 6000 values out of its box every second; allows a client station to accept at least 2000 values every second; allows an unlimited (but, configured) number of names to be defined in the system; and supports the registration of 50,000 globally known names without requiring a reconfiguration. All server stations are able to support at least 30 client stations for continuous data feed without requiring a reconfiguration. All client stations support at least 30 server stations for continuous data feeds without requiring a reconfiguration.

5.5.1.6 APIs

[0165] Within the illustrated control system **10** environment an API are provided for all of the above services. This API are delivered as a Java Bean so that the underlying communications mechanism may be replaced without breaking all of the applications that use the service.

5.5.1.7 Maintenance and Upgrades of Messaging Services

[0166] The illustrated control system **10** defines and enforces a mechanism for supporting interoperability of messages from stations at different versions. A mechanism like interface inheritance is provided, i.e., a message carries in it enough “type” information that the receiving application knows how to “decode” the message even if the receiving application is at a lower level. This adds fields to messages if necessary, but not remove/replace old fields.

5.5.1.8 Integration of Non-Native Based Systems

[0167] The illustrated control system **10** supports integration of non-Java applications using appropriate technologies, e.g., CORBA, COM/DCOM/ActiveX, and emulation of APIs. Physical gateways are permitted, but these gateways provide only a minimal amount of configuration.

5.5.1.9 Object Location Service Implementation

5.5.1.9.1 The All Java Approach

5.5.1.9.1.1 Description

[0168] JINI and JavaSpaces technologies permit object location within the illustrated system. JINI technology allows

native devices to register names in a globally available database. JINI clients can find this database and perform searches against it. JavaSpaces allow a device to register important information in a globally available bulletin board. JavaSpaces provide four simple services: posting, removal, query by example, and notification. The notification feature informs prospective users of a class of objects when such objects are added and removed from the JavaSpace. A JavaSpace handles services like system monitoring.

5.5.1.9.2 The Enhanced Object Manager Approach

5.5.1.9.2.1 Description

[0169] As an alternative, an enhanced object manager (EOM) approach to address location and connection of data clients and servers can be utilized as follows:

[0170] 1) A LDAP compliant service are created and used to store pathnames and related information, e.g., IP addresses, Port Numbers, internal indexes, etc.

[0171] 2) The EOM API provides get/set and read/write support similar to that provided by the OM today.

[0172] 3) When provided a name, the EOM looks in the local copy of the LDAP database.

[0173] 4) If the name is found and the data is local, the request is passed to the correct data provider (the EOM, the CIO database, or the AOS database) who returns the value and other data. The EOM handles sending updates on change driven data.

[0174] 5) If the name is found and not local, the EOM sends the request to the remote EOM which then makes a local query. The EOM handles sending updates on change driven data.

[0175] 6) If the name is not found locally, the location request is passed to the master LDAP database. If it is found there, the address information is passed back to the local LDAP database (which is updated) and steps 4) and 5) are repeated.

[0176] 7) If the name is not found remotely, the request is posted for a period of time. If it is resolved in that period, the requesting station is notified. If it is not resolved, the master LDAP database drops the name and notifies the requesting station.

5.6 Critical Applications

5.6.1 Time Synchronization

[0177] SNTP (Simple Network Time Protocol) is used to keep operator stations and similar non-process data producers in sync. This protocol is as accurate as 1 ms in a carefully controlled environment, but 50 ms is more common on shared networks. Controllers requiring high accuracy are placed on tightly controlled networks or equipped with an interrupt used to coordinate time updates. In either case, a master timekeeper tied to a well known reliable data source—GPS clock or broadcasted time from the national atomic clocks. This master time keeper generates the interrupt and SNTP. It is built with a highly reliable internal clock so that loss of time feed does not result in significant drift within the Mean Time to Repair (MTTR).

