



Engines for Global Connectivity

communicator

autumn 1999

From the President

VDSL (Very High Speed Digital Subscriber Line) technology is likely to emerge in the near future as the technology of choice for high-speed access to the home, as well as access within business premises. This is because VDSL promises to bring access speeds of from 13Mb/s to 51Mb/s to the user.

Because of the distance constraint, VDSL requires that the optical fiber connection extends to an ONU (Optical Network Unit) close to the premises, thus shortening the length of the twisted pair connection to the subscriber's premises. The VDSL concept poses a technical challenge for VLSI device manufacturers – particularly for those focusing on physical layer technology. But within a few years, this challenge will be overcome, with VDSL becoming quite

continued on page 2

Data and Buffer Management in ATM Systems

The explosive growth of the Internet and the emergence of multiservice platforms continues to feed demand for ever greater bandwidth. To meet these growing demands, original equipment manufacturers (OEMs) are seeking highly integrated application specific standard products (ASSPs). These integrated silicon solutions are attractive to OEMs from the perspective of both faster time-to-market and product differentiation. As a result, the core of the complexity of communications systems has shifted down to the ASSP vendors.

A significant area of challenge in ASSP design is packet data and buffer management. In this area, ASSP vendors are meeting the rising tide of data flows by developing management schemes that coordinate data transfer using matrices of buffers. Effective optimization of such management schemes requires minimizing the required system memory and processing resources, while simultaneously maximizing the flow of data.

Successful coordination means avoiding both data starvation on

logical channels and buffer overruns, thus maximizing utilization of network bandwidth and minimizing data loss.

A description of the design of one such data and buffer management scheme, the scatter and gather mechanism in an Asynchronous Transfer Mode (ATM) environment, serves to illustrate one solution strategy.

Scatter and Gather

ATM networks offer users unique services, including high-speed interconnectivity, service integration, quality of service (QoS) guarantees, and bandwidth on demand.

The variable-length higher-layer protocol data units (PDUs) are supported over ATM through their mapping to the service data unit (SDU) of the Convergence Sub-layer (CS) of the ATM Adaptation Layer (AAL). The mapping process requires buffers to facilitate protocol processing.

However, the variable length of higher-layer PDUs makes it difficult, in the absence of accommodation, to effectively

continued on page 2

In this issue

**Scatter/
gather
mechanism
aids
resource
management**

**Third
Quarter
financial
report**

Product list

From the President...continued

cost-effective. At that point, VDSL might threaten the growth of ADSL, as well as cable modem deployment.

Large scale deployment of ADSL, cable modem, or VDSL technologies in the subscriber access network will alleviate the bandwidth bottleneck currently experienced in surfing the Internet.

As the deployment of ADSL and cable modem technology in the access network ramps up and consumers begin increasing their bandwidth usage, the demand for even greater bandwidth and services will prompt service providers to undertake investment in infrastructure improvement that will be required for VDSL deployment. This means that newer and higher capacity equipment will be needed in both the access and the transport networks.

ATM is likely to be the protocol of choice for VDSL as it is currently for ADSL, stimulating demand for ATM VLSI devices. Similarly, access and transport infrastructure upgrade will create demand for SONET/SDH VLSI products, with capability for handling multiple protocols – both ATM and IP, for instance. All this bodes well for communications VLSI solutions suppliers.

TranSwitch management is committed to making sure that our customers have access to VLSI products that can handle seamless transition to multiple services and protocols to cope with the demand for higher bandwidth in access and transport networks. ♦

Santanu Das

**Best Wishes to all in 2000
–TranSwitch Corporation**



Congratulations to Sun Microsystems' Manager of Workgroup Services, Dan Schumaker, right, the grand prize winner of the EE TIMES and TranSwitch Communications Challenge. Vice President of Business Development for TranSwitch, Bob Pico, left, presented the grand prize iMac computer workstation to Dan at Sun's Massachusetts offices.

Data and Buffer Management...continued

utilize buffer space in the mapping process as the buffer size has then to be chosen so as to support the largest PDU size possible. The effect of poor buffer-space utilization under certain traffic conditions is buffer over-run, which leads to overall degradation in the QoS offered to traffic.

Therefore, efficient use of buffer space without significantly affecting packet latency becomes a concern for the system designer. The scatter and gather mechanism provides a framework for achieving balance among these parameters.

