Credit-Flow-Controlled ATM over HIC Links in the ASICCOM "ATLAS I" Single-Chip Switch

Manolis Katevenis†, Panagiota Vatsolaki†, and Vagelis Chalkiadakis††

† Institute of Computer Science (ICS), Foundation for Research and Technology – Hellas (FORTH),
Science and Technology Park of Crete, Vassilika Vouton, P.O.Box 1385, Heraklion, Crete, GR 711 10 Greece
Telephone: +30 (81) 391664, 391665, Fax: +30 (81) 391661, E-mail: {katevenis,vatsolaki}@ics.forth.gr
also with the Dept. of Computer Science, University of Crete, Heraklion, Crete Greece.

†† INTRACOM S.A., Development Programmes Department, P.O. Box 68, Peania, GR 190 02, Greece.
Telephone: +30 (1) 686 0486, Fax: +30 (1) 686 0312, E-mail: vhal@intranet.gr

ABSTRACT: ATLAS I is a single-chip gigabit ATM switch. Reliable and efficient communication with quality-of-service guarantees is essential in real-time systems. ATLAS I routes data cells reliably by implementing credit-based flow control on its IEEE Std. 1355 "HIC" links. Quality-of-service and efficiency are supported by implementing three levels of priority and by eliminating head-of-line blocking with *multi-lane* credit flow control.

1. Introduction

Switches are the building blocks for high-speed networks in which high performance is achieved using point-to-point connections. Asynchronous Transfer Mode (ATM) [LeBo92] has been adopted as the main high speed technology in both wide and local area networks. ATM has been architected so that it can provide *guaranteed quality-of-service* connections, which is an essential capability in the real-time world. The *ASICCOM* [1] project –the "Gigabit Switching" task project of the European Union *ACTS* (Advanced Communication Technologies and Services) Programme– is developing *ATLAS I* (ATm multi-LAne Switch I): a general-purpose, single-chip gigabit ATM switch, with credit-based flow control, multiple priority levels, and multicasting [KaSV96]. ATLAS I uses the OMI-developed IEEE Standard 1355 "HIC" technology for its links.

This paper presents a brief overview of the ATLAS I single-chip ATM switch, and discusses in more detail how ATM runs over the HIC links, and how *multi-lane* credit flow control is implemented on top of the standard single-lane HIC flow control. Credit flow control is essential in order for the network not to drop data, and multi-lane flow control is essential in order to avoid

[1] "ATM Switch for Integrated Communication, COmputation and Monitoring" (ASICCOM). The ASICCOM Consortium consists of industrial partners (INTRACOM, Greece; SGS THOMSON, France and Italy; BULL, France), telecom operators (TELENOR, Norway; TELEFONICA, Spain), and research institutes (FORTH, Greece; SINTEF, Norway; Politecnico di Milano, Italy; NCSR Democritos, Greece).

head-of-line blocking and provide the priorities and quality-of-service guarantees that are needed in real-time systems.

