



**ATM Forum
Technical Committee**

**Multiplexed Status Mode
(MSM3) for
UTOPIA Level 3**

af-phy-0142.000

March 2000

© 2000 by The ATM Forum. This specification/document may be reproduced and distributed in whole, but (except as provided in the next sentence) not in part, for internal and informational use only and not for commercial distribution. Notwithstanding the foregoing sentence, any protocol implementation conformance statements (PICS) or implementation conformance statements (ICS) contained in this specification/document may be separately reproduced and distributed provided that it is reproduced and distributed in whole, but not in part, for uses other than commercial distribution. All other rights reserved. Except as expressly stated in this notice, no part of this specification/document may be reproduced or transmitted in any form or by any means, or stored in any information storage and retrieval system, without the prior written permission of The ATM Forum.

The information in this publication is believed to be accurate as of its publication date. Such information is subject to change without notice and The ATM Forum is not responsible for any errors. The ATM Forum does not assume any responsibility to update or correct any information in this publication.

Notwithstanding anything to the contrary, neither The ATM Forum nor the publisher make any representation or warranty, expressed or implied, concerning the completeness, accuracy, or applicability of any information contained in this publication. No liability of any kind shall be assumed by The ATM Forum or the publisher as a result of reliance upon any information contained in this publication.

The receipt or any use of this document or its contents does not in any way create by implication or otherwise:

- Any express or implied license or right to or under any ATM Forum member company's patent, copyright, trademark or trade secret rights which are or may be associated with the ideas, techniques, concepts or expressions contained herein; nor
- Any warranty or representation that any ATM Forum member companies will announce any product(s) and/or service(s) related thereto, or if such announcements are made, that such announced product(s) and/or service(s) embody any or all of the ideas, technologies, or concepts contained herein; nor
- Any form of relationship between any ATM Forum member companies and the recipient or user of this document.

Implementation or use of specific ATM standards or recommendations and ATM Forum specifications will be voluntary, and no company shall agree or be obliged to implement them by virtue of participation in The ATM Forum.

The ATM Forum is a non-profit international organization accelerating industry cooperation on ATM technology. The ATM Forum does not, expressly or otherwise, endorse or promote any specific products or services.

NOTE: The user's attention is called to the possibility that implementation of the ATM interoperability specification contained herein may require use of an invention covered by patent rights held by ATM Forum Member companies or others. By publication of this ATM interoperability specification, no position is taken by The ATM Forum with respect to validity of any patent claims or of any patent rights related thereto or the ability to obtain the license to use such rights. ATM Forum Member companies agree to grant licenses under the relevant patents they own on reasonable and nondiscriminatory terms and conditions to applicants desiring to obtain such a license. For additional information contact:

The ATM Forum

Worldwide Headquarters

2570 West El Camino Real, Suite 304

Mountain View, CA 94040-1313

Tel: +1-650-949-6700

Fax: +1-650-949-6705

Contact:

John R. Mick, Chair	Raphael Janczyk, Editor	Bernd Jaeger, Editor
Phone: +1 408 654 6721	Phone: +49 89 722-28387	Phone: +49 89 722-21208
Email: Mick@idt.com	Email: georg-raffael.janczyk@icn.siemens.de	Email: bernd.jaeger@icn.siemens.de

Table of Contents

1. Introduction 5

2. Multi-PHY Operation with multiple Clav Signals for Parallel Polling 5

 2.1 Back-to-back polling 5

 2.2 Clav-to-PHY Port Mapping 6

 2.3 Transmit Interface 7

 2.4 Receive Interface 9

1. Introduction

This document is an addendum to UTOPIA Level 3 and describes parallel polling of a large number of PHYs during 1 cell cycle by means of multiple optional CLAV signals. The suggested polling scheme tries to make use of the maximum number of polling cycles that are available during 1 cell cycle in order to minimize the number of additional optional Clav signals

