UM10944

Software user manual for SJA1105TEL

Rev. 1 — 25 January 2017 **User manual**

Document information

Info	Content
Keywords	SJA1105TEL, Ethernet, software registers
Abstract	This user manual describes the configuration (including the static configuration interface), register structure and mapping of the SJA1105TEL 5-port automotive Ethernet switch.

Revision history

Rev	Date	Description
1	20170125	first issue

Contact information

For more information, please visit: **http://www.nxp.com**

For sales office addresses, please send an email to: **salesaddresses@nxp.com**

1. Introduction

This software user manual describes the configuration of the SJA1105TEL 5-port automotive Ethernet switch. Topics covered include the static configuration interface and format, the register structure and mapping of the IP blocks. This document should be read along with the SJA1105 data sheet, available from NXP Semiconductors.

2. Functional overview

Figure 1 shows the building blocks that make up the SJA1105TEL. The base addresses of the core, CGU, RGU and ACU are given in Table 1. The dataflow followed by a single received frame as it passes through the switch is described in Section 2.1 to Section 2.3.

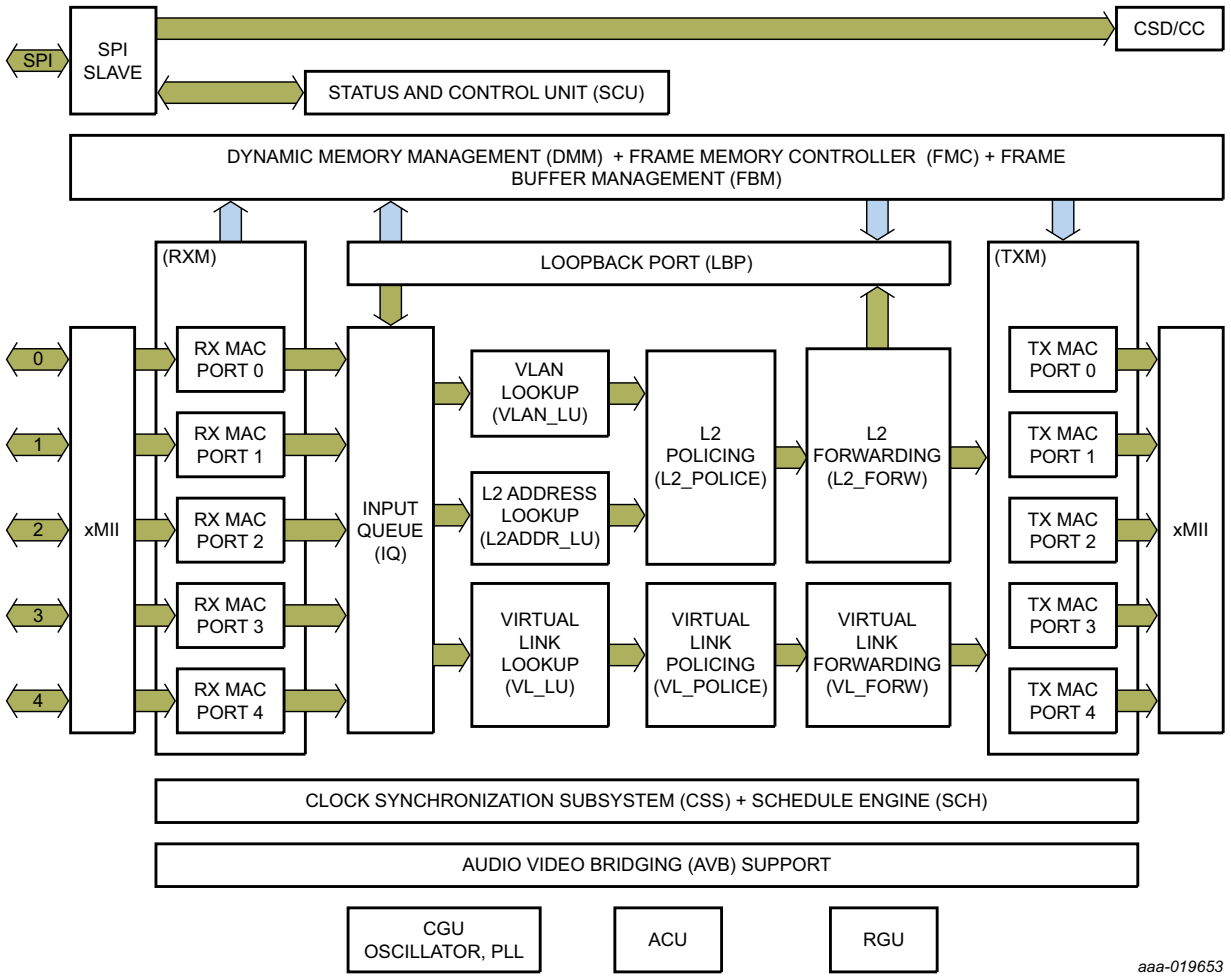

Fig 1. Block diagram of SJA1105TEL

Table 1. SJA1105TEL memory map

Name	SPI base address	Description
Core	00000000h	switch core, ingress, forwarding and egress configuration
CGU	00100000h	clock generation unit to control Oscillator, PLLs and clocking
RGU	00100400h	reset generation unit
ACU	00100800h	auxiliary configuration unit

2.1 Ingress stage

A frame is received from a neighboring PHY or MAC on one of the available ports. The xMII block passes received data to the Receive MAC (RX MAC) connected to the reception port. The RX MAC performs low-level checks on the frame data and reports any CRC or MII errors detected to the status and control unit. The frame is immediately discarded when a low-level error is detected. A frame that passes all low-level checks is stored in frame memory in 128 byte segments. The RX MAC captures an ingress timestamp, extracts meta information from the frame and forwards it to the Input Queue (IQ). If a VLAN tag was not embedded in the received frame, the RX MAC block assigns a configured Port VLAN ID and a configured Port VLAN Priority to the frame. The IQ module stores the frame meta information in a deterministic order and passes it to the forwarding stage for further processing. If multiple frames are received at the same time on different ports, the processing order is determined by the port numbers; a frame received on a port with a lower ID is processed before a frame received on a port with a higher ID. Note that this only influences the frame order on the egress stage if multiple concurrently received frames are forwarded to the same destination port.