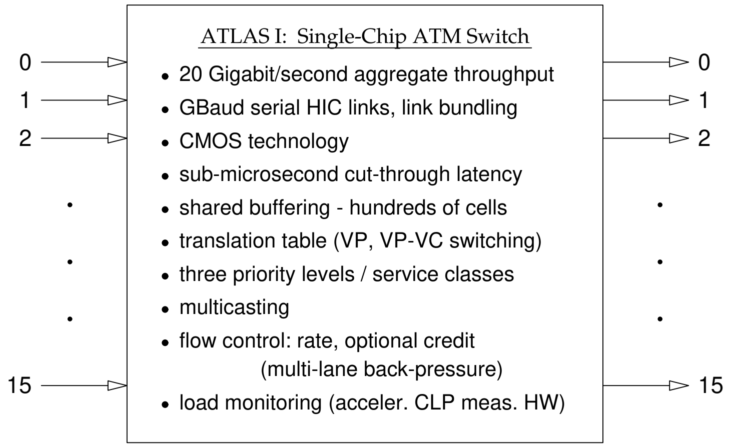

Figure 1: ATLAS I chip overview

ATLAS I is a general-purpose building-block chip for high-throughput and low-latency networks. To support a wide range of information transport services, the chip provides a number of features and mechanisms, as illustrated in figure 1. ATLAS I is a single-chip 16×16 ATM switch, with point-to-point serial links running at 622 Mbits/s each. When ports are bundled together, links of 1.25, 2.5, or even 5.0 Gbits/s can be synthesized. ATLAS I is currently under design, and will be fabricated by SGS Thomson, Crolles, France, in 0.5 μm CMOS; the clock frequency will be 50 MHz. Its intended use is in systems ranging from wide area (WAN) to metropolitan (MAN), local (LAN), and desktop (DAN) area networking. It is appropriate for supporting a mixture of services ranging from real-time, guaranteed-quality-of-service to best-effort, bursty and flooding traffic. These networks can provide a range of services from telecom to multimedia and multiprocessor applications.

2. Credit-Based Flow Control

In switch-based networks, contention for outgoing links of switches may lead to buffer overflows, as illustrated in figure 2; these are handled in either of two ways. Some networks allow packets/cells to be dropped, and use *end-to-end* congestion control protocols [Jaco88], [Ohsa95]. Other networks use *credit-based (back-pressure) flow control,* on a hop-by-hop basis, so as to never drop packets. ATLAS I supports both methods. In networks (or service classes) where cell dropping is allowed, ATLAS I provides ATM Forum standard *EFCI* (Error Forward Congestion Indication) [Ohsa95]. In addition, ATLAS I implements *optional, multi-lane* credit flow control, which has many important advantages as explained below.

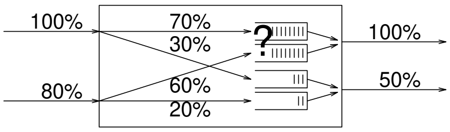

Figure 2. Example of output contention in a switch: too much traffic (130 %) is destined to the upper output, resulting in buffer overflow; this can be handled by dropping cells or by throttling down the upstream neighbors.

2.1 "Back-Pressure" (Credit) Flow Control

Credit flow control, also called back-pressure, or hop-by-hop window, or virtual channel (circuit) flow control, dates back to Tymnet [Rind77], GMDNET [Gies78], and Transpac [SiDa79], in packet switched networks; see also the survey [GeKl80]. Under credit flow control, a cell is only transmitted to the downstream neighbor if the transmitter knows that buffer space is available for it at the receiver. This is implemented as illustrated in figure 3, using a count of available receiver buffer space. This "credit count" is decremented every time a cell departs and is incremented when a credit token is received; credits are sent back every time a unit of buffer space becomes available.

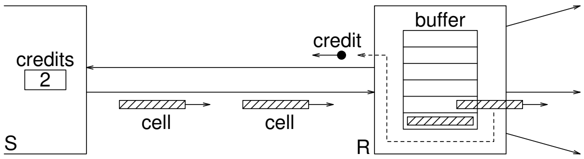

Figure 3. Back-pressure (credit) flow control: the credit count says whether buffer space is available in the downstream switch.

For a given link peak throughput, the receiver buffer space that is necessary and sufficient is equal to this *throughput* times the *round-trip time* (RTT) of data and credits. This buffer space, which is reminiscent of the braking or accelerating distance of a vehicle, is equal to the maximum number of cells and credits that can be in transit over the link. When the link operates at full speed and then *R*'s output suddenly stops, all these cells must be held in the given buffer space. Conversely, when *R*'s output suddenly starts, *R* must have that many cells available to sustain full speed operation until new cells have had the time to arrive from *S*.

Compared to end-to-end congestion control protocols, back-pressure has the advantage of never dropping cells; thus, retransmissions are not needed. Transmission can start right away at top speed, being automatically throttled down to the correct level only if and when this is necessary. These result in the minimum possible transmission time for a given message. Never dropping cells means that expensive network resources (link capacity) are not wasted; low-priority traffic can "fill in" the available throughput, resulting in full link utilization.

2.2 Head-of-Line Blocking in Single-Lane Back-Pressure

In plain (single-lane) back-pressure, there is *no selectivity* of what kind of cells (which VCs) it is desirable to stop and which ones to let go. This creates a behavior similar to the phenomenon of *head-of-line blocking* in input queueing [KaHM87], leading to performance degradation. An analogy in every-day life is single-lane streets (where the name came from): cars waiting to turn left block cars headed straight. If one uses the single-lane credit flow control described above, and *indiscriminate sharing* of buffer space, then bursty traffic creates the equivalent effect to head-of-line blocking, as illustrated in figure 4.