Figure 1 represents the protocol

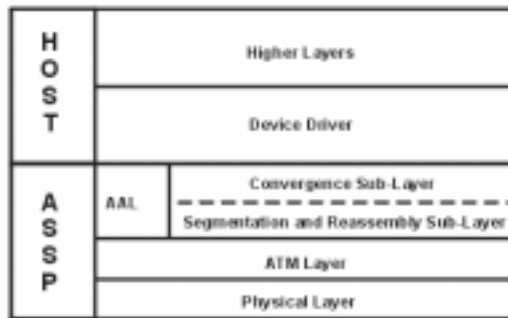


Figure 1. Protocol Layer Model.

layer model relative to the two distinct domains (host and ASSP) considered in this discussion of the scatter and gather mechanism.

The mechanism performs the distribution (scatter) and collection (gather) of data into and out of one or more buffers, as it is transmitted to or received from the network (Figure 2). The buffer space is allocated in the host during system initialization as part of the device driver. Considering the simplest case of homogeneous buffers (i.e., buffers of the same size), the amount of space to be allocated is proportional to the buffer size and

the number of buffers required. The allocated memory is shared between the direct-memory-access capable ASSP and the host to avoid interruptions to the host CPU when data transfer is needed by the ASSP. The buffers are identified by buffer descriptors, which are linked to form lists of FREE buffer pools for




Figure 2. Two-way Scatter/Gather Mechanism.

use in transmit and receive operations.

The mechanism uses buffers on demand during the scatter process and recycles them after the gather process is completed (Figure 3). The host-based buffer manager, a component of the device driver, regulates the use of FREE buffers to provide them to the scatter process, and maintains lists of INUSE buffers to recycle them when the gather process is completed.

Either the host (when transmitting data to network) or the ASSP (when receiving data from network) performs data scattering while the other performs data gathering. The ASSP combines the ATM cell reassembly and the scatter processes when transferring received CS PDUs to the host. To aid in the cell reassembly process, the buffer manager assigns the ASSP a pool of FREE buffers (the receive FREE buffers pool) during system initialization. Once assigned, the ASSP uses this pool in the cell reassembly process. However, it depends on the buffer manager for buffer recycling; that is, used buffers are recycled back to the FREE pool when the host gathers PDUs.

The buffer manager maintains a pool of FREE buffers (the transmit FREE buffers pool) for use by the scatter process at the host to deliver data to the ASSP for network transmission. Buffers that are in use are tracked through virtual circuit (VC) based linked lists and are maintained by the buffer manager. When buffers used are no longer needed by the gather process at the ASSP, they are

returned to the transmit-side FREE buffer pool through the buffer manager.

The time-event diagrams (Figure 4 a,b) presents the procedural interactions between the protocol layers in sending (a) and receiving (b) PDUs via the scatter and gather mechanism.

System Model

In a typical finite buffer scenario, the buffers have to be recycled at an arbitrary rate that prevents latencies

smaller than the PDU size). The host recycles the buffer back to the FREE buffer pool for the device's use after its contents are extracted. The buffers are recycled in the order they are given to host.

The server is characterized busy if PDUs are being reassembled, and idle if all the buffers are occupied. The average buffer recycle rate is typically slower than the average service rate of the server and, therefore, it controls the PDU service rate.

Using the model in Figure 5, the time-domain expressions for the parameters of interest to the

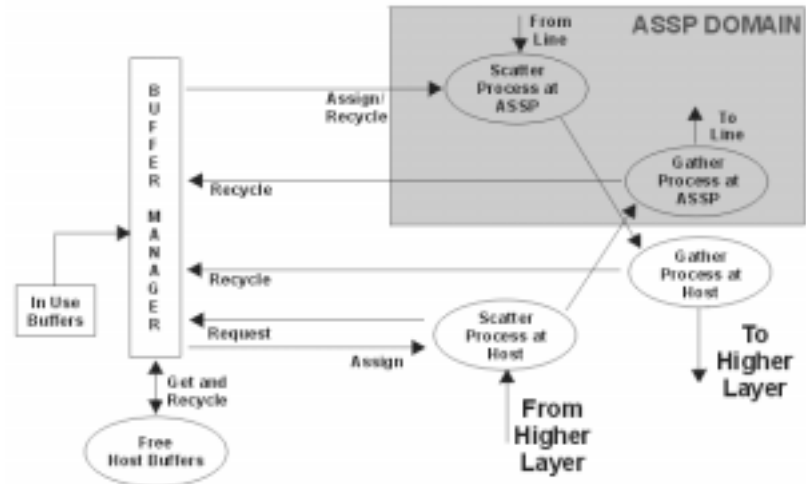


Figure 3. Buffer Management and Data Flow.

and loss. Given the number of buffers allocated (N) and average buffer consumption rate (C buffers/sec), one can arrive at the estimate of required average buffer recycle rate (R buffers/sec). Conversely, given the average buffer recycle and consumption rates, one can estimate the number of buffers required to sustain the arriving data traffic.

