SYNQNET®
THE HIGH PERFORMANCE UPGRADE TO ETHERNET

William H. West & D. Marshall Matheson
Danaher Motion
Santa Barbara, CA

Abstract – *High performance motion control over a digital network demands tightly managed timing regimens to ensure synchronous and real-time updates across multiple axes. SynqNet was specifically designed to support high performance centralized control systems and offers additional benefits including self-healing fault tolerance operation, simple discovery-based configuration, and high noise immunity. In addition, SynqNet is supported by multiple drive vendors and OEMs who deliver a wide array of cost competitive products. Taking all of these things into consideration, SynqNet offers a unique performance/price advantage, which is why according to ARC Advisory Group, SynqNet has demonstrated the fastest adoption rate to date historically of any motion network. When choosing a motion network the key factors to focus on are very simple. Does it deliver the performance you need? Is it reliable and safe? How easy is it to use? And what is the true cost of ownership? This article reviews these four areas and the technical background to help facilitate the readers understanding of digital motion networks.*

I. INTRODUCTION

All networks are not created equal. High performance motion networks demand tightly managed timing regimens to ensure synchronous and real-time updates across multiple axes. While Ethernet offers adequate performance for general purpose applications that use distributed control, it is generally too slow for more demanding situations. In such cases, a fast synchronous network is required to connect a centralized motion processor to multiple servo axes. SynqNet was specifically designed to support high performance centralized control systems and offers additional benefits including self-healing fault tolerance operation, simple discovery-based configuration, and high noise immunity. In addition, SynqNet is supported by multiple drive vendors and OEMs who deliver a wide array of cost competitive products. Taking all of these things into consideration, SynqNet offers a unique performance/price advantage, which is why according to ARC Advisory Group, SynqNet has demonstrated the fastest adoption rate to date historically of any motion network.

When choosing a motion network the key factors to focus on are very simple. Does it deliver the performance you need? Is it reliable and safe? How easy is it to use? And what is the true cost of ownership? This article reviews these four areas and the technical background to help facilitate the readers understanding of digital motion networks.

II. HIGH PERFORMANCE MOTION CONTROL

Where Ethernet runs out of steam

Ethernet has become the ubiquitous office network platform, and it seems that it will continue its triumphant progress into industrial automation. In the early days of digital networks, as seen from the industrial automation sector perspective, the lack of bandwidth, determinism, and high latency led to distributed processing solutions being offered to the market. For motion systems, the result is intelligent servo drive products that are used to interpolate between the irregular and infrequent data points transmitted over the network.

For high performance and multi-axis applications, such methods are not sufficient. They require a different control model where the processing is performed centrally. The ±10V analog torque interface has until recently been the industry standard used in centralized motion control, but new networking solutions promise to bring valuable change.

SynqNet is the first commercially available 100BaseT (IEEE802.3) network that offers all the performance advantages of a centralized control model, together with enhanced performance, fault tolerance, reliability, and diagnostic features. SynqNet has been developed by the California-based motion control experts, Motion Engineering, Inc.—a Danaher Motion Company—and has been quickly adopted by leading drive companies like Yaskawa, Advanced Motion Controls, Panasonic, Glentek, Sanyo Denki, Trust Automation, and Soonhan Engineering.

High performance motion control systems depend upon a number of key technology components that work seamlessly together in an integrated fashion. A perfect control system must transmit the desired motion profile into the movements of one or more servo axes. Achieving this often requires translation from XYZ 'space' coordinates to machine or 'joint' coordinates using some form of a kinematic model. Additionally, any mechanical imperfections such as non-linearity or the effects of axis cross coupling should be compensated for to optimize machine performance.

Kinematic models and compensation techniques are not new concepts and they rely upon a central motion processor to perform fast, accurate matrix computations based upon multiple inputs generating multiple outputs. The term MIMO (Multiple In Multiple Out) is often used to describe this generic class of control system and the software control model. The exact type of inputs, outputs, and matrix computations will vary by application and the proprietary know-how of the machine builder.