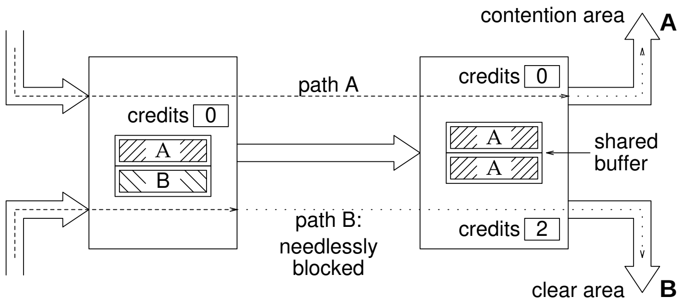

Figure 4. Indiscriminate (single-lane) back-pressure and buffer sharing: a burst of cells destined to *A* has filled all available buffer space in the downstream switch, and was then stopped by back-pressure due to contention further down the path to *A*. As a result, a cell destined to *B* is needlessly blocked in the upstream switch; poor link utilization results.

2.3 The ATLAS I Multi-Lane Back-Pressure

Multi-lane (VP/VC-level) credit-based flow control solves the above performance problem by selectively allowing unblocked traffic to proceed, while preventing congested connections from occupying all available buffer space thus delaying all other traffic as well. In wormhole routing, multi-lane back-pressure was used to prevent deadlocks [DaSe87], or to improve performance [Dally90]; multi-lane wormhole routing has been implemented in hardware in Torus [DaSe87] and iWarp [Bork90]. For ATM, VC credit-based flow control extensions have been proposed and studied for long-distance (WAN) links; e.g. see [KuBC94], [OzSV95]. In case of contention, VC-level credits combined with round-robin servicing allow the fair sharing of the available network capacity among all competing flows [HaGa86] [Kate87] [KaSC91]. DEC's GIGAswitch/ATM [Dec94] is a commercial ATM switch that implements VC credit flow control. Our earlier switch, Telegraphos [KVES95], also had a simple form of multi-lane back-pressure.

The credit flow control protocol of ATLAS I is reminiscent of, but simpler than those of [OzSV95] and [KuBC94], since the ATLAS I protocol is adapted to short links and single-chip hardware implementation. ATLAS I credits operate on the granularity of *flow groups.* A flow group is defined as a set of connections, sharing one or more links in their path through the network, whose cells need never or can never overtake each other, i.e. whose cell order is preserved over the length of their common path. Ideally, each flow group on each network link should consist of all connections whose destination is one, common, *single-threaded* device. When this definition leads to an impractically large number of flow groups, either a compromise is suggested −group together connections that go to nearby destinations−, or a higher level of back-pressure signals may be employeed to appropriately guide the multiplexing (selection) of cells at their point of entry into the flow group.

The ATLAS I chip is optimized for short-distance links; thus, each flow group can have at most 1 credit, for 1 cell's worth of buffer space. Given that the round-trip time on links up to several meters long, including switch logic, is no more than one cell time (about 680 ns), one credit per flow group is enough. This restriction also simplifies queueing, since there is no need to remember cell order when at most one cell per VC can be present inside the switch at any given time. Figure 5 shows the ATLAS I credit protocol. For each link, the downstream switch (statically) allocates a buffer pool of size L. In order for this buffer space not to overflow, the upstream switch maintains a credit count (per output port), called *pool credit* (*poolCr*), corresponding to the available pool space at the receiver. This buffer is shared among cells according to their flow group. As shown in figure 5, separate credit counts, *fgCr*, are maintained for each flow group, initialized to 1. Since the buffer pool contains L cell slots, and since each flow group is allowed to occupy at most 1 of these slots, L of the flow groups can share this buffer pool at any given time; thus, L is the *number of lanes.*