2. Multi-PHY Operation with multiple Clav Signals for Parallel Polling

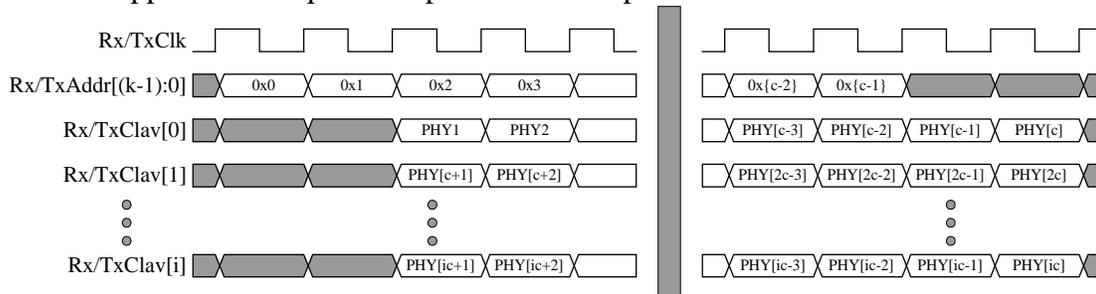
To obtain PHY status from a large number of PHY ports during 1 cell cycle, a mode with multiple Clav signals for parallel polling is defined as **optional (O)**.

In this mode the ATM layer device receives PHY port FIFO status information parallel on a variable number of Clav signals that are optionally provided for this capability. The behaviour of the Clav signals is corresponding to that described in chapter 4.3.1.

2.1 Back-to-back polling

The multiple Clav mode described in this section is very similar to the mechanism described in chapter 4.3.2 with the exception that more than 1 Clav signals are used to increase the number of PHY ports that can be polled during 1 cell cycle. Therefore the total amount of PHY ports is partitioned into several polling groups. Each polling group has its own Clav signal for PHY status indication. The lower k bits of the Address[n:0] is used to address up to $c=2^k$ PHYs inside the polling group while the higher j bits ($j=n-k+1$) of the Address[n:0] is used during PHY selection to address up to $g=2^j$ polling groups with each polling group being equivalent to one Clav signal. In applications where all ports are required to be polled in one cell time, 'c', i.e. the number of polling cycles per polling group, would be less or equal to 12 in case of a 32 bit wide UTOPIA Level 3 interface. The address order as well as the priority algorithm of the polling are application specific.

Examples of the Clav-to-PHY port mapping depending on the used buswidth and the number of supported PHY ports are provided in chapter 2.3.



c = number of polling cycles per polling group
 g = number of polling groups (equivalent to number of Clav signals)
 i = g - 1
 0x{c-2} = hex value of 'c-2'

Figure 1.1: Back-to-back Polling with multiple Clav signals

2.2 Clav-to-PHY Port Mapping

The Clav-to-PHY Port mapping is application-specific and is not specified by this document.

Depending on the UTOPIA interface buswidth, which can be 8bit, 16bit or 32bit a different number of polling cycles is available during one cell cycle. Worst case is the buswidth of 32bit with a cell format of 52byte and back-to-back cell transfer. In order to minimize the number of pins the maximum number of 12 polling cycles for that case is used. For the other bus widths the number of polling cycles are multiples of 12, which is 24 polling cycles for the 16bit wide UTOPIA interface and 48 polling cycles for the 8bit wide UTOPIA interface.

Table 1.1 defines the partitioning of PHY ports into polling groups for applications up and equal to

- 48 ports in case of a 32bit wide UTOPIA Level 3 bus,
- 96 ports in case of a 16bit wide UTOPIA Level 3 bus,
- 192 ports in case of a 8bit wide UTOPIA Level 3 bus.

This restricts the variety of combinations but is useful for compatibility reasons. For applications with port numbers above those listed before the partitioning of PHY ports into polling groups is left unconstrained.