2.2 Forwarding stage

Once a valid MAC-level frame has been forwarded by the ingress stage, the forwarding stage applies several higher-layer checks on the frame and extracts the forwarding information.

The VLAN Lookup (VLAN_LU) block reads the VLAN information configured for the VLAN ID associated with the frame. If a VLAN tag is embedded in the frame, the block checks if the reception port is configured to be a member of this VLAN. If it is not, the frame is dropped and reported to the status and control unit. It also checks if the VLAN associated with the frame is configured for mirroring or retagging and determines which egress port it should be transferred to.

The Address Lookup (L2ADDR_LU) block extracts forwarding information from the source MAC address and VLAN ID to be used with future frames addressed to this MAC address and VLAN ID combination. The VLAN ID is ignored during this process if shared address learning is activated. This block also looks up the destination MAC address and combines it with the VLAN ID to determine the forwarding information for the frame.

The Policing (L2_POLICE) block meters the incoming frame rate. The switch can be configured to drop packets if the maximum frame rate is exceeded.

The Forwarding (L2_FORW) block uses the information obtained from the other blocks to determine the set of ports to which the frame is forwarded. The switch can be configured to limit the number of egress ports accessible to frames received on a specific ingress port. For example, it is possible to direct that any frame received on a particular ingress port is only forwarded to a specific egress port, regardless of the forwarding information

provided. This block also determines the VLAN priority to be embedded in frames forwarded by the switch as well as the egress priority queue in which a frame is stored on a per priority and per port basis. It also determines if the mirroring port shall be included in the set of ports to which the frame is forwarded, based on the information configured for port-based and VLAN-based ingress and egress mirroring. The L2_FORW block also reserves the required memory space in the partition assigned to the frame by the policing module.

The Loopback Port (LBP) replicates frames that are configured to be retagged based on the associated VLAN configuration. A different VLAN ID is embedded in the replicated frame. It is associated with the same source port as the frame that triggered the replication when fed to the forwarding stage. It follows the configured forwarding rules for this MAC address and VLAN configuration.

2.3 Egress stage

The egress stage recomposes the frame from the data stored in the frame memory and the information gathered by the forwarding stage. It also performs the one-step transparent clock update for IEEE 1588 event frames which have the one-step bit set in the frame header. The Transmit MAC (TX MAC) assigns the frame to the priority queue determined by the forwarding stage. It monitors the number of frames stored in the priority queue. If the maximum number allowed has been exceeded, the frame is dropped and an error condition is signaled to the status and control unit. The TX MAC also performs priority selection based on the strict-priority algorithm and considers whether the credit-based shaper assigned to the priority queue is in the transmission-allowed state.

3. SPI interface

All memory, control and status registers can be accessed via the Serial Peripheral Interface (SPI). The device operates as a slave device in transfer mode 1 with CPOL = 0 and CPHA = 1. Both master and slave must operate in the same mode.

The SJA1105TEL expects a frame format in which the access type, address and data are encoded in a single SPI transaction. The format must conform to the SPI framing described in Section 3.1 to Section 3.3. The device uses a double word addressing scheme.

3.1 Write access

A write access consists of a 32-bit control phase followed by a data phase of up to 64 × 32 bits. The 21-bit address is encoded in control bits[24:4]. The access type is encoded in the MSB, control bit[31]. Both control and data phases are transmitted from MSB to LSB.

Bit[31] is set to 1 to indicate a write operation. A data phase of at least 32 bits, but no more than 64 × 32 bits, is transmitted after the control phase. Both control and data phases are mirrored to SDO during a write operation. Unused control bits must be logic 0.