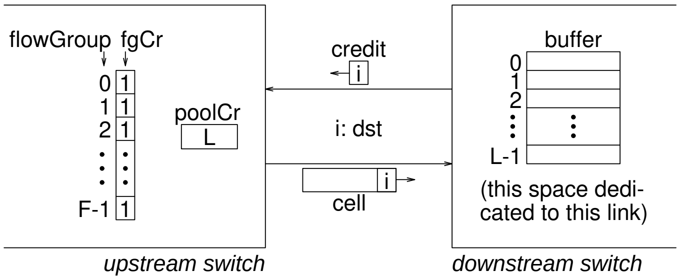

Figure 5. ATLAS I back-pressure: F flow groups sharing L lanes on a link. A cell can only depart if it can get a pool credit (so as to ensure that there is buffer space for it at the other end) *and* the (single) credit of its flow group (so as to ensure that identically destined cells do not consume excessive buffer space).

A cell belonging to flow group i can depart if and only if $fgCr[i] > 0$ (i.e. $fgCr[i] == 1$) and $poolCr > 0$. When it departs, $fgCr[i]$ and $poolCr$ are decremented by 1, each (i.e. $fgCr[i] \leftarrow 0$). When that cell leaves the down-stream switch, a credit carrying its flow group ID, i, is sent upstream. When credit i is received, $fgCr[i]$ and $poolCr$ are incremented by 1, each (i.e. $fgCr[i] \leftarrow 1$). The pool credit counts of all upstream neighbors of a switch (for the links leading to this switch) are initialized to values whose sum does not exceed the total buffer space available in the switch.

Credit flow control is optional on ATLAS I, so ATLAS-based *standard-ATM* networks can disable it. On the other hand, gigabit ATM LANs (including multiprocessor "NOW" and desktop area networks), where credit flow control is quite popular, can take advantage of ATLAS's credit features. Also, large ATM switches that use ATLAS I as building block for their switching fabric can use internally ATLAS's credit flow control in order to provide the good performance of output queueing at the low cost of input buffering; at the interfaces to the external links, the credit information can be converted to standard rate information.

3. Priority Classes, Shared Buffering, and Output Queues

Figure 6 shows the three service classes supported by the ATLAS I chip. They are differentiated by their level of priority; each connection (VP-VC pair) belongs to one of these classes. The classes provided are the interesting combinations of two parameters: policed or non-policed traffic, and back-pressured or non-back-pressured connections.

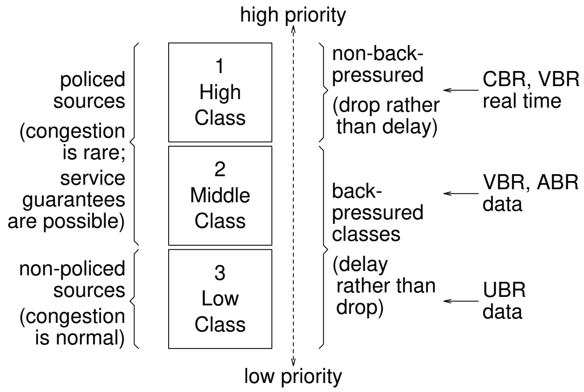

Figure 6. The three service classes of ATLAS I; whenever a cell of a higher priority class exists and is ready to depart over a link, such a cell will always be selected, regardless of the possible presence of lower priority cells.

Figure 7 shows the logical model of the ATLAS I operation. The bodies of incoming cells are stored in the shared data buffer. This buffer uses the pipelined memory architecture [KaVE95] [Efth95], which is particularly suitable for VLSI switch buffers. Besides its compactness and speed advantages, this architecture automatically provides virtual cut-through without needing any additional control circuits for that purpose. The address where each body is stored is used as an identifier (pointer) for the cell in the control (management) section of the chip.

In the control section, each incoming cell header first goes through the (shared) translation/routing table, which determines the service class, the translated VP/VC identifier for the cell, and the bit-mask (16 bits) of outgoing link(s). Cells that are allowed to depart from the switch are called *ready;* when contention for the outgoing link(s) allows it, they depart. The identifiers of (pointers to) the ready cells are maintained in *Ready Cell Queues (RQ).* There are three sets of ready queues, one per service class. For each service class, there is one ready queue per outgoing link for unicast cells, plus one queue shared among all links for the multicast cells. The cells in each queue are served (i.e. they depart) in FIFO order. Selection among the queues is made according to priorities. When a cell departs, a credit is sent back to the upstream neighbor where this cell came from.

