THE CODEX 6000 SERIES OF INTELLIGENT NETWORK PROCESSORS

G. David Forney, Jr. and James E. Vander Mey
Codex Corporation
15 Riverdale Avenue
Newton, Massachusetts 02195

Introduction

Codex Corporation has recently announced a product family, the 6000 Series of Intelligent Network Processors, intended as building blocks for intelligent data networks. In response to the editor's request for articles describing "nifty new products", we summarize briefly in this note the functional characteristics, architecture, and typical applications of these devices.

As is well known, mini- and microprocessors are getting cheaper and better, which has led to the distribution of intelligence from the classical central CPU to specialized devices at dispersed locations. This general trend includes both the dispersion of applications processing to intelligent terminals, cluster controllers, and the like, and the emergence of intelligence within the communications subsystem itself in the form of such devices as intelligent multiplexers, concentrators, packet switches, and front-end communications processors.

Not infrequently communications and applications intelligence are found within the same physical device, such as a typical present-day mini-based concentrator. However, there are good reasons for breaking out the communications intelligence into a separate class of devices, including:

* An intermix of terminals/applications of differing manufacture, age, and protocol can be supported, thus achieving economies of scale in communications facility usage without requiring homogeneous equipment.

* Since communications requirements are much less variable than applications requirements, economies of manufacturing scale can be achieved through prepackaged devices that can be used in a large number of networks and that require no customization by individual users.

* Devices specialized to the communications function can be designed to provide functions that would be impractical to achieve in more general-purpose machines.

The 6000 Series was designed to provide communications intelligence in a modular, turnkey, low-cost package. In its simplest configuration, it can serve as an efficient statistical multiplexer with error

control between two points, at data rates as low as 2400 bps. More advanced configurations are suitable for multinode networks at data rates up to 56 Kbps. The principal applications are expected to be to private data networks whose links are predominantly 9600 bps or less; however, these devices are also suitable for public networks of similar size or as feeder nodes for large common-user networks.

Functional Characteristics

The initial products in the 6000 Series (Models 6030 and 6040) are designed to provide transparent data networks interconnecting heterogeneous collections of terminals and computers. Transparency means that terminals and computers hooked into a 6000 network appear to be interconnected by direct communications links, and require no special programming or interfacing. All external interfaces are EIA or similar standard types. The communications discipline can either be asynchronous (start-stop) or character-oriented synchronous (e.g., IBM Binary Synchronous), which covers practically all terminals now installed.

Within the network, statistical multiplexing and data compression are used to minimize data bandwidth and thus communications facility costs. In addition, the network assures end-to-end error control. Providing these capabilities without excessive overhead and with delays comparable to those in a TDM (e.g., $\leq$50 msec) at line speeds as low as 2400 bps required development of a new intranetwork link protocol, some of whose features are:

* Characters from all active sources are assembled into a common frame, rather than into individual frames as in packet-switched networks. Thus buffering delays are of the order of a frame length (typically, 30 msec), rather than a message length.

* A full-duplex go-back-N (typically, N=7) ARQ scheme, similar to HDLC/ADCCP/SDLC, is used for error control with an explicit NAK to minimize retransmissions. Only 20 bits of overhead per frame are used, including a 16-bit cyclic redundancy check (CRC), in comparison to a minimum of 40 bits in HDLC.

* All characters are internally translated into variable-length codewords by tables assigned to each port. At minimum, therefore, only the information bits from each port need be sent (omitting start, stop, and parity bits); if something is known about character statistics, variable-length prefix (Huffman) codes can be used to further reduce the average number of bits per character.

A variety of communications options such as automatic speed recognition, automatic echo, automatic call answering, special interfaces, and the like are offered.

Finally, the 6000 Series is designed to facilitate centralized network management and control. From a central site, an operator can monitor conditions anywhere in the network, receive alarm messages, initiate diagnostics such as loopbacks and sending of test messages, and reconfigure the network topology, port characteristics, and so forth. The same capabilities can be put under computer control by designation of any 6000 port as a control port.

Hardware Architecture

The 6000 Series utilizes a multi-microprocessor architecture in order to match processing power to throughput requirements at network nodes of varying sizes. The processors are symmetric, and any number from one up to the physical limit of eight can be used with no software modification. The processor modules are based on the Motorola 6800 8-bit NMOS LSI microprocessor, which operates at a 1 MHz cycle rate. Each processor module adds 4800 to 9600 bps of full-duplex throughput capacity, depending on the configuration and amount of data compression.

