

November 1978

**Dual-port RAM Hikes
Throughput in Input-output
Controller Board**

**Craig Kimmie and Michael Maertz
Electronics**

Dual-port RAM hikes throughput in input/output controller board

On-board random-access memory, accessible from system bus, makes input/output controller subsystem look like just another memory board to the host microprocessor

by Craig Kinnie and Michael Maerz, Intel Corp., Santa Clara, Calif.

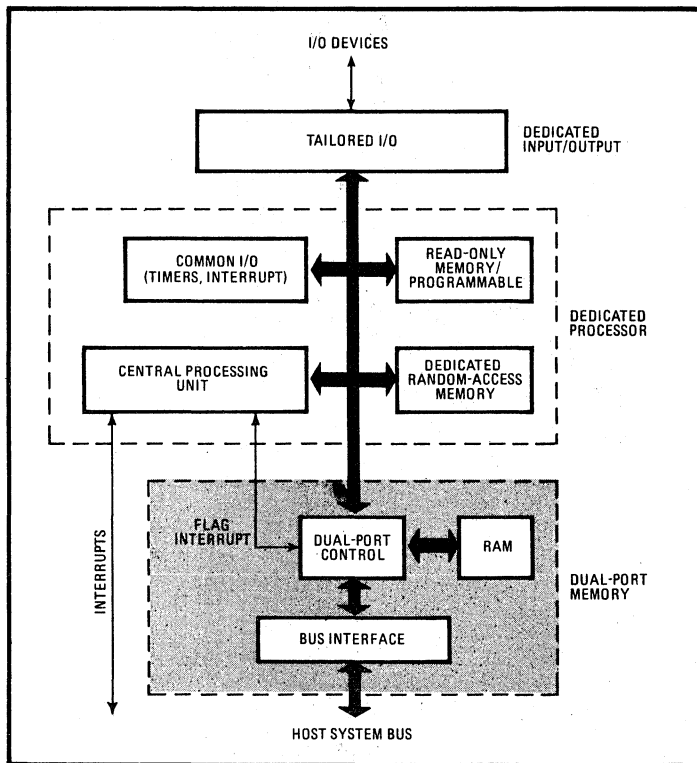
□ Input/output controllers based on microprocessors step up throughput in microcomputer systems by relieving the host processor of tedious, time-consuming control tasks—and a new design concept that increases the processing capability of this subsystem promises to hike throughput even more. It will cut the host intervention needed to transfer data and to run the controller.

In this configuration, all communications between the host processor and the controller are handled through a section of dual-port memory that resides in the controller subsystem. This setup allows more efficient transfer of large blocks of data from the I/O device to the system without contention over access to the system bus. It also

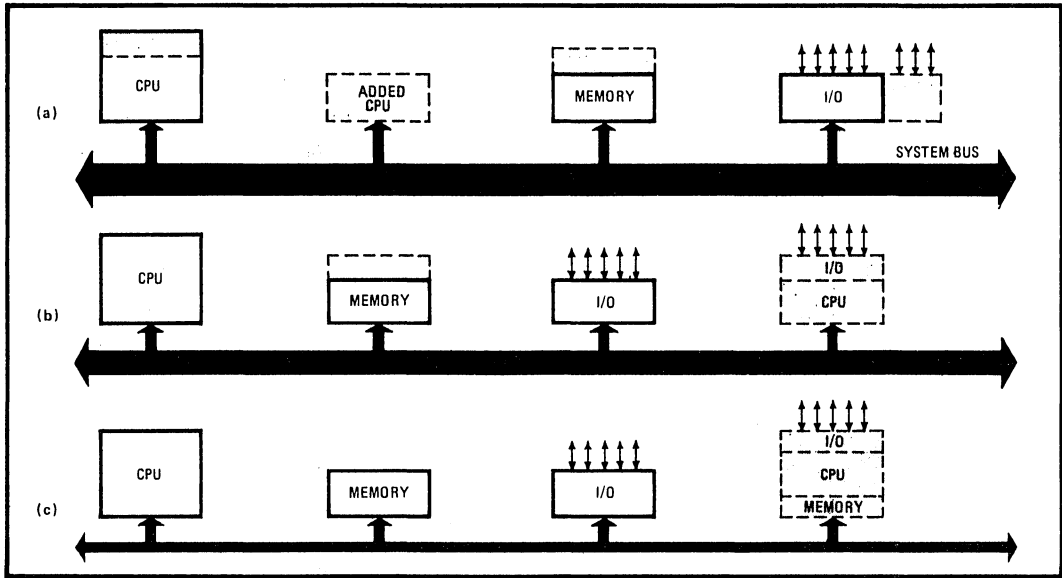
simplifies interprocessor communications because the subsystem controller appears to the host processor simply as an additional RAM board.