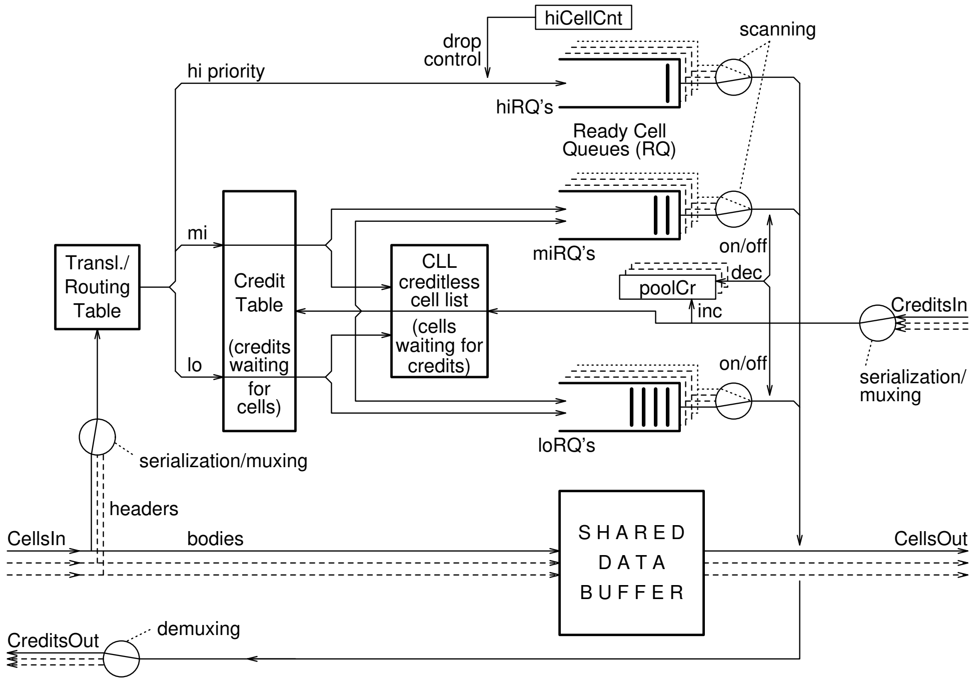

Figure 7. ATLAS I queue model; for simplicity, only 3 of the 16 links are shown. The Ready Cell Queues hold pointers to the cells that are allowed to depart; there is a separate queue for each outgoing link and for each level of priority.

Cells that belong to the high service class are not subject to credit flow control (back-pressure); thus, they are either enqueued directly to the proper ready queue, or, on buffer overflow, they are dropped. All high-class cells share *one common partition* of the total data buffer capacity. As it is known, this shared-queueing buffer allocation results in the minimum possible cell loss probability for a given buffer memory space [HlKa88]. The size of the data buffer partition that is dedicated to high-class cells is determined by the network manager, and can be changed at run-time. Note that this is only a *logical* partition, i.e. high-class cells are physically stored anywhere in the data buffer, but their total number is controlled so as to remain below a given upper limit.

Arriving cells that belong to the middle and low service classes must acquire a credit for their flow group before they can get enqueued in the proper ready queue. Control first consults the *Credit Table;* those credits that are needed by the cell and are found in the table are taken from there. If all credits needed by the cell are found, then the cell is enqueued in the proper ready queue. Else, the cell is recorded in the *Creditless Cell List (CLL),* where it waits for the remaining credits to arrive.

Incoming credits are examined one at a time. Each of them increments the proper pool credit counter, and then it is given to the creditless cell list (CLL). If there are cells waiting for this credit in that list, one of them is selected (pseudo-) randomly and gets the credit; if this was the last credit that the cell was waiting for, the cell becomes ready and gets enqueued to the proper ready queue. Else, if no cells wait for this credit in CLL, the credit is recorded in the credit table.