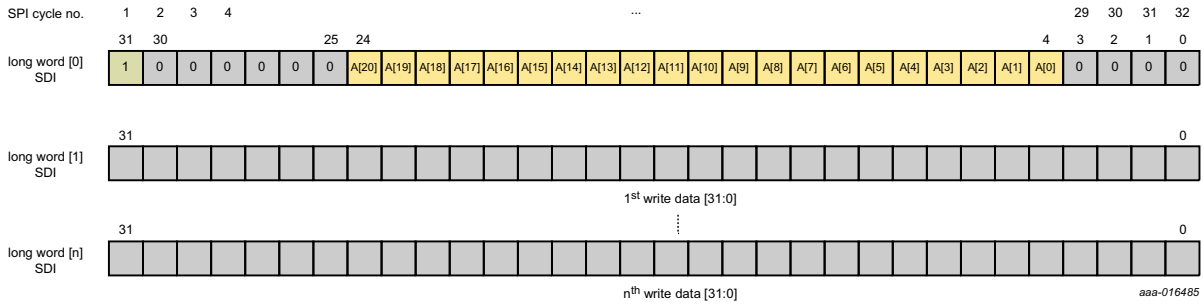

Fig 2. SPI write access to the switch on SDI

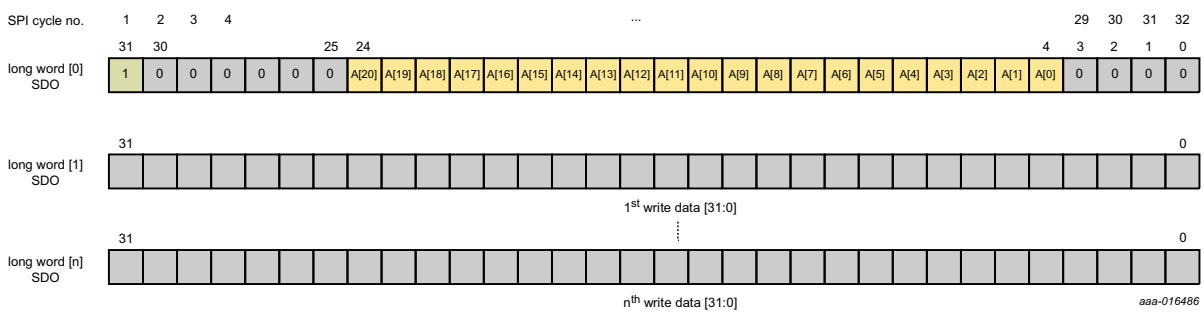

Fig 3. SPI write access, data is read back on SDO

3.2 Read access

A read access is similar to a write access. The access bit (bit[31]) is 0 to indicate a read operation. Bits[30:25] contain the number of 32-bit double words to be read from the device. The device shifts out the corresponding data in the data phase. As with a write access, the address phase is mirrored to SDO. Unused control bits must be logic 0. A read count value of 0 (RC = 0) specifies a read of 64 consecutive words.

Remark: When CGU registers are read, a 64 ns delay must be inserted between the control and data phases to allow the CGU to retrieve the data. Alternatively, the access can be performed at a frequency below 17.8 MHz.

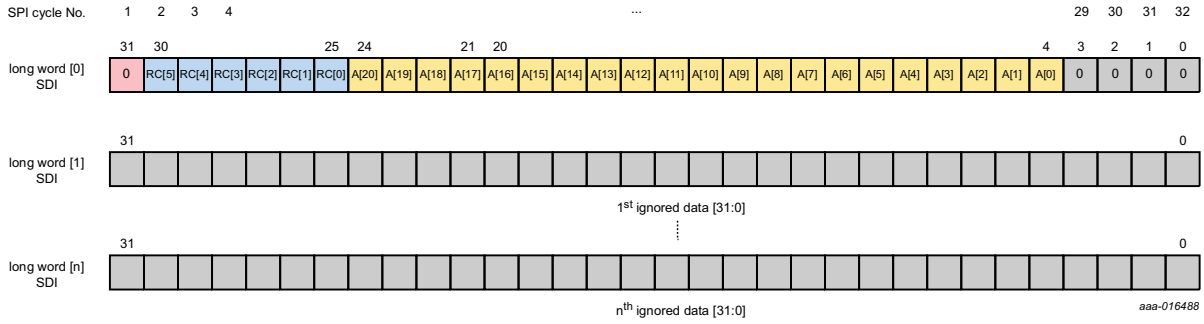

Fig 4. SPI read access to the switch on SDI

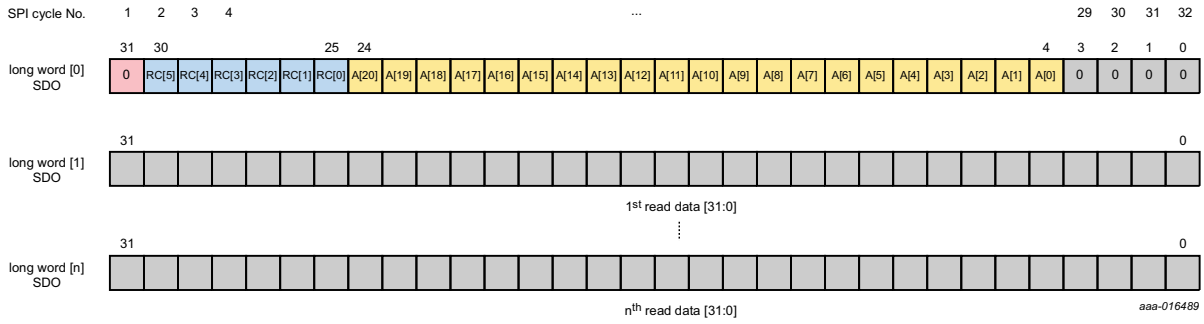

Fig 5. Data is read out (and address information mirrored) on SDO

3.3 SPI read/write timing

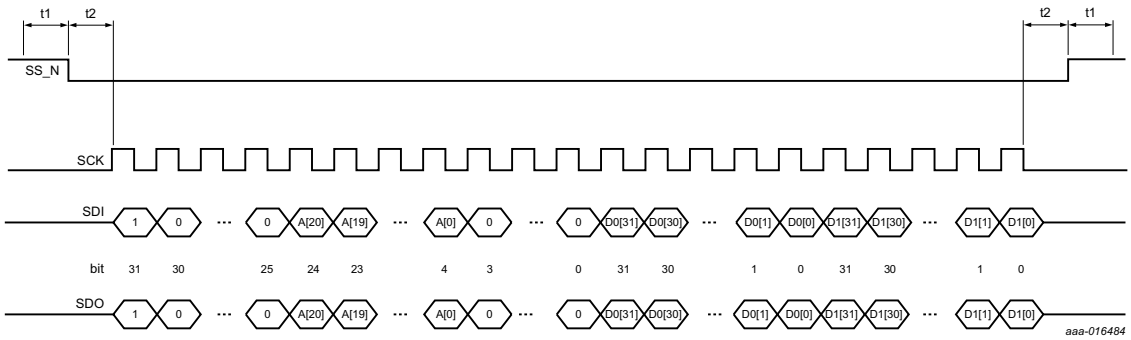

t1: The SPI slave select signal (SS_N) must be stable for at least $0.5 \times t_{clk}$ at the beginning or end of an SPI read or write operation before being asserted/de-asserted.

t2: After the SS_N signal has been asserted at the beginning of an SPI read or write operation, the SPI clock signal (SCK) must be stable for at least $0.5 \times t_{clk}$ before being asserted. At the end of an SPI read or write operation, the SPI clock signal (SCK) must be stable for at least $0.5 \times t_{clk}$ before the SS_N signal is de-asserted.