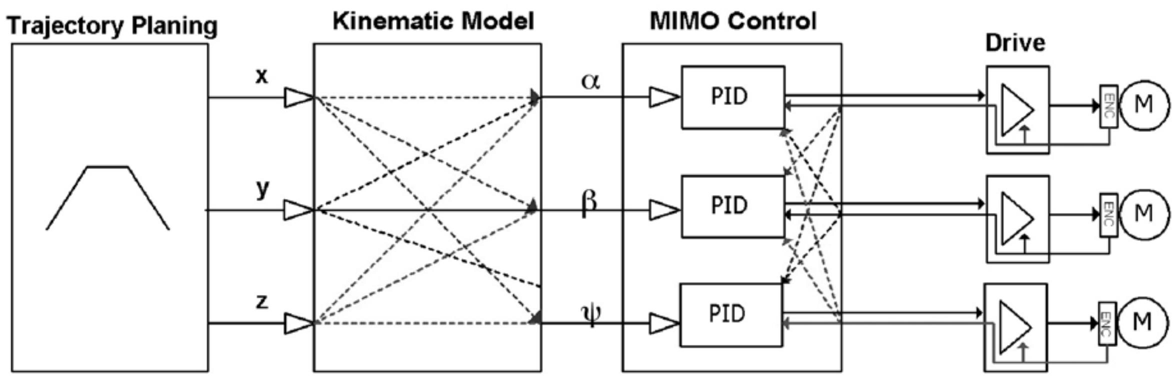

Figure #1, MIMO Model.

Whatever the final software control model might be, it is important that the total servo cycle time is minimized. The shorter the cycle time, the tighter the control system and the higher the performance of the machine mechanism. For fast point-to-point moves or accurate path motion, the cycle time becomes a significant factor in machine performance.

Network Latency & Transmission Rate Define Machine Performance

Modern control systems take multiple demand inputs and multiple feedback inputs, such as the actual torque and velocity, from each axis to compute new target data for each motor. For effective high performance control, demand data and feedback data must be transmitted synchronously, with very short cycle times and low latencies. Any transmission delay represents a phase delay in the control system, which limits the achievable gain and the effective response time of a machine.

Cycle time includes time to acquire feedback data, time to perform matrix computation, and time to transmit new target data. The key technological elements required to achieve this are a fast synchronous network~~,~~ and a fast processing engine. MIMO control cannot be performed on distributed control systems that have decentralized processing power~~,~~ and relatively slow connections between network devices.

	SynqNet	Ethernet (IP/TCP/UDP)	Profinet IRT	SERCOS	Powerlink
Min. Cycle Time/Latency	**<25μs**	1..4ms	1ms	250μs	400μs*

*Protected mode with custom hardware and RTOS

Figure #2, Network Latency Examples.

Real-time Operating Systems
SynqNet works across multiple operating platforms including a variety of real-time operating systems (RTOS) and extension to Windows. But a RTOS is not a requirement to achieve high servo update rates with a high axis count. Because SynqNet utilizes a dedicated hardware control scheme, time critical cyclic operation is guaranteed. Other Ethernet based motion networks such as Powerlink and EtherCAT® must use custom hardware with the addition of a RTOS to achieve a reasonable level of dedicated motion performance outside of basic I/O handling. In addition, the TCP/IP stack is disabled with both Powerlink and EtherCAT during run-time to avoid the added overhead and non-determinism inherent in the protocol. By comparison, SynqNet utilizes a background run-time service channel that can execute during synchronous operation.

Synchronous Networks Offer Low Skew & Jitter for Improved Path Accuracy
All networks, including any digital system, rely on sampling of data at some discrete time based upon a clock reference. When independent systems with independent clocks are connected together, as is the case with a network, the natural and random variances in clock frequency can present a challenge. Many engineers are familiar with the concept of 'beating' when two high frequency sources, closely matched but not exactly, beat at the difference frequency. In an aircraft with multiple engines, such beating is audible as a throbbing noise. Digital control systems are no different and in a collision-free network, the beating ('Jitter' in networking parlance) arises primarily from differences between the local clocks at the master and slave nodes. Electromagnetic interference can also contribute to jitter in a real-world network. This jitter is transmitted directly to the path motion, thus for accurate path motion, it is necessary to have minimal jitter. For a single axis, jitter can result in erratic control behavior such as variation of velocity or oscillatory final position error. For multi-axis systems, the results are more severe.