5.6.2 Alarm and Message Management

[0178] Alarm and message manage requires that all process alarming to be based on the top-level composite’s attributes;

that acknowledgments be handled at the composite level; that PCOs, system monitors, and the native human interface a reliable message protocol to send the message to a message management system (MMS); that the MMS is used by end-users to view the current alarm state and the alarm history; and that the MMS use the plant historian to permanently store and archive messages.

[0179] To meet the control market's requirement for a message management system, the illustrated control system **10** provides a message management support service (MMSS) in the native devices that are used to deliver messages to a Message Manager. The service operates as follows.

[0180] When a client of the MMSS connects to the MMSS server, the MMSS server send any messages in its internal alarm queue that are older than the time of the connection to the client; transmits the alarm state at the time of the connection to the client, and transmit any new messages subsequent to the connection. If the connection to the MMSS client fails, there is no impact on the operation of the control station or the control network; minor delays associated with detecting the loss of connection are generally viewed as acceptable. The MMSS support at least two (2) and, preferably, four simultaneous clients.

[0181] The alarm state of a PCO attribute is defined to include block attribute identification (full PATH attribute name, at least); the bad I/O status of the block for I/O attributes; an indication of the alarm status for each alarm type defined for the attribute; the priority, criticality, and acknowledgment status of the attribute; the on/off status of the attribute, i.e., is it updating; the status of the alarm inhibit and alarm disable or alarm option parameters, and the time tag of each alarm event stored within the block, e.g., the into alarm time, the acknowledged time, and the return to normal time. This information is recovered from all currently active alarms.

5.6.2.1 Configuration

5.6.2.1.1 Process Alarms

[0182] When a composite is configured, the application engineer specifies which attributes of the PCOs within the composite are to be annunciated. When a composite is placed on line, it notifies the object location service that it exists. A message management service are notified of the new composite and arrange for notification from the device.

5.6.2.1.2 System Alarms and Messages

[0183] No user configuration is required to get these alarms into the MMS. The sources of such messages arranges for the appropriate notification.

5.6.2.1.3 Native Operator Actions

[0184] No user configuration is required to get these alarms into the MMS. The sources of such messages arranges for the appropriate notification.

5.6.2.2 Native Message Management Service

[0185] The MMS are centralized and redundant. It provides a complete alarm state for the plant at any time including startup. It retains a message history for as long as configured through the use of the plant historian. Users view the alarms through to the Alarm Manager.

6 Native Hardware SubSystems

6.1 Field Devices

6.1.1 Overview

[0186] The native field devices are coupled to the network **48** via an IP network, preferably, Ethernet and, still more preferably, powered Ethernet. Powered Ethernet delivers the power, along with conventional Ethernet data. To this end, the native field devices use one of two wiring schemes, with each scheme supplying power to the field device: (a) standard four wire 10Base-T wiring, or (b) industry standard two-wires used for 4-20 ma based transmitters. The devices also provide a single Ethernet interface (or redundancy, where desired in the cabling or the field device); and connect to a non-redundant powered hub with the following features: (a) non-redundant connections to the IFDs; and (b) support for redundant connection to a plant-wide network. Where used, the powered Ethernet implementation leaves the physical and logical Ethernet interfaces in place. This means that except for adding a new wiring type, the physical layer, the data link layer, the use of CSMA/CD, and 10Base-T connectors are preserved.

[0187] The illustrated embodiment utilizes powered Ethernet hubs to provide communications and power to its intelligent field devices. These IFDs may be transmitters, actuators, Networked I/O, or PCO platforms. The powered hubs are preferably, stackable, DIN rail-mountable, support connection to a redundant network, and support both 10 Mbps and 100 Mbps Ethernet.

[0188] A further understanding of the operated of the powered Ethernet network and of the circuitry used within the native field and control devices to draw power therefrom may be attained by reference to U.S. patent application Ser. No. 09/444,827, filed Nov. 22, 1999, entitled POWERED ETHERNET FOR INSTRUMENTATION AND CONTROL, and by reference to the corresponding PCT Patent Application US00/10013, filed Apr. 14, 2000, the teachings of which are incorporated herein by reference, and a copy of which is attached as an appendix hereto.