Fig 6. SPI write access timing

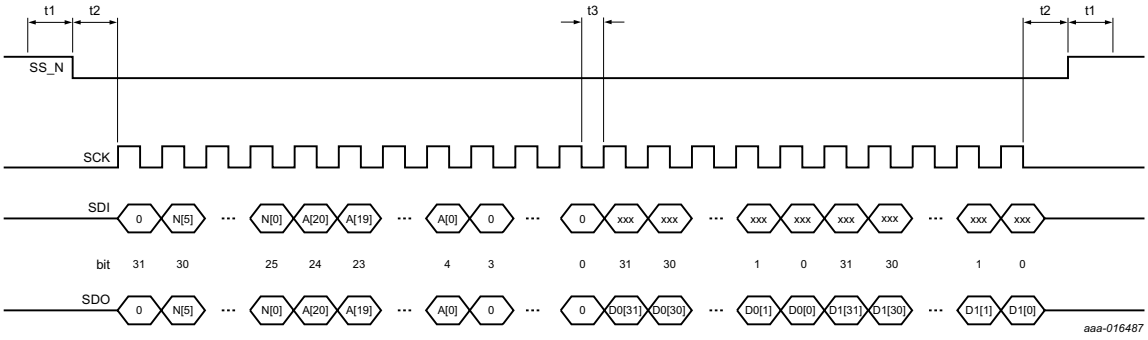

t3: The SPI clock signal (SCK) must be stable for at least half a clock period between the reception of long-word[0] on the SDI input and the transmission of long-word[1] on the SDO output.

Fig 7. SPI read access timing

4. Ethernet switch core interface control

Two distinct interfaces are used to configure the switch core. When the device is powered up, it expects to receive an input stream containing initial setup information over the configuration interface. The initial configuration data sets the port modes, sets up VLANs and defines other forwarding and quality-of-service rules. Once the device is operational, it can be reconfigured at runtime over the programming interface (see Section 5).

This section explains the loader format, the individual configuration blocks (tables) and associated fields. A host microcontroller must upload a valid configuration stream every time the SJA1105TEL is reset or power-cycled. The CONFIGS flag in the Initial device configuration flag register (Table 26) is set once the device has been configured successfully.

4.1 Loading configuration data

Configuration information for the switch core must be loaded at start-up, using the generic loader format as described in Section 4.1.1. The configuration area starts at address 0x2 0000. The entire configuration area is write only. A read access to any address in this area returns arbitrary data.

The configuration data is divided into a number of blocks via the SPI interface, as described in Section 4.2. The blocks can be loaded in any order. The SPI interface is described in Section 3.

4.1.1 Generic loader format

Data is loaded into the configuration area as a continuous stream of 32-bit data. The load operation is initiated by writing the device ID (0x9E00 030E) to the configuration address space at address 0x2 0000. The format for subsequent write operations is illustrated in Figure 8. The configuration data blocks, listed in Table 2, are loaded in turn. The first double-word after the device ID contains the block ID; the second double-word contains length of the first data block to be loaded (i.e. the number of data double-words, excluding the checksum). This is followed by the CRC checksum and the data.

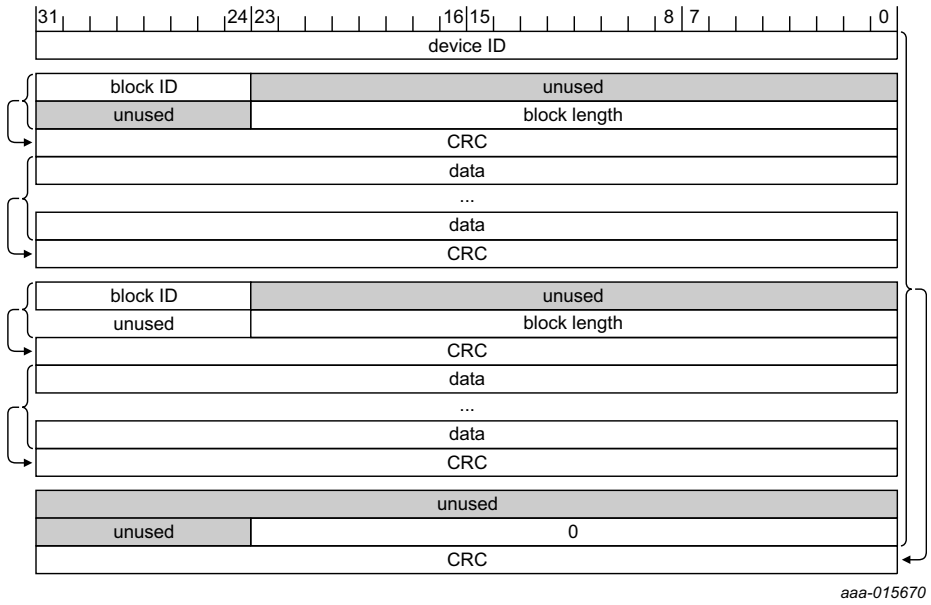

Fig 8. Generic loader format

The data blocks can be loaded in any order. For the VLAN Lookup table, for example, the first 8 bits of the first double-word would contain the block id (0x07). The last 24 bits of the second double-word would define the number of entries (or data double-words) to be loaded. The VLAN Lookup table supports a maximum of 4096 entries, so the block length would be a value between 0x00 0000 and 0x00 2000.

Once all the configuration blocks have been successfully loaded, subsequent write operations are ignored.

A block length of 0 signals the end of the configuration file and a global CRC is expected to follow. Fields labeled 'not used' (e.g bits 0 to 26 in the VLAN Lookup table; see Table 12) are not interpreted by the IP and may be set to any value. However, the values assigned to 'not used' fields must be reflected in the checksum.

Checksums are calculated as CRC-32 Ethernet checksums with the lower bytes of each double-word included first in the CRC calculation. See IEEE 802.3-2015, clause 3 for details on how CRC checksums are calculated for Ethernet frames.

4.2 Switch configuration tables

This section describes the contents of the configuration tables. Configuration data is split into separate configuration blocks as shown in Table 2. These blocks must be loaded using the generic loader format. Blocks can be loaded in any order and a configuration block may be split into several loader format blocks.