Skew is a constant delay of a data transmission between transmitter and receiver or between network nodes. It's caused by the propagation delay of the cable (approx. 1µs/100m) and delays in internal logic circuitry. For high performance motion control, skew becomes relevant and the network has to be capable of measuring the skew and compensating for it.

Skew introduces a constant phase shift between network nodes. Thus, coordinated axes do not receive a simultaneous set of command values. For example, assume that you want to carry out a fast circular interpolation with two coordinated axes (e.g. material cutting). In this case, one axis receives a sinusoidal and the second axis a co-sinusoidal command profile. Skew (constant phase shift) will mutate the circle into an ellipse as shown in Figure #3. Additional jitter would also add some distortion to the shape.

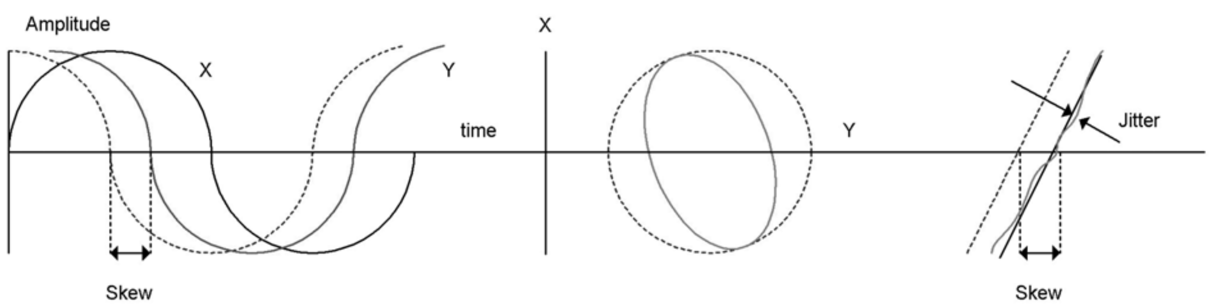

Figure #3, Effects of Jitter and Skew in Motion Control

SynqNet Offers Far Superior Skew & Jitter Reductions to Ethernet
SynqNet limits the jitter to less than 1µs by using phase locked loop techniques to synchronize the independent clocks of each network slave to the network master. This provides superior performance to other non-synchronized networks such as TCP/IP or IP/UDP based Ethernet networks that can cut jitter down to only 20µs using time stamps. Ethernet protocols introduce additional overhead burdens that limit typical cycle times and latencies to 1ms of longer. While this level of performance may be adequate for many general automation applications, it is not adequate for high performance motion control systems.

SynqNet limits skew to 40ns using special algorithms that measure the system skew and compensate for it in hardware. Jitter and skew are guaranteed for any number of nodes or network traffic conditions.

Protocol Efficiency is Essential for Motion Control

Ethernet has been designed to transmit long data packages. A data frame according to the IEEE802.3 specification consists of 28 control bytes and at least 46 data bytes. This protocol is oversized for typical industrial motion applications. Usually the data needs of a node (device) are small (fewer than 46 bytes). To reduce the cycle time and latency, SynqNet has optimized the data frame on layer 2. Instead of at least 74 bytes, a SynqNet frame consists of at least 24 bytes. This is a key advantage of SynqNet over Ethernet, enabling faster and more predictable performance.