4. Multi-Lane Back-Pressure ATM over HIC Links

To implement a multi-gigabit-per-second switch chip, one has to solve first the problem of providing that much bandwidth through the chip's periphery (pins). ATLAS I uses the high-speed parallel-serial conversion "STRINGS" macrocell that has been developed by BULL within OMI [MCLN93]. The CMOS version of this macrocell is capable to operate at Gigabaud speed on the serial side, which is the signaling frequency required to provide 622 Mbps of data rate, given the 8b/12b encoding used by the circuit.

4.1 IEEE Standard 1355 HIC Links

The ATLAS I links run ATM on top of IEEE Std. 1355 HIC/HS as the physical layer. Each ATM cell is carried as a sequence of 53 data bytes, followed by 1 end-of-packet (EOP) control character. Idle characters are not allowed inside a cell, but an arbitrary number of idle's can appear between cells. Although ATM today is most frequently run on top of SDH/SONET (STM1/STM4), ATM is a network technology that is independent of physical layer, and has been run at least on UTP-3, T1/T3, E1/E3, STM1/STM4, and FDDI-physical. HIC (IEEE 1355) is our link protocol of choice, because ATLAS I is optimized for short-distance links, as found inside large ATM switch boxes for WANs and in local and desktop area networks.

To yield a net data rate of 622.08 Mb/s (1.4672 Mcells/s), HIC/HS uses a signaling rate around 1 GBaud, because each byte is encoded using 12 bits rather than e.g. 8 bits in SDH/SONET. The advantage of this encoding, however, is that no external crystals and no internal PLLs (expensive circuits) are needed for each of the 16 incoming links of every ATLAS I chip. HIC is much simpler than SDH/SONET to implement in hardware; in fact, 16 SDH/SONET interface circuits would not fit on the single chip of ATLAS I.

A second reason for using HIC is that on an SDH/SONET link all cells must be aligned on cell-time boundaries (idle periods last for an integer number of cell times), whereas cells on a HIC link can have arbitrary byte-aligned starting times. Thus, HIC reduces the average (cut-through) latency of the switch (under light load) by half a cell time (340 ns), which is important in local area applications. And thirdly, for short distance links and correspondingly small buffer spaces, the transmission of a flow control credit (section 2) cannot cost as much as an entire cell, neither can it be delayed until it gets bundled with other credits inside a cell; SDH/SONET leaves no room for credit transmission outside cells, while HIC/HS allows credits to be encoded in a straightforward way, on arbitrary byte-aligned times, using dedicated control character(s).

4.2 Two Levels of Back-Pressure

IEEE Std. 1355 HIC provides single-lane back-pressure flow control, while ATLAS I implements optional multi-lane back-pressure *on top of* this existing first level, as illustrated in figure 8. In this schematic, the central part is one direction of the standard IEEE 1355 link. Only right-bound data and left-bound credits are shown; the real link multiplexes these with right-bound credits and left-bound data, correspondingly. The receiver includes a FIFO queue, *rq*, so that the receiving application is free to pull data out of the link at whatever rate it desires. The transmitter stops transmitting when it runs out of *single-lane* credits, *scr*, so as to not overflow *rq*. The transmitter FIFO queue, *xq*, isolates the instantaneous rate of the sending application from the line rate.

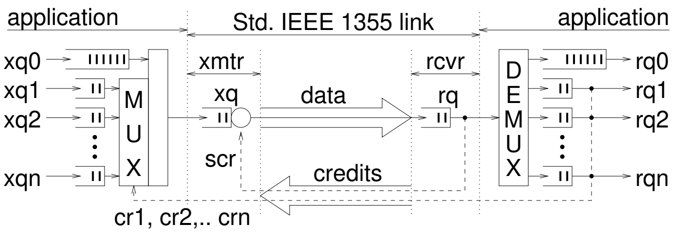

Figure 8. Two levels of back-pressure: the standard HIC single-lane credits, *scr*, control the transfer of information from *xq* to *rq*, while multi-lane credits transparently pass through the HIC link and control the transfers between the multiple logical queues of ATLAS I (one per flow group).