Each entry is composed of an integer number of 32-bit double words, padded at the LSB. These padding bits are reserved and should be filled with zeros. An entry in the VLAN Lookup table (Table 12), for example, consists of two double words (64 bits). The upper 37 bits (63:27) are used to store valid configuration data; the lower 27 bits (27:0) are padding bits. An entry in the MAC Configuration table is 224 bits long (seven double words, see Table 14) and has a single padding bit (bit 0).

Entries provided first are written to the lower addresses in the respective table. Unused bits are located the lower end of each entry. An entry in the VLAN Lookup table, for example, contains 37 data bits, with 27 unused bits (see Table 12). The lower 27 bits of the first data double-word received contain the unused bits; the upper 5 bits contain the lower 5 bits of the data entry (bits 31 to 27). The second data double-word received contains the upper 32 bits (bits 63 to 32). This format is repeated for each entry in the table.

Table 2. Configuration tables

Table name	Block ID	Loading mandatory?
Schedule table	00h	no
Schedule Entry Points table	01h	yes, if Schedule table is loaded
VL Lookup table	02h	no
VL Policing table	03h	yes, if VL Lookup table is loaded
VL Forwarding table	04h	yes, if VL Lookup table is loaded
L2 Address Lookup table	05h	no
L2 Policing table	06h	yes, at least one entry
VLAN Lookup table	07h	yes, at least the default untagging VLAN
L2 Forwarding table	08h	yes
MAC Configuration table	09h	yes
Schedule Parameters table	0Ah	yes, if Schedule table is loaded
Schedule Entry Points Parameters table	0Bh	yes, if Schedule table is loaded
VL Forwarding Parameters table	0Ch	yes, if VL Forwarding table is loaded
L2 Lookup Parameters table	0Dh	no
L2 Forwarding Parameters table	0Eh	yes
Clock Synchronization Parameters table	0Fh	no
AVB Parameters table	10h	no
General Parameters table	11h	yes
Retagging table	12h	no
xMII Mode Parameters table	4Eh	yes

4.2.1 Schedule table

Table 3 shows the layout of entries in the Schedule table. A schedule can host up to 8 periods of arbitrary length. These periods are referred to as subschedules. The user may decide to use any number of subschedules but, if a schedule is enabled, subschedule zero must be one of, or the only, active subschedule. The number of subschedules defined, as well as whether the schedule is enabled, is determined by the settings of the Schedule Entry Points table as discussed in Section 4.2.2.

If the schedule is disabled, writing to this configuration block has no effect (and can be omitted). Entries for a particular subschedule must be provided in back-to-back write accesses and must be ordered according to their appearance on the timeline. The order in which the subschedules are provided is arbitrary (but it determines the contents of the Schedule Entry Points table, Table 4). Entries of subschedules with lower indices must be provided prior to entries of subschedules with higher indices. The entries are referenced by the ADDRESS field in the Schedule Entry Points table as well as by the SUBSCHEIND fields of the Schedule Parameters table (Table 15) where the reference equals the ordinal number used to load the respective entry decremented by one (so the first entry of the Schedule table is referenced as 0). The table contains up to 1024 entries. This table is compulsory if entries are provided for the Schedule Entry Points table.

Table 3. Schedule table (block 00h)

Bit	Symbol	Description
63:54	WINSTINDEX	Defines the index in the VL Forwarding table referred to by the WINST flag of the trigger. It is only used when WINST set.
53	WINEND	When set, indicates that the reception window of the entry of the VL Forwarding table indexed by VLINDEX ends here.
52	WINST	When set, indicates that the reception window of the entry of the VL Forwarding table as indexed by WINSTINDEX starts here.
51:47	DESTPORTS	Defines the ports (1 bit per each port) that the respective trigger event applies to. Bits at lower bit positions are assigned to ports with lower port numbers.
46	SETVALID	This flag is used to mark the first trigger to apply on a sampled non-time-triggered input. it allows the delivery order of a sampled non-time-triggered input to be fixed off-line. If such a delivery order is not needed, the flag is set for all triggers having the same VLINDEX value. For time-triggered input, the flag is usually set for all entries (since input and transmit triggers are in phase anyway). The TXEN flag of an entry must be set for this flag to have an effect.
45	TXEN	If this flag is set, the current entry triggers dispatch of an output VL as indexed by the VLINDEX of the entry.
44:36	RESMEDIA	Contains an ‘enable’ flag at the highest bit position and one flag per priority at the output port, where the flag at the lowest bit position is assigned to priority 0. If the ‘enable’ flag is set, the switch stops processing all priority queues whose respective flags are in the lower bit positions of this field for all Ethernet ports that have their respective flag set in DESTPORTS; it enables processing of all priority queues whose respective flags are cleared in the lower bit positions of this field for all Ethernet ports that have their respective flag set in DESTPORTS. The reservation state of ports not having their respective flag set in DESTPORTS does not change. Media reservation is processed individually for each subschedule and a priority queue at a specific port remains suspended as long as at least one subschedule has a reservation pending for this priority at this port. Transmission of locally sourced protocol control frames cannot be blocked by media reservation. At times that the schedule is stopped (not synchronized), all media reservation will be removed until the first trigger after integration that enables media reservation.