In addition to the microprocessor modules, the mainframe contains a bipolar microprogrammed microcontroller based on the Intel 3000 Series bit-slice microprocessor LSI chips, which operates at a 6 MHz rate. This microcontroller executes "superinstructions" which implement: a task dispatching scheme which minimizes interrupts and context-switching; a bus contention algorithm which minimizes conflicts on the mainframe bus; I/O transfers between the mainframe and the ports; and other utility functions.

These processors communicate with each other via a global memory, which is provided in 8K and 16K byte modules. ROM (program) and RAM (data buffer) memory may be freely intermixed up to a maximum of 96K in Model 6040. The RAM memories use 4K NMOS chips and achieve effective cycle times of 125 nsec by automatic 4-way interleaving. A common synchronized mainframe bus links the processors and memories.

The external I/O ports are implemented two-to-a-card using Western Digital ASTRO communications controller LSI chips. Up to 32 ports can be supported in a 7" high, 19" wide port nest. The intranetwork communication protocol is supported by a special port module, which performs in hardware functions like flag insertion/detection, zero stuffing/stripping, CRC generation/checking, and channel buffering. The ports interface to the mainframe on a separate I/O bus via the microcontroller.

An optional operator's console provides a 32 character alphanumeric display and 18-button keyboard, as well as system indicators. It can be used for hardware diagnostics as well as network management and monitoring.

A more complete description of the architecture appears in [1].

Comparison to Alternatives

The data network building blocks now in use that are most comparable to the 6000 Series products are time division multiplexers (TDMs) and concentrators.

One obvious and fundamental comparison is economic: typical 6000 prices fall between TDM and concentrator prices, roughly (very) twice a conventional TDM and half a concentrator for the same number of ports.

TDMs have typically been used to reduce line costs while maintaining transparency. The statistical multiplexing and data compression in the 6000 typically yield two to four times the efficiency of TDMs (very few dedicated channels are used more than half the time). This can be translated into substantial savings in line costs whenever a TDM cannot carry all offered traffic over a single analog line (i.e., when aggregate traffic exceeds 9600 bps), or whenever line tariffs are quasi-linear with data rate, as is currently true in Bell's DDS. In addition, the 6000 processors offer error control, multinodal configurations, centralized network management, growth capabilities and other features which may justify their use even when line cost savings alone do not.

Concentrators have typically been implemented on general-purpose minicomputers, augmented with communications software packages and hardware. They are invariably message-oriented and therefore must be custom-programmed for each application in order to properly recognize and process messages, whereas in the 6000 no user programming is required. Like the 6000, they typically provide error control and statistical multiplexing; sometimes they perform some byte-oriented data compression. Generally concentrators use more overhead per frame and bits per character than the 6000, but they may be able to make up the difference in line efficiency if they provide enough localized applications intelligence to reduce significantly the total volume of data exchanged with the central processor. In summary, a concentrator costs more and generally has less communications intelligence than a 6000, but may be able to overcome these disadvantages by providing applications intelligence, with additional costs for customized programming, maintenance, etc.

Finally, readers of this Review will want to know how a 6000 network compares to a packet-switched network such as ARPANET. First, it is necessary to mention that as of this writing no switching capability is offered in the 6000 family, unless its ability to be reconfigured in seconds by the operator could be considered a sort of manual circuit-switching capability. The current structure is, however, amenable to expansion to support circuit switching with automatic adaptive routing, to provide a 'virtual circuit' service. When these capabilities are available, it seems to us that the 6000 approach will be providing the flexibility and reliability that are the central virtues of packet-switching networks without their requirements for wideband 50 Kbps lines (to keep delays acceptable), their excessive overhead, and their requirement for

special interfaces. We would expect that for smaller networks the 6000 approach would be the economically attractive one, while for large networks that can reasonably utilize wideband links a synthesis of the two approaches will emerge.

Conclusion

The 6000 Series processors represent a new type of data network building block that currently provides the transparency of a TDM with the efficiency, error control, and network capabilities of a concentrator, at a price intermediate between them. In the next few years we expect to see the emergence of a variety of devices of this type with continually expanding capabilities.

Reference

1. J. E. Vander Mey and L. J. Krakauer, "The Architecture of a Multiple-Microprocessor Communications Processor," 1976 Zurich Conference on Data Communications, March 1976.