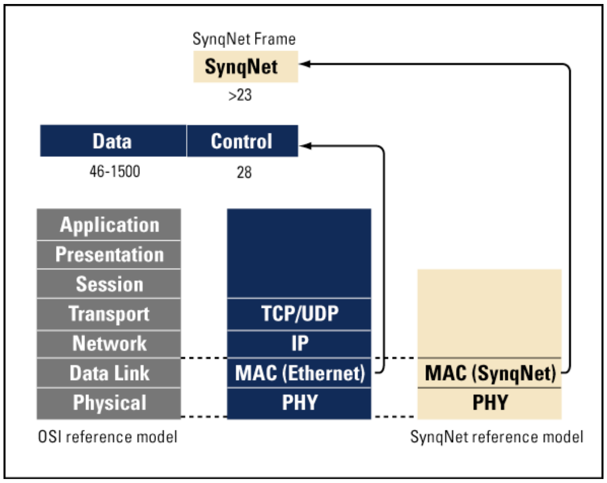

Figure #4, SynqNet Reference Model.

Synchronous Full-Duplex Communications Deliver Deterministic Performance

Standard Ethernet relies upon a single pair of wires to transmit and receive data. Access to the wire is managed by a well-established mechanism known as multiple access collision detect (MACD). As the name suggests, multiple devices on the network try to access the same piece of wire. If two devices try to talk at the same time, a collision occurs and the device stops transmitting, or ‘backs off’, only to retry later after some random time. Such a mechanism is inherently non-deterministic and as the number of devices on the network increases, the collision time increases almost exponentially resulting in a rapid degradation of performance. For office networking, and general automation, the lack of determinism works just fine, however, for most serious motion applications, alternate solutions are required.

SynqNet was designed to eliminate the MACD mechanism. It uses a synchronous method (hence the name) to transmit data on a regular time-scheduled manner to every device. Independent receive and transmit wire pairs (full-duplex) are used to avoid data collisions, and deliver a deterministic date rate of 2 x 100Mbit. The result is cycle times as short as 25µs for 4 axes. In addition, SynqNet has a configurable packet structure that allows for cycle times as low as 10µs.

II. SAFETY AND RELIABILITY

“Self-Healing” Fault Tolerant Operation Increases Safety and Availability

SynqNet can be configured in either a string or a ring topology. The ring topology offers convenient wiring and tolerance to cable breaks within a SynqNet system. “Self-Healing” fault tolerance refers to an ability to operate after an actual cable break, loose connection, or complete fault of any node or nodes. As an example, if two out of five nodes fail, SynqNet is still able to control the remaining three nodes, flag the application, and then execute alternative motion parameters. A closed ring ensures that there is always a redundant data path for transmitted data through the entire ring. SynqNet uses this redundant path as a secondary data channel.

In the event that a wiring segment fails, SynqNet hardware re-routes the data path within two servo cycles and the network connection remains available (Figure #5). At the same time, the application will be informed about the event and event location, allowing the machine to respond in a manner appropriate to the specific situation. For example, a machine can be programmed to finish a move sequence that would otherwise cause a dangerous or expensive collision of independent or interlocked machine axes.