Table 3. Schedule table (block 00h) *...continued*

Bit	Symbol	Description
35:26	VLINDEX	Defines the VL Forwarding table index (as discussed in Section 4.2.5) the trigger refers to. The contents of the field are arbitrary when the TXEN flag of the entry is not set. If the TXEN flag of the entry is set, the VLINDEX field must contain a value smaller than the number of entries defined for the VL Forwarding table.
25:8	DELTA	This parameter defines by how much the current trigger event precedes the next trigger event of the same subschedule in multiples of 200 ns. The user must ensure that no two entries in this, or another, subschedule, ever fire at the same time. To avoid the former, a value of zero is not allowed for this field. Schedule analysis is needed to prevent the latter (this analysis must take the contents of the Schedule Entry Points table into account).
7:0	not used	

4.2.2 Schedule Entry Points table

Table 4 shows the layout of an entry in the Schedule Entry Points table. This table establishes the link between the synchronization algorithm and the schedule. If the user does not load this table, the schedule will not be active (whether or not the user loads the Schedule table). Each entry point into the schedule consists of 8 entries in the Schedule Entry Points table. The entries included in an entry point must be sorted in ascending order (i.e. those having smaller values must be loaded before those with larger values) according to their respective DELTA values. Undefined entries (if the subschedules are not all used) must be provided after the valid entries. The entry point must contain exactly one valid entry for each active subschedule (the table thus provides space for up to 2048 entries). No two entries of an entry point can contain identical DELTA values and zero is not allowed. The schedule starts when the local clock of the synchronization algorithm wraps with the entry point as indexed by the integration cycle number contained in the PCF used for integration.

Table 4. Schedule Entry Points table (block 01h)

Bit	Symbol	Description
31:29	SUBSCHINDX	This field defines the subschedule the respective entry point refers to. The value provided here must identify an active subschedule as defined by the SUBSCHEIND field of the Schedule Parameters configuration block. Any two entries of an entry point assigned to an active subschedule must contain different values for this field. An arbitrary value may be provided for inactive subschedules.
28:11	DELTA	This field defines the delay before this entry fires in multiples of 200 ns. Any two entries of an entry point assigned to an active subschedule must contain different values for this field. A value of zero is not allowed. An arbitrary value may be provided for inactive subschedules.
10:1	ADDRESS	This field provides the index in the Schedule table holding the event to fire at the respective time. This index must be within the range assigned to the subschedule as identified by the SUBSCHINDX field of this entry and defined by the respective SUBSCHEIND field of the Schedule Parameters configuration block. An arbitrary value may be provided for inactive subschedules.
0	not used	

4.2.3 VL Lookup table

Table 5 and Table 6 show the layout of an entry in the VL Lookup table. The table establishes the link between the stream identifier and the entry of the VL Policing table. The layout depends on the setting of the VLLUPFORMAT flag in the General Parameters configuration block (see Table 22). If the user does not load this table, all critical input traffic is dropped. The table has 1024 entries.

If VLLUPFORMAT is set to 0, the entries in the VL Lookup table must be sorted in ascending order (i.e. the smallest value must be loaded first) according to the following sort order: MACADDR, VLANID, PORT, VLANPRIOR. If VLLUPFORMAT is set to 1, the entries in the VL Lookup table must be sorted in ascending order according to the following sort order: VLLD, PORT.

Table 5. VL Lookup table when VLLUPFORMAT = 0 (block 02h)

Bit	Symbol	Description
95:91	DESTPORTS	This field contains the set of destination ports to which a frame matching this entry is forwarded if ISCRITICAL is cleared.
90	ISCRITICAL	When this field is set, the configured entry is treated as rate-constrained or time-triggered; if this field is cleared, the configured entry is a static configuration of a best-effort flow and is treated as best-effort.
89:42	MACADDR	This field contains the destination MAC address to be associated with the respective table index.
41:30	VLANID	This field contains the VLAN ID to be associated with the respective table index position.
29:27	PORT	This field contains the number of the input port the respective stream (as identified by MACADDR, VLANID and VLANPRIOR) is allowed to access. A stream may be allowed on any number of ports. If allowed on more than one port, it will have dedicated entries within the VL Policing table configuration block for each eligible source port.
26:24	VLANPRIOR	This field contains the VLAN Priority to be associated with the respective table index position.
23:0	not used	

Table 6. VL Lookup table when VLLUPFORMAT = 1 (block 02h)

Bit	Symbol	Description
95:91	EGRMIRR	All traffic matching the stream identified by the VLLD and PORT fields of the entry, and routed to any of the ports having its flag asserted in this field, is routed to the mirror port (as defined by the MIRRORPORT field of the General Parameters configuration block).
90	INGRMIRR	If this flag is set, all traffic matching the stream identified by the VLLD and PORT fields of the entry is routed to the mirror port as defined by the MIRRORPORT field of the General Parameters configuration block.
89:58	not used	
57:42	VLLD	This field contains the VL ID to be associated with the respective table index position.
41:30	not used	
29:27	PORT	This field contains the number of the port the respective VL ID is allowed to access. A VL ID may be allowed on any number of ports. If allowed on more than one port, it will have dedicated entries in the VL Policing table configuration block for each source port.
26:0	not used	

4.2.4 VL Policing table

Table 7 and Table 8 show the layout of an entry in the VL Policing table. This table provides timing and sizing rules for critical traffic. The rules to be applied depend on the type of traffic and, consequently, the table entries have different layouts depending on whether they are used to police time-triggered traffic or rate-constrained traffic. The table has 1024 entries.

Table 7. VL Policing table - time-triggered VLs (block 03h)

Bit	Symbol	Description
63	1	A value of 1 at this bit position indicates that the entry defines a time-triggered VL.
62:52	MAXLEN	This field defines the maximum length of frames of this entry in bytes including all Ethernet overhead (6-byte destination MAC address, 6-byte source MAC address, 2-bytes EtherType field, 4-byte frame checksum). The maximum allowed value for this field is 2043.
51:42	SHARINDX	Contains the index in the VL Forwarding table to be used with this entry. This index is usually the index of the entry itself. All entries for time-triggered VLs having identical values set for this field share memory space. They can be used to implement a pick-first-valid redundancy mechanism on time-triggered VLs on input. The dispatch policy is determined by the entry in the VL Forwarding table indexed by this field.
41:0	not used	

Table 8. VL Policing table - rate-constrained VLs (block 03h)