UTOPIA Interface Width	No. of Pol-ling Cycles (c)	No. of PHY Ports (p)	Min. No. of Pol-ling (Clav) Groups (g)	No. of Addr. Bits for Addr[n:0] (n+1)=(j+k)
Acc. to UL3	Defined	Appl. specific	$g = \text{ceil}(p/c)$	$n = \text{ceil}[\log, g] + \text{ceil}[\log, c]$
32	12	16	2	$5 = 1+4$
		32	3	$6 = 2+4$
		48	4	$6 = 2+4$
16	24	16	1	$6 = 1+5$
		32	2	$6 = 1+5$
		48	2	$6 = 1+5$
		64	3	$7 = 2+5$
		96	4	$7 = 2+5$
8	48	16	1	$7 = 1+6$
		32	1	$7 = 1+6$
		48	1	$7 = 1+6$
		64	2	$7 = 1+6$
		96	2	$7 = 1+6$
		128	3	$8 = 2+6$
		192	4	$8 = 2+6$

ceil(x)=ceiling(x): next integer value greater or equal than x

Table 1.1: Partitioning of PHY ports into polling groups

For simplicity of illustration, the following examples assume that all ports would need to be polled during each cell time, although this is not a requirement of this specification for applications below STS-1.

As a specific example consider in Table 1.2 the case of a 32bit wide UTOPIA Level 3 interface supporting 48 ports (port1 ... port48). In this case $n+1=j+k=6$ address bits are necessary. The upper $j=2$ bits are used to address 4 polling (Clav) groups of 12 ports each while the lower $k=4$ bits are used to address the 12 ports inside a polling (Clav) group. (This is the defined partitioning but not mapping for ATM carried over a SONET STS-48 signal that is channelized down to the STS-1 level.)

	A5, A4	A3, A2, A1, A0											
		0xB	0xA	0x9	0x8	0x7	0x6	0x5	0x4	0x3	0x2	0x1	0x0
Clav[0]	00	12	11	10	9	8	7	6	5	4	3	2	1
Clav[1]	01	24	23	22	21	20	19	18	17	16	15	14	13
Clav[2]	10	36	35	34	33	32	31	30	29	28	27	26	25
Clav[3]	11	48	47	46	45	44	43	42	41	40	39	38	37

Table 1.2: Clav-to-PHY Port Mapping for 48 ports on a 32bit UL3 interface

2.3 Transmit Interface

The ATM layer device can send a cell to a PHY port only when the PHY port has indicated that it is ready to receive at least one cell. The PHY device must send Transmit Cell Available information to the ATM layer device.

This is illustrated in Figure 1.2 for the case of a ‘32bit wide UTOPIA Level 3 interface with 48 PHY ports’. The TxClav[0] to TxClav[3] signals indicate the status of the PHY transmit ports two clock cycles after the lower $k=4$ bits of the port’s address (TxAddr[3:0]) have been selected by the ATM layer device. The decode-response timing between the TxAddr and the TxClav is therefore 2 clock cycles.

TxAddr[5:0] during the clock cycle before asserting the TxEnb signal will select the PHY port for reception of the next cell. This will be decoded by the PHY device and the specified port will be ready to receive cell data from the ATM side at the next clock cycle.

The shifted polling cycle (0x0 ... 0xB) starts 3 clock cycles before SOC and ends 3 clock cycles before the end of the currently transferred cell to fulfill the decode-response time of 2 clock cycles for the TxClav signals as well as for selection of the next port.

In the clock cycle before clock edge 4 for example PHY port 16 is selected by TxAddress ‘010011 = 0x13’ (according to Table 1.2) which means port 4 of polling (TxClav) group [1]. During the time of a cell transfer all 48ports are polled parallel in 4 polling groups through TxClav[0] to TxClav[3]. It is assumed that the priority algorithm selects PHY port 48 (which is TxAddress ‘111011 = 0x3B’) for the next cell transfer to show the worst case for the decode response times.