Although this concept allows the subsystem to remain dedicated to its I/O control function and to assume a subservient role to the host processor, it has more processing power than previous generations of such controllers. Hence it has been dubbed the intelligent-slave concept by Intel, which applies it in the iSBC 544 intelligent communications controller.

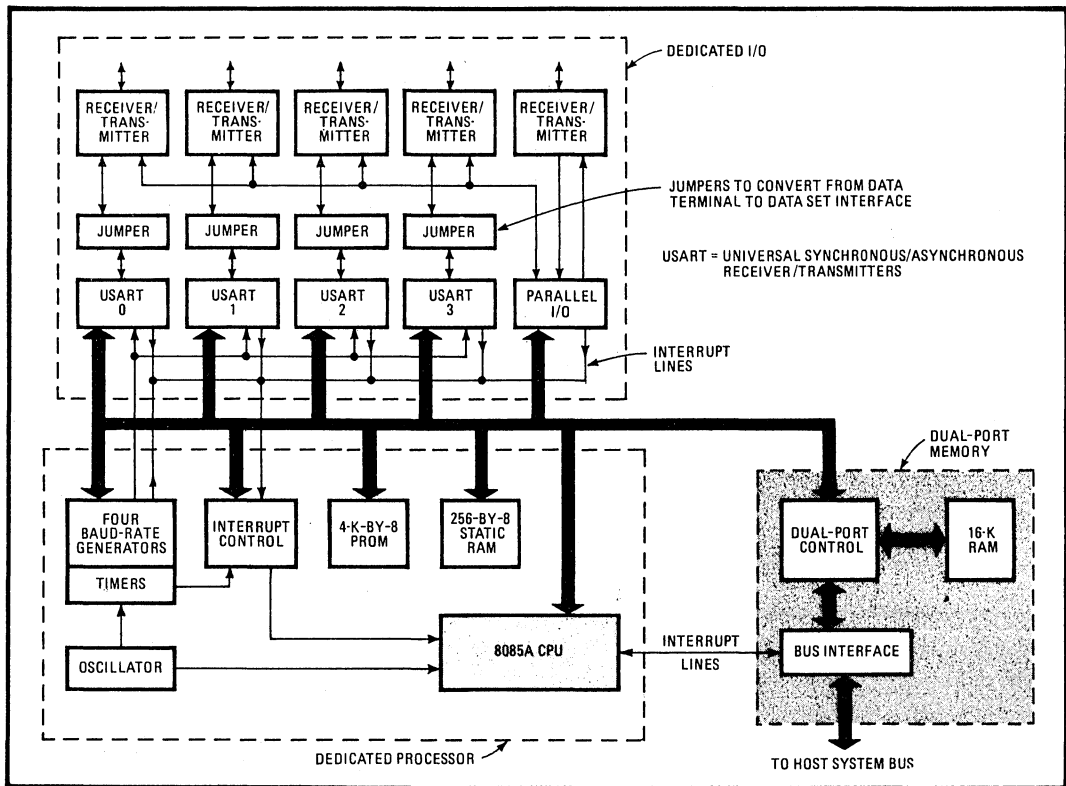
The new subsystem architecture is divided into three major sections: dedicated I/O, dedicated computer, and dual-port memory (Fig. 1). The dedicated input/output,



1. Heart of memory. New controller architecture includes the dedicated input/output circuitry and dedicated processor of an intelligent peripheral controller, but its heart is the dual-port random-access memory.



2. Performance advantages. In adding a real-time task to an existing real-time system, the load on the system bus is significantly reduced over the traditional multitasking approach (a) or the intelligent controller approach (b) by the intelligent-slave controller approach (c).



3. The 544. Based on the 8085A microprocessor with 4 kilobytes of PROM and 16 kilobytes of RAM, the subsystem is designed as a communications controller with four synchronous/asynchronous buffered serial I/O channels, and a 10-bit parallel I/O interface.

