

The Next Generation Networking Paradigm: Producer/Consumer Model

As today's source/destination-based networks cannot offer the required functionality or accommodate increased traffic, system capabilities and productivity improvements are restricted. Consequently, a new network model - one that provides more functionality, makes more efficient use of bandwidth, and increases information flow, all while reducing traffic on the wire - is needed.

In a discussion of what is needed in the new network model regarding diagnostics, explicit and I/O messaging and throughput, the producer/consumer network model is revealed as the only model available today that can meet the control environment's demanding requirements and allow for future migration. The paper concludes with a discussion of the benefits of the producer/consumer network model, including Multicast and two one-way I/O trigger mechanisms: change-of-state and cyclic I/O production.

INTRODUCTION

If there's one thing we've all learned over the past decade, it's that users are demanding more from their control systems, and consequently, from the networks that tie the system together. Users want better diagnostics available over the network, less downtime, and reduced installation and maintenance costs. At the same time, they are demanding improved throughput.

With increased functionality comes more traffic and data on the wire. Today's networks, which are source/destination based, cannot offer the required functionality and accommodate increased traffic, thus restricting system capabilities and productivity improvements.

Increased demands on networks have forced the evolution of a new network model - one that provides more functionality, makes more efficient use of bandwidth, and increases information flow, all while reducing traffic on the wire.

Unfortunately, much of the discussion to date about networks has focused on baud rates, protocol efficiency, and physical characteristics (i.e. type of wire used). In reality, it's more complex than that. Available diagnostics, messaging types and throughput must all be considered when evaluating a network.

The most important factor affecting these capabilities in the control environment is the network model.

The source/destination model used for the past two decades can no longer meet today's network needs. The only model available today that can meet these demanding requirements and allows for future migration is the producer/consumer network model. Here's why:

DIAGNOSTICS

Networks provide a convenient way to retrieve diagnostics from devices. Device-specific information,

such as detection of a photoelectric sensor's low margin due to a dirty lens can be communicated over the network to the control system during run time. The network delivers the diagnostic to the system operator interface, alerting plant personnel to the problem. The lens can be cleaned at a convenient time before there is a glitch in the process. Trouble-shooting a device, reading its fault codes, updating data logs -- all while not impacting the remote I/O control data exchange among other nodes -- is a must.

EXPLICIT AND I/O MESSAGING

Explicit messages, used for device configuration and diagnostics, are extremely flexible, with the data field carrying protocol information and instructions for service performance. For example, a message would be able to write new presets to five timers in a controller, or to execute a self-test. Explicit messages are used for uploading and downloading programs, modifying device configurations, and data logging, trending and diagnostic functions. Nodes must interpret each message, perform the requested task, and generate responses. These types of messages are highly variable in both size and frequency.

I/O messages on the other hand are implicit in nature. The data field contains no protocol information, only real-time I/O control data. The meaning of the data has been predefined and processing time in the node is minimized. An example of an I/O message is a controller sending output data to an I/O block, and the I/O block responding with its input data. Such messages are low overhead, short, vary frequently and require high performance.

In the past, manufacturers have had separate networks to deal with the very different requirements of these two messaging types. A network used for I/O control cannot tolerate the variability introduced by explicit messaging. Allen-Bradley's blue-hose duo, DH+/RIO and Siemens' Profibus FMS/ Profibus DP are examples of

this situation.

Today's users are demanding both functions on the same wire. And today's smarter devices need the functionality provided by both messaging types. Yesterday's source/destination networks cannot deal with these modern demands.

THROUGHPUT

Ultimately it's the throughput required by the application that determines what type of network model is required. Throughput is the rate at which input data from devices can be delivered to all nodes that need it and the resulting output data (decisions) can be delivered to all the devices that need it. Nodes include sen-

cally limited to that function alone to obtain the necessary repeatability and throughput for control.

Peer-to-peer networking goes beyond master/slave, providing considerably more flexibility. But as a result most networks that support peer-to-peer use explicit messaging.

PC-based programming and configuration of controllers uses explicit peer-to-peer messaging. PC-based MMI also use explicit peer-to-peer messaging. As additional MMI units are added, the network load increases dramatically as each unit typically will read all the same variables out of each controller as does the prior units so an operator can get the same alarms, trends, and graphics from multiple locations.



Figure 1.

sors, operator interfaces, controllers, data loggers, alarm monitors, actuators, etc. It is determined by baud rate, protocol efficiency, and most important of all, the network model, or delivery method. Let's briefly discuss each.

Baud rate is raw speed. It's unfortunate that this is often the most used measure of performance because it's the most misleading. Not only that, but with today's new networks, it's the least important of the three throughput factors.

Protocol efficiency - data bytes (the payload) versus the total bytes in the packet - typically expressed as a percent, is a measure of the network protocol overhead. While important, it is not nearly as significant as the data delivery and exchange method (network model) used. If a particular information exchange takes two or more packets on the wire as compared to one, the fact the one protocol has 25 percent greater efficiency becomes meaningless.

To keep nodes from dominating the wire, most peer-to-peer networks use some sort of token rotation algorithm. While these algorithms have been enhanced over the years to be more "fair, the basic flexibility that makes it attractive makes its use for peer-to-peer interlocking between controllers very problematic. Response times vary considerably for any given message, depending on load and on how "far away" one is from the token holder when there is a need to speak.