Address ‘0xB’ is valid on the address bus between clock edge 13 and 14. PHY port 48 responds 2 clock cycles later (between clock edge 15 and 16) by asserting TxClav[3]. After another 2 clock cycles selection of PHY port 48 is done via deassertion of the ‘TxEnb*’ signal on clock edge 18 and TxAddr[5:0] = ‘0x3B’ on the address bus between clock edge 17 and 18.

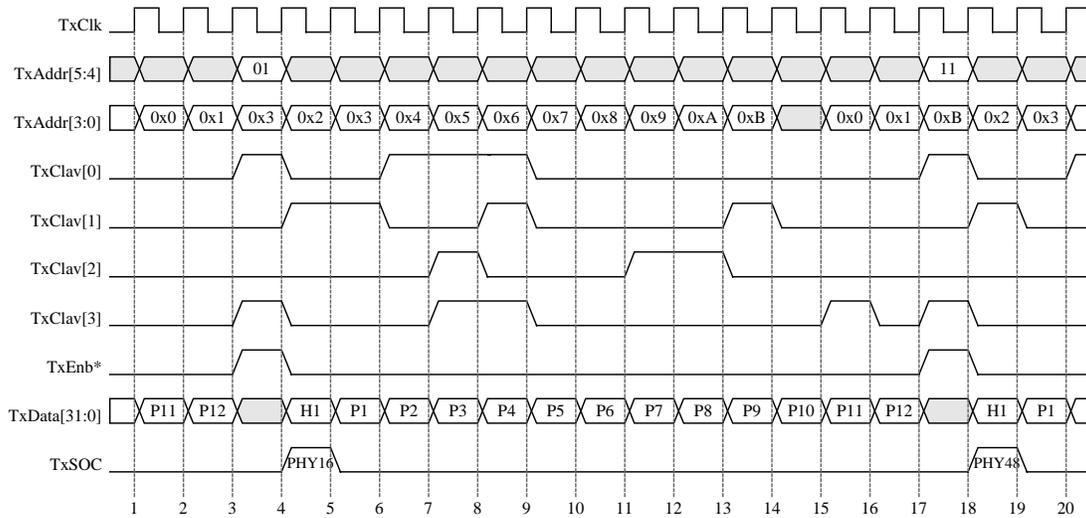


Figure 1.2: Transmit Decode and Selection for Parallel Polling

Back-to-back transfer of cells is possible when two cells have to be sent to the same PHY port and this PHY port indicates that it can receive the second cell. This is done by assertion of the TxClav when polled. In the case of back-to-back transfer the ATM layer device implicitly reselects the PHY port by leaving the TxEnb* asserted during the next to the last cycle of the cell. The second (or subsequent) cell is transferred immediately after the previous one and the TxSOC is asserted to indicate the start of cell. This is illustrated in Figure 1.3.