Dual-port RAM also shows up in new single-board computer

The concept of a dual-port read-write memory used in the iSBC-544 communications board is also employed in another new Intel product: its latest single-board computer, the iSBC-80/30. A dual port makes the 80/30's random-access memory directly accessible by the on-board 8085A central processing unit via internal busing without tying up the external system bus, the Multibus. At the same time, it also makes the RAM directly accessible by any other boards, like direct-memory-access controllers or other one-board computers that may be tied to the Multibus.

Moreover, the 80/30 adds its dual-port bus to the earlier iSBC computers' pair of buses: an internal bus, which hooks the CPU to peripheral chips and read-only-memory program storage and the system bus, over which the CPU and other boards communicated with RAM. Eight bits wide, the new bus is connected to a pair of buffer registers that coordinate, thus making the RAM accessible either by the internal bus or the system bus.

The objective is throughput: the CPU has priority in access to the on-board RAM. But since the access is not over the Multibus system bus, which might be tied up, there is no waiting. From the viewpoint of other system boards, the system bus is accessible a greater percentage of the time.

With the incorporation of 16 kilobytes of memory on the 80/30, Intel had little choice but to move to the dual-port, triple-bus architecture. The reason is that few system designs require more than 16 kilobytes, so in many applications all boards will be demanding access to the 80/30's memory over the Multibus. The CPU had better have priority to its RAM, through its own private line, lest the queue for the system bus bog down throughput.

The 80/30 also packs lots of extras, in addition to the total 16,384 bytes of read/write memory built with 2116 16-K dynamic RAMs. A pair of ROM sockets provide 4,096 bytes of program storage if fitted with 16-K erasable programmable read-only memories like the 2716. When pin-compatible 32-K erasable PROMs are available, program storage can be extended to 8 kilobytes.

Also on board is a socket for Intel's universal peripheral interface chip, the 8041 (or 8741 erasable-PROM version), which can function as a slave processor to drive peripheral devices. An 8251A universal synchronous/asynchronous receiver/transmitter is included for serial communications, and the 80/30 also boasts three 16-bit programmable timers. The 24 programmable input/output lines are brought out to sockets that accept quad line-drivers or -terminators for interfacing.

Ray Capece

consisting of the necessary peripheral chips, timers, buffers, and interface integrated circuits, tailors the controller to the application's I/O requirements.

The dedicated computer consists of a general-purpose microprocessor, electrically programmable read-only memory, dedicated RAM, timers, interrupt logic, and the decode and chip-select logic. The size and speed of the central-processing unit can be tailored to match the requirements of the dedicated I/O section.

The dual-port memory is the heart of the architecture and sets it apart from traditional approaches to intelligent controllers and multiprocessing. Passing all commands and data between the system and the controller's processor through this memory offers a number of significant advantages.

First, the dedicated computer's performance can be optimized for its applications. Its software always operates at full speed, since all required memory and I/O resources are immediately accessible on the board, without indeterminate delays caused by other system activity on the bus. This accessibility is especially important in real-time systems, since it allows the controller's performance to remain constant even though system bus activity may change.

Secondly, the architecture presents a consistent and convenient interface between the host CPU and all the controllers in the system, regardless of function. Because the controllers' dual-port RAM looks to the host CPU like just another location in system memory, the hardware and software problems associated with connecting multiple processors together are reduced to interfacing a number of identical intelligent memory locations.

Also, the architecture offers a degree of protection for the data in memory. The subsystem computer and soft-

ware can only alter that portion of system memory that resides in its own dual-port memory section. In contrast, traditional intelligent controllers have access to the entire system RAM and, should a malfunction occur, can destroy all of that memory.

System performance advantages


Because all processing assigned to the new controller's CPU takes place off the system bus, its architecture offers important performance advantages to the system. These advantages come from the appearance of the processed data blocks in system memory without consuming any system resources or bus time.

The advantages of this approach are best demonstrated by comparing it to alternative means of adding a real-time task to an existing real-time system. In this case, the new task requires additional CPU, memory, and I/O resources.

The traditional multiprocessor approach of Fig. 2a expands CPU resources in one of two ways: software utilization of reserve capacity in the existing processor, or adding another processor. In either case, memory and I/O increments generally will be required.

The primary disadvantages of this approach are the increased complexity of the system software and the increased load on the system bus. Both will slow the existing real-time system unless it has been designed with adequate reserve. The system bus must also provide sufficient capacity for the incremental memory-execution and data-transfer operations. This additional bus load will also require that the primary real-time task can tolerate CPU delays due to bus contention.