The ATLAS I *application* that we build on top of this standard link transfers cells among multiple logical queues: from xq_0, xq_1, xq_2,... xq_n, to rq_0, rq_1, rq_2,... rq_n; cells from transmit queue i go to receive queue i. A selection/multiplexing circuit selects one of the candidate transmit queues every time, and sends a cell from it; candidate queues are only those for which it is known that space exists in their respective receive queue. (xq_0 is treated differently: it contains high-priority non-back-pressured cells). This knowledge is achieved using (multi-lane) credits, cr_1, cr_2,... cr_n, as explained in section 2.3. Multi-lane credits are sent by the downstream application, whenever cells depart from rq_1, rq_2,... rq_n.

Multi-lane credits are encoded as a sequence of bytes, starting with a control character, *MLCR,* that identifies this as a multi-lane credit, and followed by data bytes which carry the flow group number that must receive the credit. The MLCR control character is distinct from the control characters that are used by the standard HIC/HS devices (characters 0, 1,... 6, or 125), so MLCR passes transparently through HIC/HS links. In ATLAS I, multi-lane credits can be interleaved any time in the normal link traffic, except that they are not allowed to appear inside the header of an ATM cell, at most one multi-lane credit is allowed inside the body of an ATM cell, and multi-lane credits must be spaced from each other by at least half a cell time.

Conclusions

Priorities and multi-lane credit (back-pressure) flow control are important mechanisms in providing efficiency and quality-of-service guarantees in networks that carry real-time information. ATM, extended with credit-based flow control, is an appropriate technology for such high-speed networking. The ASICCOM project is developing ATLAS I, a single-chip ATM switch that incorporates all the above features.

Acknowledgements

ATLAS I is being developed within the ASICCOM project, funded by the European Union ACTS Programme. This is a collective effort, with many individuals making valuable contributions. We would like to thank all of them, and in particular Dimitrios Serpanos, Costas Courcoubetis, Luis Merayo, Helge Rustad, Evangelos Markatos, Vassilios Siris, Christoforos Kozyrakis, Chara Xanthaki, Yannis Papaefstathiou, George Kalokerinos, Kostas Kassapakis, Roland Marbot, Bjorn Olav Bakka, Geir Horn, and Jorgen Norendal.

References

[Bork90] S. Borkar e.a.: "Supporting Systolic and Memory Communication in iWarp", *Proc. of the 17th Int. Symp. on Computer Architecture,* ACM SIGARCH vol. 18, no. 2, June 1990, pp. 70-81.

[DaSe87] W. Dally, C. Seitz: "Deadlock-Free Message Routing in Multiprocessor Interconnection Networks", *IEEE Trans. on Computers,* vol. C-36, no. 5, May 1987, pp. 547-553.

[Dally90] W. Dally: "Virtual-Channel Flow Control", *Proc. of the 17th Int. Symp. on Computer Architecture,* ACM SIGARCH vol. 18, no. 2, May 1990, pp. 60-68.

[Dec94] "GIGAswitch/ATM Networking Switch", *Digital Equipment Co.,* Nov. 1994, URL: file: //ftp.digital.com/pub/Digital/info/infosheet/ EC-F3924-42.ps

[Efth95] A. Efthymiou: "Design, Implementation, and Testing of a 25 Gb/s Pipelined Memory Switch Buffer in Full-Custom CMOS", Master of Science Thesis, University of Crete, Greece; Technical Report FORTH-ICS/TR-143, Institute of Computer Science, FORTH, Heraklio, Crete, Greece, November 1995; URL: file://ftp.ics.forth.gr/tech-reports/1995/ 1995.TR143.Design_Implementation_25Gbs_PipelinedMem_Switch_Buffer.ps.gz

[GeKl80] M. Gerla, L. Kleinrock: "Flow Control: A Comparative Survey", *IEEE Trans. on Communications,* vol. 28, no. 4, April 1980, pp. 553-574.

[Gies78] A. Giessler, e.a.: "Free Buffer Allocation – An Investigation by Simulation", *Comput. Networks,* vol. 2, 1978, pp. 191-208.

[HaGa86] E. Hahne, R. Gallager: "Round Robin Scheduling for Fair Flow Control in Data Communication Networks", *Proc. IEEE Int. Conf. on Commun.,* June 1986, pp. 103-107.

[HlKa88] M. Hluchyj, M. Karol: "Queueing in High-Performance Packet Switching", *IEEE Journal on Sel. Areas in Communications,* vol. 6, no. 9, December 1988, pp. 1587-1597.