Figure 5 models the system in hand for the receive direction. In the host-provided buffers, the server reassembles several PDUs at a time from the interleaved cell stream received from the line. The host is indicated, with a receive, when a buffer is filled with either a complete PDU (case of buffer size greater than the PDU size) or a portion of it (case of buffer size

system designer are formulated. In arriving at the expressions the following parameters were assumed:

- a PDU is small enough to be assembled in one buffer (of a given size);
- all active virtual circuits exhibit the same traffic characteristics.

The number of buffers N to be allocated, to sustain the packet arrivals for t seconds, given the buffer recycle R and consumption C rates, can be computed from the following relation:

$$N \geq (C-R) \cdot t; (C > R).$$

The required buffer recycle rate, to

continued on page 4

Data and Buffer Management...continued

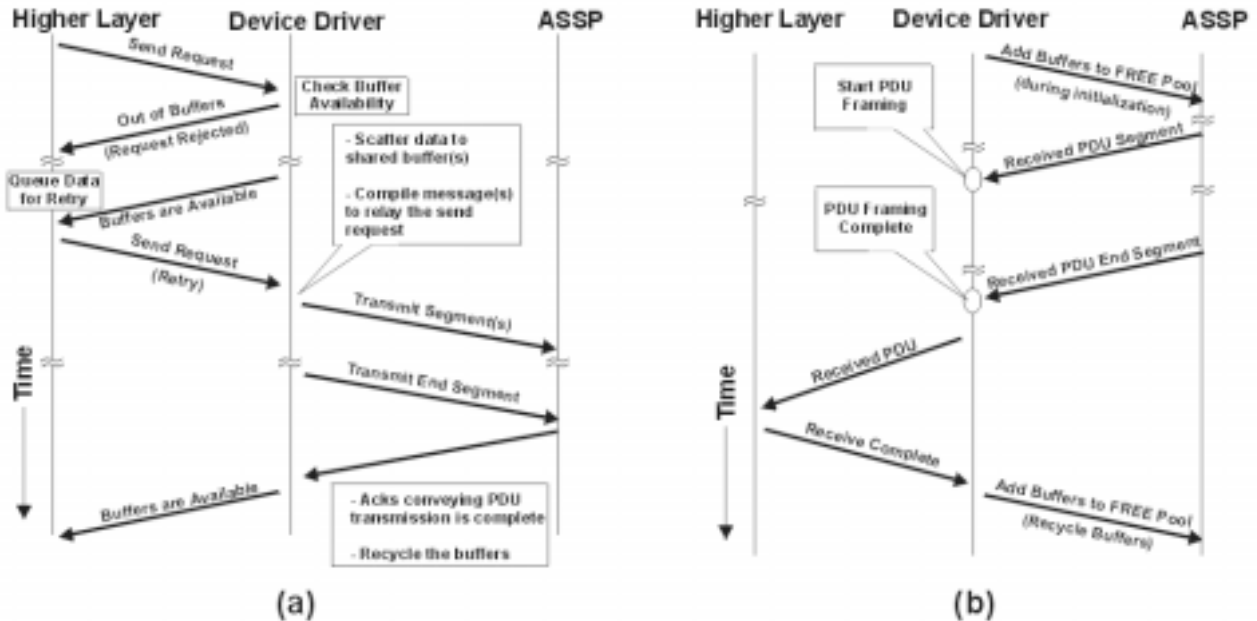


Figure 4. Sending and Receiving PDUs.

sustain the packet arrivals for t seconds given the number of buffers allocated and the buffer consumption rate can be computed from the following relation:

$$R \geq C - (N/t); (C > R).$$

Application in HDLC Environments

Similar mechanisms apply to developing packet data and buffer management in HDLC environments. The example architecture of such an ASSP (Figure 6) includes an embedded RISC processor. The architecture paves the way for the support of any of other HDLC family protocols such as the Point-to-Point Protocol (PPP) or the Link Access Procedures for D-channel (LAPD) within the scope of the device.

The highly integrated and programmable ASSP is capable of performing HDLC protocol processing functions for large number of multi-rate (nx64Kbps) and/or sub-rate (px8Kbps) logical channels and consists of significant features suitable for a range of telecommunication and data-communication applications.

The ASSP interfaces with digital

asynchronous transmission line framers at the line side and a broad spectrum of the higher-layer

space utilization. Both of the scenarios result in retransmissions at the higher layers and, thus, in poor utilization of network bandwidth. The




Figure 5. System Model (Receive Direction).

protocols through the device driver on the host side. The higher-layer protocols that are supported, for instance, are frame relay, IP, SS7, and ML-PPP.