[0189] On the software side, specific implementations of the native field devices support web based configuration, as described above. The support is supplied by including an embedded web server with a selection of pages used for configuration and maintenance.

6.1.2 Transmitters

[0190] Intelligent transmitters use the IEEE 1451 standard to communicate to their sensor(s). This facilitates use of the same module in many types of transmitters from potentially many vendors and allows the transmitter to support up to 255, perhaps externally mounted, sensing devices.

6.1.3 Positioners

[0191] The native intelligent positioners also support a powered Ethernet interface like the one used by the native intelligent transmitters. FIG. 6 shows a preferred native intelligent positioner implementation. The approach standardize interfaces and use a PCO to control the positioner as evident in the drawing.

6.2 Summary of Hardware/Software Features by Station Type

[0192] A summary of the features of native control devices follows. Blank cells indicate that the feature is not present in the indicated device.

Features	Xmitter	Positioner	Controller	Integrator	PC based Operator Consoles	Native Workstation	PC Based Web Server	Solid-state Web Server	SS DHCP Server	SS bulletin board
Low power, high performance CPU	Yes	Yes				Yes				
Top performance CPU			Yes	Yes	Yes	Yes	Yes	Yes		
RAM	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
ROM	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Flash	Yes	Yes	Yes	Yes		Yes	Yes	Yes	Yes	Yes
Sensor I/F (IEEE 1451)	Yes	Yes								
Powered Ethernet (2 wire)	Yes	Yes								
Serial I/F (RS-232/485/422/423)	Yes	Yes		Yes						
Ethernet (10Base-T)	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Wireless Ethernet						Yes		Yes		
the TCP/IP suite	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
DHCP Server							Yes	Yes	Yes	Yes
Full OS	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Java Virtual Machine	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Web Server	Yes	Yes	Yes	Yes		Yes	Yes	Yes	Yes	Yes
Device	Yes	Yes	Yes	Yes					Yes	Yes
Configurator										
Annunciator					Yes	Yes				
Keyboard										
Touchscreen					Yes	Yes				
Mouse/Trackball					Yes	Yes				
Display					Yes	Yes	Yes	Yes		
Hard drive					Yes		Yes			
Solid State Bulk Storage						Yes		Yes		
System Monitor							Yes	Yes		
Process Control Objects (PCOs)	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes		
System Management Objects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Omnibus Messaging Services	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes		
Profibus PA Support	Opt.	Opt.		Yes						
Profibus DP Support	Opt.	Opt.		Yes						
HART Support	Opt.	Opt.		Yes						
FoxComm Support	Opt.	Opt.		Yes						
Fieldbus HSE Support	Opt.	Opt.		Yes						
Fieldbus H1 Support	Opt.	Opt.		Yes						
DeviceNet	Opt.	Opt.		Yes						
LONworks	Opt.	Opt.		Yes						
Modbus				Yes						
Modbus+ ABDH				Yes						
I/A Series Nodebus				Yes						
I/A Series Fieldbus (FoxComm)	Opt.	Opt.		Yes						
Application Development Environment							Yes			
Historian							Yes			
Time Synchronization							Yes	Yes		
PIMS							Yes			
Message Management							Yes			

[0193] Described above are apparatus and methods that achieve the goals of the invention. It will be appreciated that the embodiments described herein are examples and that other embodiment incorporating changes thereto also fall within the scope of invention. Thus, by way of example, it will be appreciated that the invention can be implemented using a virtual machine environment other than JVM and using byte-code other than Java bytecode. By way of further example, it will be appreciated that the apparatus and methods taught herein can be applied to a range of control application, in addition to process control.