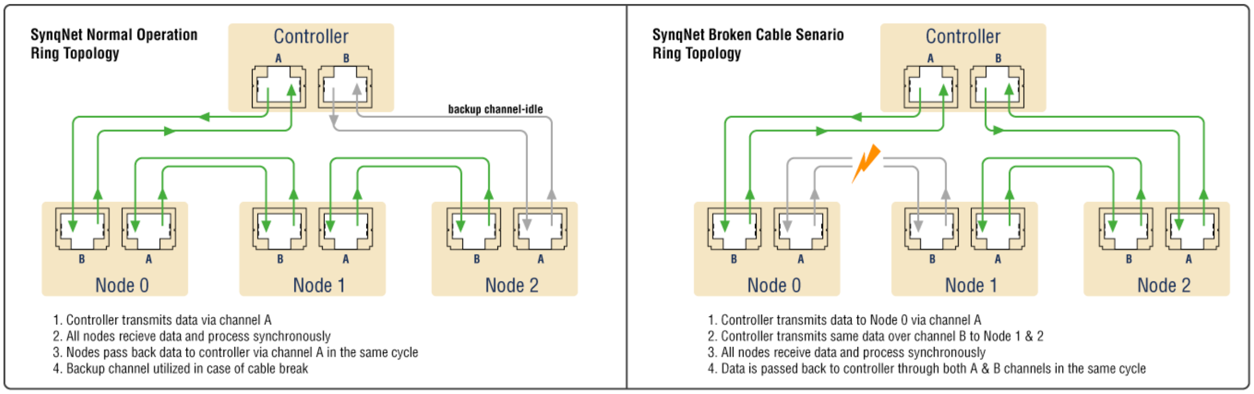

Figure #5, SynqNet Fault Tolerant Topology

In addition, each node has its own watchdog timer so even if the host or whole network fails, each node can react in a predictable and safe manner for a smooth and controlled shut down. To predict possible network failures, SynqNet includes transmission error counters at every node. Any abnormal increase in error count can be used to alert the application software and localize the potential problem before it becomes a catastrophic failure. SynqNet utilizes 100BaseT CRC error checking.

SynqNet is the only commercially available network that offers this level of safety and reliability. Fault tolerance is important in all digital network systems, especially high value applications and medical applications where redundancy is crucial for cost savings and/or safety.

Electrical Isolation of SynqNet (100BaseT) Delivers Solid Reliability

Both Ethernet-based networks and 1394-based networks are designed to accommodate a large number of nodes. When nodes are distributed around a machine or plant, they are often referenced to different ground domains, which introduces ground noise and circulating currents. As an example, IEEE 1394 cables provide a dc connection between these grounds and can cause a ground loop and ground currents will adversely affect IEEE 1394-network reliability. Effects include degradation of data signals and excessive EMI from the cable, which translate to erroneous, possible dangerous motion performance, or system shut down. If the ground currents are high enough, system components can be damaged, as well as create personal shock hazards. The IEEE 1394 network is designed to source and/or sink power to/from remote nodes to allow nodes that do not have their own power to function on the network. This feature coupled with the high-speed signaling rate required in an IEEE 1394 system makes dc cable isolation unfeasible.

In contrast, industry-standard network systems like 100BaseT (IEEE802.3) and others employ dc cable isolation using transformers or optical couplers. Since SynqNet is based on 100BaseT, the EMI problems inherent in 1394 networks are avoided.

III. EASE OF USE

Simplicity is an Economic Necessity for Motion Networks

Networks are conceptually simple, designed to transport data between devices. But the mechanics of transporting data in the real world, reliably, safely, and in a timely and synchronous fashion demand some very complex underlying technology. Well-designed networks should free the user from this complexity, allowing them to focus on system design and machine performance. SynqNet was designed with the machine builder in mind. Setup and configuration are simplified using techniques such as auto discovery of network devices, common tuning, and reporting methods.

Interoperability Defined By Software

The interoperability of networks is often misunderstood and misrepresented. For example, the IEEE 1394 standard defines an interface at the network device driver level. It does not define the software interface to a motion control application and today no 1394-automation standard exists to resolve multi-vendor interoperability problems. As a result, 1394 is available from multiple vendors, yet there is no common software API, making multi-vendor interoperability impractical if not impossible. The machine designer is effectively locked into a specific vendor offering closed 1394 drives and controls. SERCOS adopts a different approach, using some standard or mandatory elements, and some proprietary elements. As a result, to achieve true interoperability demands intimate knowledge of the proprietary elements of every networked component.

In contrast, SynqNet was designed with a common software API for all network devices, from multiple vendors. SynqNet products from leading US and Japanese suppliers offer both standard and custom motion products. The API is available as a set of powerful C/C++ or ActiveX motion libraries.

Powerful Network-Ready Tools

SynqNet tools are designed to work with networked motion systems that contain components from multiple vendors. Tools for real-time data graphing, network configuration and management, mechanical characterization, and optimization are available for Windows platforms, and can also be used across any TCP/IP socket connection.

Centralized Software Configuration Control using Firmware Download

SynqNet has the ability to interrogate firmware revisions and to perform firmware downloads to every device on the network. This simplifies the process of configuration management of software, firmware, even FPGA images, and provides an efficient method for implementing machine upgrade packages, or installing spare components of unknown configuration status in the field.