As in the ATM environment, the challenge involved in such an architecture is avoiding data starvation on link channels, in the transmit direction, and buffer over-runs in the receive direction. The data starvation could occur as a result of the PCI bus latencies and leads to rate adaptation buffer under-flows, which in turn leads to aborted HDLC frame transmissions. The buffer over-runs could occur as a result of the poor buffer-

efficient utilization of the buffer space addressing the trade-offs involved can be achieved by employing the scatter and gather mechanism of buffer handling.

Summary

The scatter and gather mechanism is a valuable building block for system designers as they support the expansion and development of the Internet, the Public Switched Telephone Network, and Wide Area Corporate Networks into a complex of increasingly convergent multi-service networks.

ASSPs that contain more capabil-

continued on page 5

Data and Buffer Management...continued

ity, that is, an increased percentage of the total system's function, effectively off-load development time for system designers and so decrease time-to-market. And at the same time, ASSPs' increased intelligence—programmability—enables differentiation of services.

However, to truly make profitable use of these increased abilities, limited system resources must be managed effectively at each level. For data handling, the scatter and gather mechanism is a valuable strategy whereby to attain such effective management. ❖

Suggested for Further Reading

For more information on this and related topics, go to www.transwitch.com/buffermgmt.

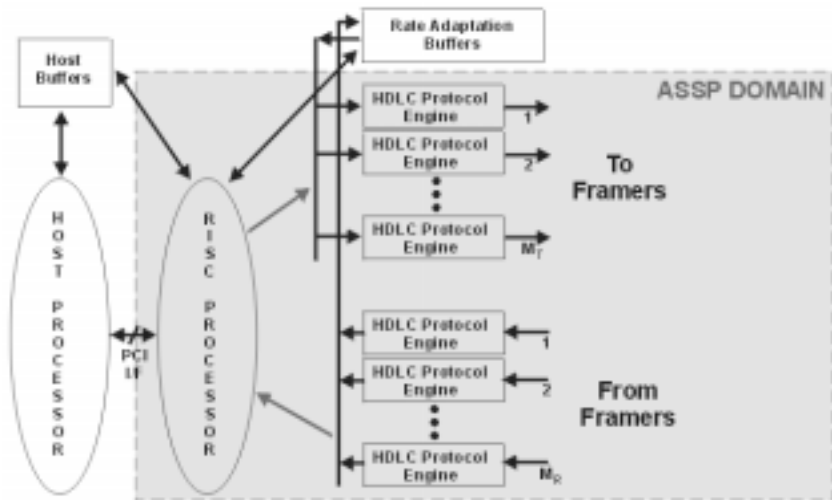


Figure 6. The Distributed Processing Architecture.

Fall Roundup

Record Third Quarter

For the quarter ending September 30, 1999, TranSwitch reported that total revenue showed a 51 percent increase over the third quarter of 1998. In addition, the Company reported diluted earnings per share of \$0.36, for the third quarter of 1999. This included a one-time positive adjustment of \$4.5 million in total to the provision for taxes. Absent this positive tax

adjustment and had the tax rate been in line with the five percent tax rate as in the first two quarters, diluted earnings per share would have been \$0.21, versus diluted earnings per share of \$0.08, in the third quarter of 1998.

“Our current revenue growth trend and current profitability, our success in penetrating the three high-growth end-markets—namely, the Public Network Infrastructure,

Internet Infrastructure, and corporate Wide Area Network (WAN)—our significant corporate relationships, and our solid balance sheet all give us the confidence that we will conclude 1999 with a significant growth in revenue and profits over 1998 and begin the new millennium with positive anticipation,” commented Dr. Santanu Das, President and CEO of TranSwitch Corporation. ❖

Siemens Design Win

In September, TranSwitch's ASPEN-based *Cellbus* architecture was selected by Siemens for its next-generation broadband access system, XpressLink, “. . . because we believe this is the best, field proven, switch fabric solution for the access market,” said Albrecht Baur, President Broadband Access Division at Siemens Information and Communication Networks.

“The programmability of ASPEN,” said Baur, “will provide XpressLink with significant additional capabilities

such as in-the-field re-provisioning for standards updates, feature enhancements and differentiated services.”

Through programmability, ASPEN can support full ATM layer processing, Voice over IP (VoIP) and Voice over ATM (VoATM) gateways, ATM adaptation layer (AAL2) trunking, frame relay and TCP switching, and frame relay to ATM interworking.

Further confirmation of the success of the *Cellbus* family of multiservice-programmable devices was clear in the new strategic technology partnership between TranSwitch and pri-

vately held Onex Communications Corporation.