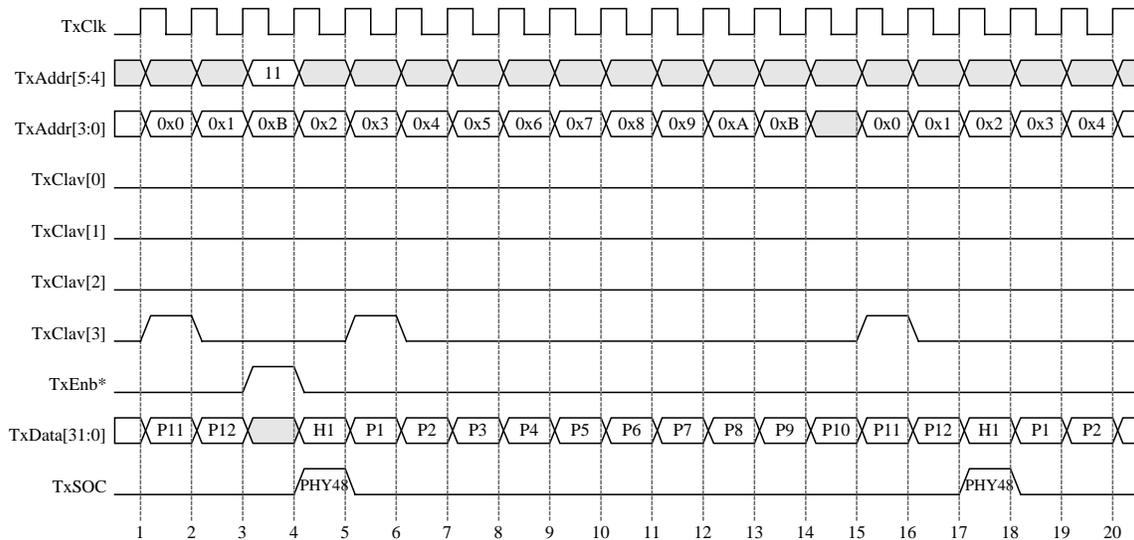


Figure 1.3: Back-to-back transmission of cells for Parallel Polling

2.4 Receive Interface

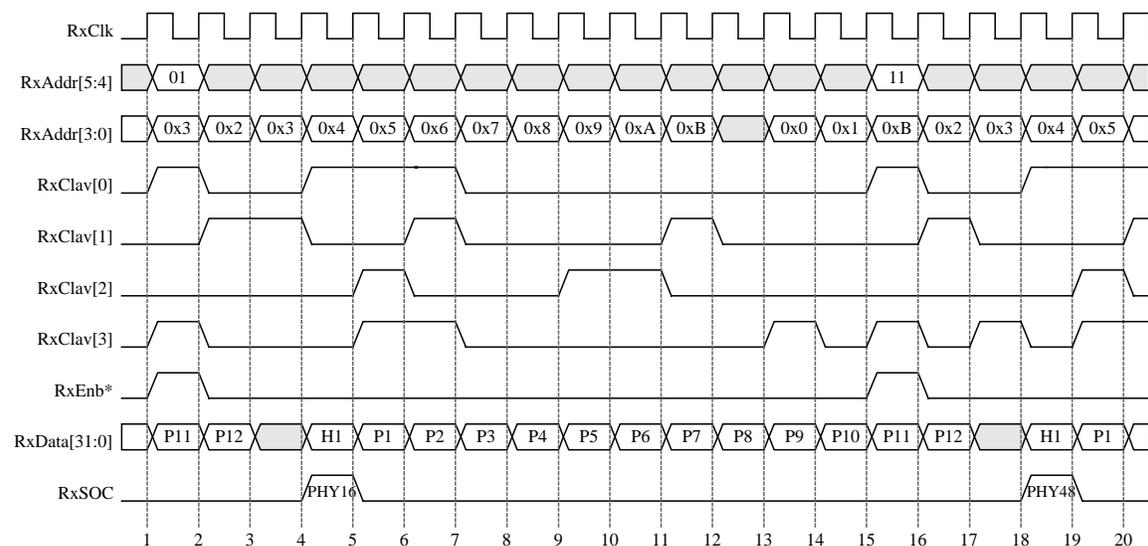
The ATM Layer can receive a cell from a PHY port only when the PHY port has indicated that it is ready to transmit at least one cell. The PHY device must send Receive Cell Available information to the ATM layer device.

This is illustrated in Figure 1.4 for the case of a '32bit wide UTOPIA Level 3 interface with 48 PHY ports'. The RxClav[0] to RxClav[3] signals indicate the status of the PHY receive ports two clock cycles after the lower k=4 bits of the port's address (RxAddr[3:0]) have been selected by the ATM layer device. The decode-response timing between the RxAddr and the RxClav is therefore two clock cycles.

RxAddr[5:0] during the clock cycle before asserting RxEnb* signal will select the PHY port which will transmit the next cell across the UTOPIA interface. This signal will be decoded by the PHY device and the specified port will be ready to transmit cell data two clock cycles after the falling edge of RxEnb. The decode-response timing between the RxEnb and the RxData is therefore two clock cycles.