Frequently low-end electronic operator interface (EOI) units will be found on I/O networks, basically replacing simple push button, pilot light and meters. But as each EOI device is added, an additional load of typically the same data new node with a different destination address is added to the network. While variability isn't a factor because of the fixed nature of such loads, the increase in data load slows response time for all nodes, including the real I/O. It's not just EOI that's causing excess network loading. As I/O devices get



Figure 2.

Network model. Every control vendor has its own favorite networks, whether it be Data Highway Plus, Remote I/O, Profibus FMS, Profibus DP, Interbus-S, ASI, Modbus Plus, GeniusLan, or Lonworks. All these networking options have the same thing in common. They follow the legacy source/destination network model. A typical packet is shown below.

In master/slave implementations of this model, the source field is usually not present, as the master is the only source and all responses from slaves are for the master. This master/slave polling is inherently a one-to-one data exchange. It is typically used for the exchange of real time control data (I/O messaging). When used for I/O exchange, such networks are typi-

smarter, the extra diagnostic and configuration data can absorb considerable bandwidth.

Whether master/slave or peer-to-peer, destination-oriented networks waste considerable bandwidth sending the same data set to multiple nodes. Trying to do coordinated control like sending a new setpoint to different drives in a synchronized manner is very difficult, as data arrives at each drive at a different time.

The new network paradigm: The producer/consumer network model

To manage the growing need for data, smarter devices and better control, new networks that simply increase the baud rate or number of nodes only postpones the

inevitable. What is needed is a whole new network model that is designed to manage today's control issues.

That new model is producer/consumer. With producer/consumer networks, messages are identified by content. If a node needs data, it will "accept" that identifier and consume it.

Multicast. Because data is identified by its content; if a node needs that data, multiple nodes can consume the same data at the same time from a single producer. Nodes may be synchronized more precisely while achieving more efficient bandwidth usage. The source of data has to produce the information only once. Additional EOI and MMI can be added without increasing network traffic, since they can consume these same messages. And nodes can produce more than one data set, each using a unique identifier.

Multicast is inherently impossible with source/destination networks, although attempts have been made. Some have added a third field for a group destination and then reserved node numbers for group destination. Others allow a node to carry more than one node number. But these are all band-aid approaches, desperately trying to extend the exhausted legacy source/destination model.

Producer/consumer also allows for two new powerful I/O triggers, in addition to traditional polling. Polling is born out of the source/destination model, and is inherently a two message bi-directional transaction (originator sends output data, and receiving node responds with input data). This transaction is repeated as rapidly as possible to minimize latency from when an input occurs and is delivered to the controller. Most polling cycles are filled with the same outputs and inputs, wasting bandwidth.

With the producer/consumer model, two more efficient and effective one-way I/O trigger mechanisms are available: change-of-state and cyclic I/O production.

Change-of-state (event-based) data production. Nodes produce data only when that data changes. There is no "network polling cycle delay," and, as a result, the data is delivered to all consumers when it changes. A background heartbeat is produced cyclically so that consumers can tell if a device hasn't changed from one that is no longer online. Change-of-state can dramatically reduce network traffic and the load-on typically needed to repeatedly receive, process and generate the same data.

Cyclic (time-based) data production. Cyclic data production involves nodes producing data at a user-configured rate. Data is updated at a rate appropriate to the node and the application. Data can be sampled and produced by sensors at precise intervals for better PID control. Controllers can collect a stock of data for operator interfaces and produce it a couple times a second, plenty fast for human consumption; thereby preserving bandwidth for nodes with rapidly changing I/O.

Both peer-to-peer and controller-to-device exchanges can be handled more efficiently with cyclic and change-of-state data production of producer/con-

sumer networks. Operator interface needs can be layered on top of I/O traffic with minimal increases to network load.

At the same time, producer/consumer networks can accommodate the flexible explicit messaging, point-to-point needs for device configuration and programming. Certain identifiers are typically specified for such traffic and nodes know they contain destination and other protocol information. These identifiers, coupled with the network access method, combine to insure that explicit messages, with their assorted larger overhead are much lower priority on the network than the I/O messages. Large uploads and downloads adjustments to configurations parameters, and diagnostic activities by users with their S/W tools are relegated to background traffic, fitted between the higher priority I/O messages.

No need for users to run both an I/O network and explicit message network through the plant. And no need for vendors to put an I/O port and an explicit message network port on devices. Is it any wonder that the newest open control networks -- DeviceNet, ControlNet, and FOUNDATION Fieldbus -- are all based on the producer/consumer model?

Will the source/destination model disappear?

Source/destination is a "hand me down" from the computer and data processing industry. While limited, source/destination systems are still well-suited for a variety of applications which do not require complex coordination and sharing of data. The flexibility and efficiency of the producer/consumer network model will allow for the expanded functionality demanded by today's applications and is well suited for tomorrow's smarter devices. In this day in age where users are demanding more (functionality, diagnostics) with less (one wire, not two), both users and vendors need a control networking strategy that works smarter -- and, consequently, a network model that works smarter and accommodates the future ■

Patricia A. Murphy is manager for Emerging Technology and Standards at Rockwell Automation Control Systems.

Murphy's responsibilities include:

- *Integrating advanced technology into Rockwell Automation Control System's Communications Business.*
- *Coordinating advanced technology projects.*
- *Participating in industry consortia such as Fieldbus Foundation, ODVA and ControlNet International.*
- *Leading business standards activity and participating in international standards committees as appropriate.*

Murphy has many years of marketing and product management experience in technology driven industries including automation, with Rockwell Automation Control Systems, and telecommunications, with Ericsson and GE.