Under the agreement, joint ownership rights in certain intellectual property developed by TranSwitch have been assigned to Onex for their product development.

This partnership will enable TranSwitch customers to develop access and transport products for converged networks covering a performance range of 1Gb/s to more than 900Gb/s, using TranSwitch CUBIT and ASPEN devices with Onex iTAP products. ❖

List of TranSwitch Products

<u>Product Name</u>	<u>Product Number</u>	<u>Product Description</u>
Asynchronous VLSI Devices		
ART/ARTE	TXC-02020/21	Advanced DS3/STS-1 Line Interface Device
DART	TXC-02030	Advanced E3/DS3/STS-1 Line Interface Device
DS3F	TXC-03401	DS3 Framer Device
DS3LIM-SN	TXC-20153	DS3/STS-1 Line Interface Module
E123MUX	TXC-03361	E1/E2/E3 Mux/Demux Device
E1Fx8	TXC-03109	8-Channel E1 Framer Device
E3LIM	TXC-20163	E3 Line Interface Module
E3RT	TXC-02053	34-Mbit/s E3 Line Interface Device
E2/E3F	TXC-03701	8- 34-Mbit/s Framer Device
HDLC	TXC-05101	HDLC Controller Device
MCHDLC	TXC-05132	Multi-Channel HDLC Controller Device
JT2F	TXC-03702	6-Mbit/s Framer Device
M13E	TXC-03303	DS3/DS1 Mux/Demux Device
MRT	TXC-02050	6- 8- 34-Mbit/s Line Interface Device
QE1F- <i>Plus</i>	TXC-03114	Quad E1 Framer- <i>Plus</i> Device
QT1F- <i>Plus</i>	TXC-03103	Quad T1 Framer- <i>Plus</i> Device
T1Fx8	TXC-03108	8-Channel T1 Framer Device
XBERT	TXC-06125	Bit Error Rate Generator/Receiver Device
SONET/SDH Synchronous VLSI Devices		
ADMA-E1	TXC-04002	2-Mbit/s to TU-12 Async Mapper-Desync Device
ADMA-T1/T1P	TXC-04001/11	1.544 Mbit/s to VT1.5/TU-11 Async Mapper-Desync Device
DS1MX7	TXC-04201	DS1 Mapper 7-Channel Device
L3M	TXC-03452	SDH/SONET Level 3 Mapper Device
L4M	TXC-03456	SDH/SONET Level 4 Mapper Device
PHAST-1	TXC-06101	SONET STS-1 Overhead Terminator Device
PHAST-3N	TXC-06103	STM-1/STS-3/STS-3c SDH/SONET OH Terminator Device with Telecom Bus 1/F
PHAST-3P	TXC-06203	STM-1/STS-3c SDH/SONET OH Terminator Device with UTOPIA Interface
PHAST-12	TXC-06112	Programmable, High-Performance ATM, SONET/SDH Terminator
QE1M	TXC-04252	Quad E1 Mapper Device with Microcode
QT1M	TXC-04251	Quad T1 Mapper Device
SOT-1/SOT-1E	TXC-03001/11	SONET STS-1 Overhead Terminator Device
SOT-3	TXC-03003	STM-1/STS-3/STS-3c Overhead Terminator Device
SYN155/155C	TXC-02301/02	155-Mbit/s Synchronizer Device
ATM (Asynchronous Transfer Mode) VLSI Devices		
ASPEN	TXC-05810	<i>CellBus</i> Access Processor Device
CDB	TXC-05150	ATM Cell Delineation Block Device
CUBIT- <i>Pro</i>	TXC-05802	ATM <i>CellBus</i> Switch Device
CUBIT-3	TXC-05804	Multi-PHY <i>CellBus</i> Access Device
SALI-25C	TXC-07625	Six ATM 25-Mbit/s Interface Controllers Device
SARA-R	TXC-05601	ATM/SMDS Reassembly Controller Device
SARA-S	TXC-05501	ATM/SMDS Segmentation Controller Device

Evaluation Boards

Contact TranSwitch Sales

Development Support Boards and Systems

Contact TranSwitch Sales

communicator © 1999 is a publication of TranSwitch Corporation. If you have comments, need more copies, or wish to reprint any portion of this newsletter, contact Mary Lombardo, TranSwitch Corporation, 3 Enterprise Drive, Shelton, CT 06484 USA. Phone: (203) 929-8810 Fax: (203) 926-9453 e-mail: info@txc.com Web: www.transwitch.com
TranSwitch is a registered trademark of TranSwitch Corporation. TranSwitch is ISO 9001 registered.