The intelligent-controller approach of Fig. 2b has gained widespread use since the advent of the micropro-



4. Memory mapping. The variable system memory addresses are always mapped into the on-board address of 8000HEX, providing software independence for the subsystem and the host.

processor. This approach combines the CPU and I/O increments onto a single module that usually includes direct-memory-access transfer logic. In some cases the execution memory for the CPU is included.

This approach lessens the bus-loading problem since the I/O data transfers and some memory-execution cycles take place off the system bus. However, both CPUs' programs will have to tolerate delays caused by increased bus contention. Increased software sophistication is the primary disadvantage of this approach, much as with the multiprocessing approach of Fig. 2a.

The intelligent-slave approach of Fig. 2c can be viewed as a logical extension to the intelligent-controller approach. Combining the CPU, I/O, and memory increments creates a single module that has a minimal impact on the existing system software and bus loading. What's more, the subsystem can operate at full capacity without regard to other system activity. It can be programmed outside the primary system and then added with minimal impact on the system software or performance.

A limitation of the approach is the inability of the subsystem to transfer data into portions of the system memory space that reside off its board. This problem is minimized by the ability of the controller's RAM to serve as a substantial portion of the entire memory space addressable by the system. In this light, the on-board processor can be viewed as having a DMA capability limited to a portion of the system's address space.

In a system with more than one of the new controllers, the system CPU handles any data that must be trans-

TABLE: EIA RS 232-C SIGNALS PROVIDED AND SUPPORTED

Carrier detect	Receive clock
Clear to send	Receive data
Data set ready	Ring indicator
Data terminal ready	Transmit clock
Request to send	Transmit data

ferred from one to another. Applications involving the transfer of large blocks of data would be best served by a central block-transfer device elsewhere on the bus.

The advantages offered by the new approach in this example of adding onto an existing system are just as applicable to a ground-up design. This modular approach to configuring real-time multiprocessing systems simplifies hardware and software design, as well as system integration.

While the primary design objective of the new architecture is operation in a multiprocessing system, it can provide significant utility as stand-alone processors. Thus these controllers incorporate a second mode of operation called the limited bus-master mode.

In this mode of operation the new controller can be used like a single-board computer as long as it is the system's only master of the bus. It can be connected to standard memory or I/O expansion boards to enhance its capability. It can even be used to drive other such controllers as long as they are used in the subsystem mode. This dual operational mode will allow the new controllers to serve a broad range of applications.

Communications first

Communications applications present complex processing requirements and an inherent real-time nature, so it is logical that a communications processor be the first of these new controllers to be marketed. The iSBC 544 intelligent communications controller can serve as a flexible front end to an iSBC system or as a cost-effective stand-alone processor configured as a terminal cluster or line concentrator. Its design (Fig. 3) incorporates an 8085A CPU, 16 kilobytes of dual-ported dynamic RAM, 4 kilobytes of PROM, programmable interrupt control, three interval timers, four programmable baud-rate generators, four synchronous/asynchronous buffered serial I/O channels, and a 10-bit parallel interface compatible with a Bell 801 automatic calling unit.

The dual-port memory block basically consists of the 16-kilobyte bank of random-access memory, which is accessible from either the system bus or the on-board processor through the dual-port controller. This memory block provides the primary means of communication between the system and the on-board 8085A. The port to the memory, which looks to the system bus like any other RAM card belonging to the system, features full 20-bit

addressing and a typical access time of 600 nanoseconds.

The interface's address-decode logic allows switching of the base address of the iSBC 544 to any 4-kilobyte boundary in the host system's address space. In addition, the user may reserve 8, 12, or 16 kilobits of the 544's memory for use by the on-board processor only. This reserved memory is not accessible from the system bus and does not occupy any system address space. The only restriction is that all of the unreserved memory reside in the same 64-K address page of the system memory.

This memory division can be a significant advantage in large 8-bit microcomputer systems. Only that portion of the controllers' memory needed to pass data between CPUs must be made accessible to the system. The remaining buffer and execution memory does not consume any system address space.

The net result is an increase in the system's overall memory capacity. For example, a microcomputer system that would usually be limited to 64 kilobytes of memory has a total memory capacity of over 190 kilobytes when driving seven 544s.