Remote Diagnostics

The availability for real-time node information enables predictive maintenance, remote diagnosis and repair regimes to be supported. For example, if the node is a SynqNet amplifier, parameters such as temperature, fault and warning conditions, configuration, drive motor, and encoder information can all be accessed remotely and in real-time by the user application.

Availability & Flexibility

The key components of a motion system include the motion processor, drives and motors, load, feedback, and the I/O. SynqNet is supported by a growing number of servo drive vendors, offering standard and custom products, single and multi-axis, ranging from 10W to 10KW. A wide choice of components provides design flexibility and ensures competitive pricing from the world’s leading drive and motor vendors. In addition, a wide range modular and custom I/O solutions are available for a range of motion application types. A bridge device is also available from multiple vendors that allow any analog ±10V type drive to be integrated onto a SynqNet network. This allows for added flexibility in machine design and I/O integration.

SynqNet is Built on Open & Cost Effective Technology

SynqNet is built upon openly available industry standard silicon, making it cost effective and simple for OEMs and drive vendors to embed the technology. No ASICs are required. PHYs and Dual-PHYs are available from numerous suppliers. The SynqNet Media Access Controller (SQMAC) requires an FPGA of only moderate size. SynqNet developer kits are available for drives, I/O and motion specific interfaces. These include reference designs and licensed FPGA images.

Network Specifications and Comparison

	SynqNet®	Ethernet (TCP/IP) (UDP)	FireWire IEEE1394b	SERCOS	Powerlink	Profinet IRT
Motion Control Requirements	Hardware Controller	RTOS + NIC Card	RTOS + 1394b card	Hardware Controller	RTOS + Special NIC or Hardware Controller	Hardware Controller
EMI Immunity	High Transformer Isolation	High Transformer Isolation	High Transformer Isolation	High Fiber Optic	High Transformer Isolation	High Transformer Isolation
Max. Inter-Node Distance	100m	100m	100m	40m	100m	100m
Min. Cycle Time Latency	<25µs	1..2ms	125..250µs	250µs	400µs	1ms
Bandwidth	100 +100 Mbit/s	100 Mbit/s	800 Mbit/s	16 Mbit/s	100 Mbit/s	100 Mbit/s
Transfer Mode	Full-Duplex	Half-Duplex	Half-Duplex	Half-Duplex	Half-Duplex	Half-Duplex
Max. Jitter	<1µs	20µs	<1µs	1µs	1µs	1µs
Fault Tolerance	Yes	No	No	No	No	No
Timing Model	Synchronous Cyclic	None	Isochronous Channel	Cyclic	Mixed Polling / Time Slicing	Isochronous Channel
Real-Time OS Required†	No	Yes	Yes	No	Yes	No
Fast I/O Capture & Time/Position Interpolation	Yes	No	No	No	No	No
Network Programming Required	No	Varies	Yes	Yes	No	No
Current Status	Shipping Globally Since 2002	Available	1394b specification complete (no commercially available motion control solutions to date)	SERCOS III still under development	Available since 2003	First products available mid 2005
Next Generation	500Mbit/s + 500Mbit/s (Full-Duplex)	500Mbit/s + 500Mbit/s (Full-Duplex)	-	100 Mbit/s	Powerlink Version II	-

RTOS – Real Time Operating System
† For given minimum cycle times given

Figure #6, SynqNet Specifications and Comparison

IV.SUMMARY

All networks are not created equal. High performance motion networks demand tightly managed timing regimes to ensure synchronous and real-time updates across multiple axes. While Ethernet and other Ethernet variants offer adequate performance for general purpose applications that use distributed control, it is generally too slow for more demanding situations. In these cases, a fast synchronous network is required to connect a centralized motion processor to multiple servo axes.

SynqNet was specifically designed to support high performance centralized control systems, and offers additional benefits including fault tolerance, simple discovery-based configuration, and high noise immunity. In addition, SynqNet is supported by multiple drive vendors who deliver a wide array of cost competitive products. All things considered, SynqNet offers a unique performance/price advantage, which is why it is gaining fast adoption by industry leaders and OEM customers.