The shifted polling cycle (0x0 ... 0xB) starts 5 clock cycles before SOC and ends 5 clock cycles before the end of the currently transferred cell to fulfill the decode-response time of 2 clock cycles for the RxClav signals, 2 clock cycles for the selection of the next port and 2 clock cycles for the beginning of the next cell transfer.

In the clock cycle before clock edge 2 for example PHY port 16 is selected by RxAddress '010011 = 0x13' (according to Table 1.2) which means port 4 of polling (RxClav) group[1]. During the time of a cell transfer all 48 ports are polled parallel in 4 polling groups through RxClav[0] to RxClav[3]. It is assumed that the priority algorithm selects PHY port48 (which is RxAddress '111011 = 0x3B') for the next cell transfer to show the worst case for the decode response times. Address '0xB' is on the address bus between clock edge 11 and 12. PHY port 48 responds 2 clock cycles later (between clock edge 13 and 14) by asserting RxClav[3]. After another 2 clock cycles selection of PHY port47 is done via the deassertion of the 'RxEnb*' signal on clock edge 16 and RxAddr[5:0] = '0x3B' on the address bus between clock edge 15 and 16. PHY begins to start the cell 2



clock cycles later (after clock edge 18) by asserting RxSOC.

Figure 1.4: Receive Decode and Selection for Parallel Polling

Back-to-back reception of cells is possible when two or more cells have to be received from the same PHY port and the ATM layer device is able to receive the second cell. The PHY layer device indicates it has another cell by assertion of the RxClav, when polled. In case of back-to-back transfer the ATM layer device does not explicitly select the PHY port as a transfer from the same port is assumed. The second (or subsequent) cell is transferred immediately after the previous one and the RxSOC is asserted to indicate the start of cell. This is illustrated in Figure 1.5.

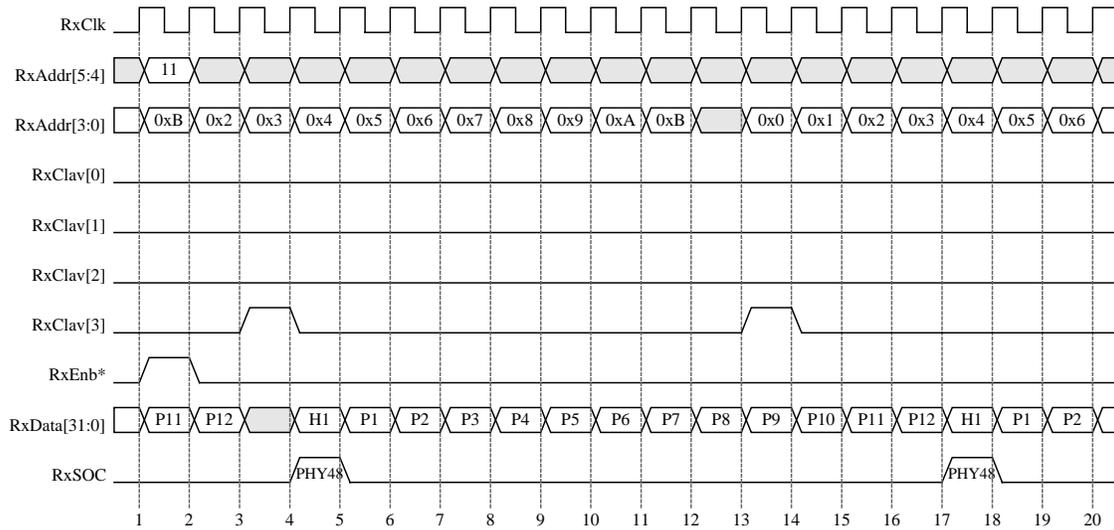


Figure 1.5: Back-to-back reception of cells for Parallel Polling