Address maps and interrupts

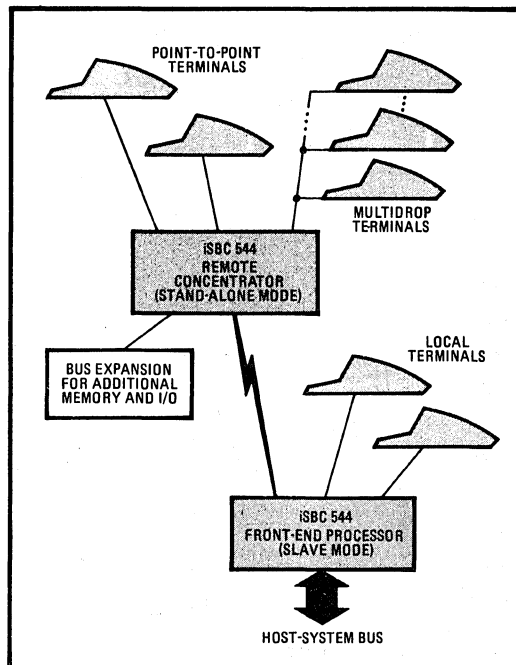
To the on-board processor, the base address of its memory is fixed at 8000HEX. Furthermore, all on-board addresses are fixed, so that multiple 544s operating on the same system bus can be running identical programs regardless of their base address on that bus. This capability necessitated the address-mapping logic to transform addresses from the system bus into the equivalent in the on-board address space starting at location 8000HEX (Fig. 4).

The address-mapping logic also implements the flag-interrupt feature. It provides an interrupt to the on-board processor whenever a byte is written into the 544's base address from the system bus, and a read from the on-board processor to the base address clears the interrupt. Since each 544 in a system has a different base address in that system's RAM, it also has a unique interrupt. This flag-interrupt capability is a key element in establishing a protocol for communications between the host CPU and the subsystems' processors.

The dual-port control logic is responsible for resolving contention over access to the memory and is designed to optimize the performance of the subsystem CPU. Unless the system bus has initiated a memory cycle before the on-board processor requests memory, that CPU runs at full speed. The maximum delay that can be encountered is one memory cycle. The arbitration logic actually reserves the memory for the on-board processor before it generates the necessary commands. This advance reserving guarantees that the on-board CPU will suffer minimum intervention from system bus accesses.

When the iSBC 544 is used in the stand-alone limited bus-master mode, the dual-port logic is disabled and the bus interface buffers are turned around to drive onto the bus. This reversal allows the on-board central-processing unit access to the memory of other subsystems or I/O expansion boards on the system bus.

The dedicated computer is built with an 8085A CPU operating at 2.76 megahertz, between 2 and 4 kilobytes of PROM and ROMs or 8 kilobytes of ROM using 2332



5. Communications applications. Two typical applications of the new iSBC 544 would be as a front-end communications processor to a microcomputer system and as a remote concentrator to a series of point-to-point or multidrop connected terminals.

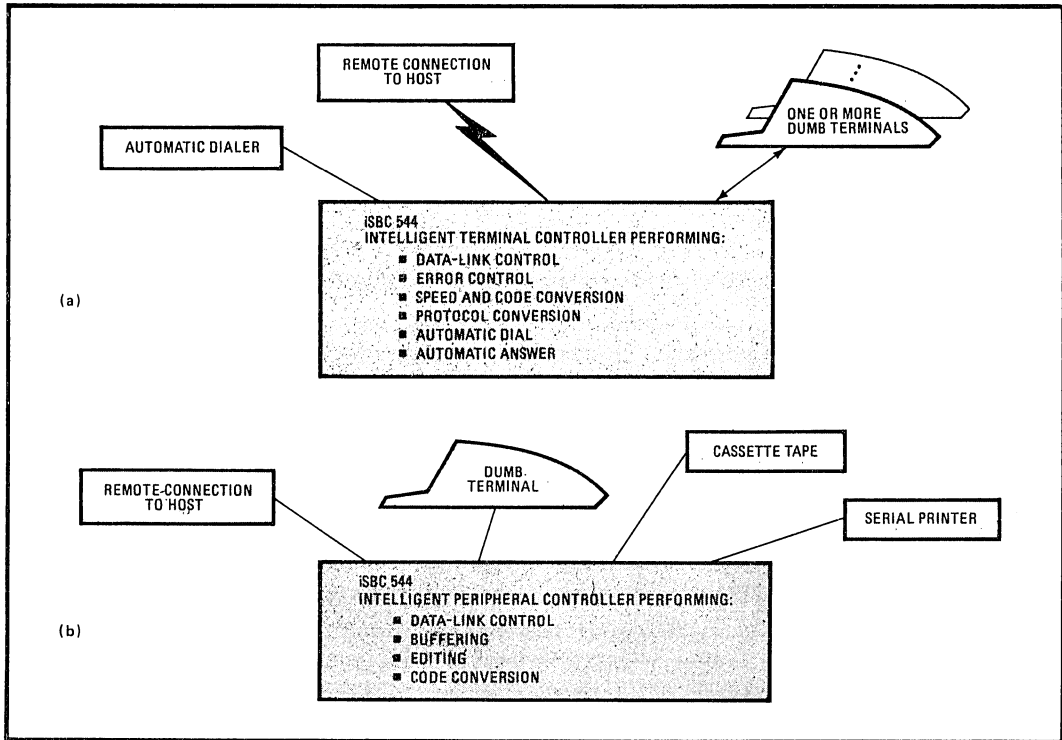
mask-programmable parts, 256 bytes of static RAM, two 16-bit and one 14-bit interval timers, and a 8259 programmable interrupt controller for individual receive or transmit interrupt inputs for each serial port.