[Jaco88] V. Jacobson: "Congestion Avoidance and Control", *Proceedings of the ACM SIGCOMM '88 Conference,* Stanford, CA USA, August 1988, ACM Comp. Comm. Review, vol. 18, no. 4, pp. 314-329.

[KaHM87] M. Karol, M. Hluchyj, S. Morgan: "Input versus Output Queueing on a Space-Division Packet Switch", *IEEE Trans. on Communications,* vol. COM-35, no. 12, December 1987, pp. 1347-1356.

[Kate87] M. Katevenis: "Fast Switching and Fair Control of Congested Flow in Broad-Band Networks", *IEEE Journal on Sel. Areas in Communications,* vol. 5, no. 8, October 1987, pp. 1315-1326.

[KaSC91] M. Katevenis, S. Sidiropoulos, C. Courcoubetis: "Weighted Round-Robin Cell Multiplexing in a General-Purpose ATM Switch Chip", *IEEE Journal on Sel. Areas in Communications,* vol. 9, no. 8, October 1991, pp. 1265-1279.

[KVES95] M. Katevenis, P. Vatsolaki, A. Efthymiou, M. Stratakis: "VC-level Flow Control and Shared Buffering in the Telegraphos Switch", *Proceedings of the Hot Interconnects III Symposium,* Stanford Univ., CA, USA, Aug. 1995. URL: file://ftp.ics.forth.gr/tech-reports/1995/1995.HOTI.VCflowCtrlTeleSwitch.ps.gz

[KaVE95] M. Katevenis, P. Vatsolaki, A. Efthymiou: "Pipelined Memory Shared Buffer for VLSI Switches", *Proceedings of the ACM SIGCOMM '95 Conference,* Cambridge, MA USA, 30 August - 1 Sep. 1995, pp. 39-48. URL: file://ftp.ics.forth.gr/tech-reports/1995/1995.SIGCOMM95.PipeMemoryShBuf.ps.gz

[KaSV96] M. Katevenis, D. Serpanos, P. Vatsolaki: "ATLAS I: A General-Purpose, Single-Chip ATM Switch with Credit-Based Flow Control", *Hot Interconnects IV Symposium,* Stanford Univ., CA, USA, Aug. 1996. URL: file://ftp.ics.forth.gr/tech-reports/1996/1996.HOTI.ATLAS_I_ATMswitchChip.ps.gz

[KuBC94] H.T. Kung, T. Blackwell, A. Chapman: "Credit-Based Flow Control for ATM Networks: Credit Update Protocol, Adaptive Credit Allocation, and Statistical Multiplexing", *Proceedings of the ACM SIGCOMM '94 Conference,* London, UK, 31 Aug. - 2 Sep. 1994, ACM Comp. Comm. Review, vol. 24, no. 4, pp. 101-114.

[LeBo92] J. LeBoudec: "The Asynchronous Transfer Mode: A Tutorial", *Computer Networks and ISDN Systems,* vol. 24, no. 4, May 1992.

[MCLN93] R. Marbot, A. Cofler, J-C. Lebihan, R. Nezamzadeh: "Integration of Multiple Bidirectional Point-to-Point Serial Links in the Gigabits per Second Range", *Hot Interconnects I Symposium,* Stanford Univ., CA, USA, Aug. 1993.

[Ohsa95] H. Ohsaki, e.a.: "Rate-Based Congestion Control for ATM Networks", *Computer Communication Review,* ACM SIGCOMM, vol. 25, no. 2, April 1995, pp. 60-72.

[OzSV95] C. Ozveren, R. Simcoe, G. Varghese: "Reliable and Efficient Hop-by-Hop Flow Control", *IEEE Journal on Sel. Areas in Communications,* vol. 13, no. 4, May 1995, pp. 642-650.

[Rind77] J. Rinde: "Routing and Control in a Centrally Directed Network", *Proc. Nat. Comput. Conf.,* Dallas, Tx, USA, June 1977.

[SiDa79] J. Simon, A. Danet: "Controle des Ressources et Principes du Routage dans le reseau TRANSPAC", *Proc. Int. Symp. Comput. Networks, Versailles, FR, February 1979.*