In view of the foregoing, what I claim is:

1-228. (canceled)

229. A device for use in a control system, comprising:

a controller having a powered Ethernet connection and at least one communication connection; and

a virtual machine environment that executes a byte code therein, the byte code configuring the controller to execute a control program for controlling filed devices within a control network;

wherein the powered Ethernet connection is adapted to provide power to the controller from the control network, and the at least one communication connection is adapted to provide for communication with at least one field device.

230. The device of claim 229, wherein the communication connection is a bus connection.

231. The device of claim 230, wherein the bus connection is adapted to provide for communication with a field device via a bus operating in substantial accord with a protocol selected from the group consisting of: Foundation Fieldbus, Foxcomm, Profibus, ControlNet, Modbus, Devicenet, HART, NodeBus, and LonWorks.

232. The device of claim 229, wherein the communication connection is an Ethernet connection.

233. The device of claim 229, wherein the controller further comprises a processor adapted to execute a control algorithm to control the at least one field device.

234. The device of claim 229, further comprising a first I/O card comprising the powered Ethernet connection and a second I/O card comprising the at least one communication connection.

235. The device of claim 229, wherein the device is coupled for communication with a control device adapted to monitor and control aspects of a process.

236. The device of claim 229, wherein the device functions at the direction of a plant-level processor.

237. The device of claim 229, wherein the at least one communication connection provides power to the at least one field device.

238. The device of claim 237, wherein the at least one field device includes a sensing element and the device includes a pair of wires providing excitation to the sensing element.

239. The device of claim 229, wherein the field device obtains power via the at least one communication connection and via the powered Ethernet connection.

240. The device of claim 229, wherein the device is coupled for communication via one or more networks.

241. The device of claim 240, wherein the device further comprises a device integrator adapted to provide communication between a plurality of the one or more networks.

242. The device of claim 240, wherein the controller is coupled for communication with one or more control devices

disposed on the one or more networks and adapted to monitor and control aspects of a process.

243. The device of claim 229, further comprising one or more networked I/O modules for providing interoperability between control systems.

244. The device of claim 229, wherein the at least one communication connection comprises at least two communication connections operating according to different protocols from one another.

245. A device for a process control system, comprising:

a powered Ethernet connection adapted to provide power to the device and a communication connection adapted to provide for communication with a field device via a communication link that provides power to the field device; and

a virtual machine environment that executes a byte code therein, the byte code configuring the controller to execute a control program for controlling filed devices within a control network.

246. The device of claim 245, wherein the device further includes a first network I/O card for providing the powered Ethernet connection and a second network I/O card for providing the communication connection.

247. The device of claim 245, wherein the device further comprises:

a processor for implementing station management software and for communicating with a control device adapted to monitor and control aspects of a process and a field device; and

a memory for storing data.

248. A control device for a system, comprising:

a powered Ethernet connection for communicating with a control network, the powered Ethernet connection providing power to the control device;

a processor that executes a control algorithm to control one or more field devices associated with the control device; and

a virtual machine environment that executes a byte code therein, the byte code configuring the controller to execute a control program for controlling filed devices within a control network.

249. The control device of claim 248, wherein the field devices are sensing devices adapted to sense any of (i) an environmental condition and (ii) a physical condition, the physical condition comprising any of temperature, pressure, flow, and position.

250. The control device of claim 248, wherein the field devices are actuators adapted to control a mechanism affecting a physical condition of an environment.

251. The control device of claim 250, wherein the mechanism comprises any of a valve, paddle, inlet, and outlet.

252. The control device of claim 248, wherein the processor further executes another control algorithm that at least one of (i) maintains the system at a desired level, and (ii) drives it to that level as a function of one or more measured values and one or more setpoints.

253. The control device of claim 248, wherein the processor further executes station management software adapted to allow the control device to report on any of health information, performance information, and status information.

254. The control device of claim 248, further comprising: a bus connection for communicating with one or more field devices via a bus operating in substantial accord with a

protocol selected from the group consisting of: Foundation Fieldbus, Foxcomm, Profibus, ControlNet, Modbus, Devicenet, HART, NodeBus, and LonWorks.

* * * * *