Bit	Symbol	Description
63	0	A value of 0 at this bit position indicates that the entry defines a rate-constrained VL.
62:52	MAXLEN	This field defines the maximum length of frames of this entry in bytes including all Ethernet overhead (6-byte destination MAC address, 6-byte source MAC address, 2-bytes EtherType field, 4-byte frame checksum). The maximum allowed value for this field is 2043.
51:42	SHARINDX	Contains the index within this table that holds the BAG and JITTER fields to be used for this entry. Usually, this index is the index of the entry. For shared BAGs, all entries in this table that share the BAG must contain the same value for this field and this value must be the index of one of the entries sharing the BAG. The values of the BAG and JITTER fields of entries that are not pointed to by any entry are not used. Unlike time-triggered VLs, the contents of this field is not used to access the VL Forwarding table. Rather, the index received from searching the VL Lookup table is used to access the VL Forwarding table for VLs being policed in rate-constrained fashion.
41:28	BAG	The bandwidth allocation gap (BAG) value to be used for this entry in multiples of 100 μs. A value of zero disables the bag check. In the latter case, JITTER must also be set to zero.
27:18	JITTER	The bandwidth allocation gap (BAG) jitter value to be used for this entry in multiples of 10 μs. The value provided for this field must not be larger (in seconds) than the value provided for the BAG field. For example, if the value of BAG was 5, the maximum allowed value for JITTER would be 50. This is deemed sufficient to police traffic received from an end system, which, according to the ARINC 664 p7 specification, exhibits a maximum jitter of 500 μs at a minimum BAG value of 1 ms. Should jitter accumulate in a multi-hop network to exceed the BAG value, policing must be disabled (by setting BAG to zero) starting at the first switch where jitter exceeds the BAG.
17:0	not used	

4.2.5 VL Forwarding table

Table 9 shows the layout of an entry in the VL Forwarding table. This table provides forwarding definitions for critical traffic. The table has 1024 entries.

Table 9. VL Forwarding table (block 04h)

Bit	Symbol	Description
31	TYPE	A value of 1 at this bit position indicates that the entry defines a time-triggered VL (i.e., frames of the VL are dispatched to destination ports in response to triggers from the schedule). A value of 0 at this bit position indicates that the entry defines a rate-constrained VL (i.e., frames of the VL are immediately dispatched to destination ports).
30:28	PRIORITY	Priority at the output port for frames matching this entry where larger values indicate higher priority.
27:25	PARTITION	VL memory partition that frames matching this entry draw from.
24:20	DESTPORTS	Defines the ports that frames matching this entry are routed to (1 bit per port). Bits at lower bit positions are assigned to ports with lower port numbers. In the case of rate-constrained VLs, these flags define the ports the frame is dispatched to when all policing checks have been passed. In the case of time-triggered VLs, this field should be set to the bit-wise-or of the DESTPORTS fields of all the Schedule table entries that have set VLINDEX to point to this entry of the VL Forwarding table and have TXEN asserted.
19:0	not used	

4.2.6 L2 Address Lookup table

Table 10 shows the layout of an entry in the L2 Address Lookup table. Parts of the table can be statistically configured prior to dynamic address learning. Unlike other configuration blocks, loading of this block must not start before the L2BUSYS flag in the status area has been cleared (see Table 28). Entries in this table share memory with entries dynamically learned during operation. However, loaded entries never time-out and cannot be replaced by learned entries even in the case of a hash conflict. Physically, the memory used to store the lookup table has 1024 entries, organized in 256 rows each having 4 columns. The computed hash value maps to the row, so each hash conflict can be resolved four times.

Table 10. L2 Address Lookup table (block 05h)

Bit	Symbol	Description
95:84	VLANID	The VLAN ID associated with this entry. VLANID is only included in the lookup process if SHARED_LEARN in Table 18 is cleared, otherwise this parameter is ignored. If SHARED_LEARN is set, the VLANID is set to zero for the hash computation.
83:36	MACADDR	the MAC address associated with this entry
35:31	DESTPORTS	Defines the ports (1 bit per port) to which frames carrying MACADDR as destination MAC address are forwarded. Bits at lower bit positions are assigned to ports with lower port numbers.

Table 10. L2 Address Lookup table (block 05h) *…continued*

Bit	Symbol	Description
30	ENFPORT	If this flag is set, MACADDR is enforced as the source MAC address on ports having their flag set in DESTPORTS; i.e., an Ethernet frame with MACADDR as its source MAC address that is received on a port other than those set in DESTPORTS is dropped.
29:20	INDEX	Contains the address in physical memory where this entry is stored. The physical address of an entry is calculated as INDEX = 4 × hash (MACADDR, VLANID) + i, where i Є {0;1;2;3}. If more than four MAC address/VLAN IDs pairs produce identical hashes, a different value for the POLY must be chosen or MAC addresses and VLAN IDs of endpoints or default VLANs of switch ports must be changed. If SHARED_LEARN is set, the VLANID portion of the hash computation is set to zero, regardless of the actual VLANID of the frame. The 8-bit CRC hash computation operates on 62 bits consisting of the 48-bit MAC address, 10-bit VLANID and 4 bits padded with 0.
19:0	not used	

4.2.7 L2 Policing table

Table 11 shows the layout of an entry in the L2 Policing table. This table defines traffic policing rules for each port individually, along with a priority value for each port and switch broadcast traffic. The table has 45 entries. Ethernet frames received on mappings for which the user has not provided an entry are automatically mapped to entry 0 (all such traffic is dropped). The entry to which an incoming frame maps is determined in the following way: if the incoming frame is classified as broadcast, the matching entry is 40 + PORT (where PORT is the physical port number between 0 and 4); if the frame is not classified as broadcast, the matching entry is 8 × PORT + VLANPRIO where VLANPRIO is the VLAN priority value associated with the frame.

The switch allows traffic from different ports or priorities to share common policing blocks. Resolving the actual policing block is a two stage process. First, the device determines the entry as discussed above (i.e. 8 × PORT + VLANPRIO or 40 + PORT). The SHARINDX field of this entry is then used to determine the policing block. This SHARINDX field can point to any of the 45 available policing blocks.