Special command-decode logic was added to the CPU to allow it to operate at maximum speed independent of other system activity. There are 21 sources of interrupt on the 544, including the separate transmit and receive interrupts for each port and separate timer interrupts. In addition to receiving an interrupt from the system, the 544 can also send an interrupt to the system bus via the 8085A's serial-output data line.

Since this controller is intended for communications applications, latched interrupts are provided directly to the CPU for loss of carrier and ring indicator for all four I/O ports. The ring-indicator and carrier-detect lines can also be monitored through the parallel port.

Dedicated I/O

The dedicated-I/O section of the 544 provides a high degree of flexibility and programmability. This results primarily from the inclusion of four 8251A universal synchronous/asynchronous receiver/transmitters. These devices are programmable for synchronous or asynchronous mode, character size, parity bits, stop bits, and baud rates. Data, clocks, and control lines are buffered with RS-232-C-compatible drivers and receivers to four 26-pin card-edge connectors. Each port is configured as a data-terminal interface, but may be converted to a



6. From slave to master. In its stand-alone mode, the 544 can operate as a bus master and be configured as an intelligent terminal controller connecting dumb terminals to a data link (a) or as a peripheral controller connecting RS-232-C-compatible units to the terminal (b).

data-set interface by changing a single jumper-plug assembly. The ports support most RS-232-C signals (those that are listed in the table).

A programmable baud-rate generator is also provided for each port. The range of baud rates available is 75 to 56 kilobits per second. The generators are implemented with 8253 programmable interval timers, which receive a jumper-selectable input frequency of 1.84 or 1.23 MHz. In addition, one of the CPU's interval timers can be converted to baud-rate operation and jumpered to any port to provide it with split-speed operation.

The 544 also provides a parallel port with four RS-232-C buffered input lines and six RS-232-C buffered output lines. This port is configured to interface to most automatic calling units but may be used as a general-purpose I/O port. It is implemented with an 8155 programmable peripheral interface that also provides the 256 bytes of static RAM and the 14-bit timer.

Applications

A likely common use of the 544 as a subsystem is as a front-end processor or terminal multiplexer (Fig. 5) in an iSBC system. The 544 performs all communications-related functions such as format control, code conversion, data-link control, error checking, data compression, and protocol management. It can handle multiple protocols, line speeds, and data formats.

All the system processor sees are the processed data

blocks that appear in system memory. An automatic dialer could be added to provide a dial-up connection to a host processor or network.

Also shown in Fig. 5 is another 544 used in its limited bus-master mode as a remote concentrator and terminal controller. The line and memory capacity of the remote concentrator can be increased by the addition of standard iSBC memory and I/O expansion boards.

The intelligent-terminal controller shown in Fig. 6a is a prime example of a 544 used in the stand-alone mode. It can connect one or more dumb terminals to a data link and provide the necessary buffering, code conversion, and data-link control. It could also connect a terminal that happens to communicate in a different protocol to a new network or to more than one network.

The iSBC 544's multiple serial lines do not have to be used for communications. They can also be used to connect RS-232-C-compatible peripherals to the terminal (Fig. 6b). In this configuration, the 544 can provide message editing and formatting, bulk storage, and hard-copy output.

As this last application suggests, the 544 is the vanguard of a family of intelligent I/O controllers that will add tremendous increases in throughput and versatility to the iSBC line of single-board computers. The basic architecture will simplify the task of developing multiprocessing hardware and software solutions that will overcome throughput limitations. □