The algorithm used for bandwidth budgeting works as follows. Each policing block contains the parameters SMAX and RATE. Initially, the bandwidth credit of an entry gets set to SMAX. When a valid Ethernet frame mapping to this entry is received, the value of the bandwidth credit is decreased by the number of bytes in the frame (including Ethernet header and checksum). At times when no traffic associated with this entry is received, the bandwidth credit gets increased by the value of RATE every 8 μs, to a maximum of SMAX. An associated frame gets dropped if the resulting value of the bandwidth credit is less than or equal to zero. This makes it possible to control the traffic rate individually for each port. In addition to the rate, each entry specifies the maximum length of frames associated with this entry and the memory partition that gets credited for this frame. This makes it possible to partition the maximum amount of frame memory available for different traffic classes.

Table 11. L2 Policing table (block 06h)

Bit	Symbol	Description
63:58	SHARINDX	This field contains the index pointing to the policing entry associated with this frame. It is a pointer to the L2 Policing table itself and can be used to merge several traffic classifications in one combined policing entry. As an example, if all incoming L2 traffic from port 0 is to be policed by policing block 0, the value of SHARINDX for entries 0 through 7 must be set to 0.
57:42	SMAX	This field contains the maximum burst size for received frames which map to this entry in bytes. Its value is used to initialize the bandwidth budget for this entry on start-up. This field defines the maximum bandwidth budget when no traffic associated with this entry has been received for a long time.
41:26	RATE	This field contains the rate at which the bandwidth budget of traffic associated with this entry is credited when the port does not receive any traffic. The budget is credited RATE divided by 64 bytes every 8 μs with a maximum value of SMAX. A port allowed to source traffic at 1 Gbit/s would thus have a value of 64000 set for this field.
25:15	MAXLEN	This field defines the maximum length of frames of this entry in bytes including all Ethernet overhead (6-byte destination MAC address, 6-byte source MAC address, 2-bytes EtherType field, 4-byte frame checksum). The maximum allowed value for this field is 2043.
14:12	PARTITION	Memory partition that Ethernet frames matching this entry will draw from.
10:0	not used	

4.2.8 VLAN Lookup table

Table 12 shows the layout of an entry in the VLAN Lookup table. This table is used to statically configure VLAN information. A table entry defines the ports that are members of a specific VLAN. It also defines the broadcast domain together with the set of ports on which a VLAN tag has to be inserted or removed on egress. The table supports 4096 entries. If no entry is loaded, the switch is initialized with default entry: VING_MIRR:0, VEGR_MIRR:0, VMEMB_PORT:0x1F, VLAN_BC:0x1F, TAG_PORT:0x1F, VLANID:0.

Table 12. VLAN Lookup table (block 07h)

Bit	Symbol	Description
63:59	VING_MIRR	All traffic tagged with VLANID and received on any of the ports whose flag is asserted in this field is forwarded to the mirror port as defined by the MIRR_PORT field of the General Parameters configuration block.
58:54	VEGR_MIRR	All traffic tagged with VLANID and forwarded to any of the ports whose flag is asserted in this field is forwarded to the mirror port as defined by the MIRR_PORT field of the General Parameters configuration block.
53:49	VMEMB_PORT	Defines the set of ports on which a frame tagged with the respective VLAN ID may be received. All bits must be set in order to deactivate VLAN-based ingress port admission.
48:44	VLAN_BC	This field restricts the broadcast domain of the specific VLAN. That means that, if a bit is cleared, a frame tagged with the specific VLAN ID cannot reach the respective port. All bits must be set to deactivate reachability limitations for certain VLANs.
43:39	TAG_PORT	Defines if a frame associated with the respective VLAN ID is transmitted untagged (the flag of these ports would be cleared in TAG_PORT), i.e., not containing an IEEE 802.1Q VLAN tag field, or transmitted with a tag (the flag of these ports would be set in TAG_PORT). As each untagged frame gets tagged on ingress with the port VLAN ID, all bits must be cleared in order to receive untagged frames at the output.
38:27	VLANID	The VLAN ID associated with this entry.
26:0	not used	

4.2.9 L2 Forwarding table

Table 13 shows the layout of an entry in the L2 Forwarding table. This table defines the mapping of ingress VLAN priority values to egress VLAN priority values as well as the mapping of egress VLAN priority values to priority queues physically available on the transmission ports. In addition, this table is used to define forwarding limitations for each ingress port.

The first five entries in the table are used for a per-port based remapping of the ingress priority values to egress priority values. For instance, the value of VLAN_PMAP in entry 0 defines the mapping of either a received or per-port assigned ingress priority value p_i to an egress priority value p_o for frames received on port 0 by assigning p_o = VLAN_PMAP[p_i]. This means that p_o will be used as the PCP (Priority Code Point) value on all egress ports forwarding the frame with a VLAN tag included (obtained by the TAG_PORT parameter in the VLAN configuration, see Table 12).

The last eight entries in the table are used for a per-egress priority-based mapping of logical priority values to physical priority queues of the different ports. For the previously obtained egress priority value p_o, the resulting mapping to priority queues on each port i is obtained by assigning qi = VLAN_PMAP[i], where qi is the priority queue used for p_o on port i. For example, to map priority value p_o = 4 to priority queue 1 on port 0 and to priority queue 2 on port 3, the value of VLAN_PMAP for entry 5 + 4 = 9 must be set to 1 for index 0 and to 2 for index 3.

Table 13. L2 Forwarding table (block 08h)

Bit	Symbol	Description
63:59	BC_DOMAIN	Only valid for the first five entries in the table. Defines the broadcast domain of the port associated with the entry. Each port is assigned a bit in this field with the LSB mapping to port 0. Broadcast Ethernet frames received from the respective port are forwarded to the ports whose flags are set in this vector. The flag of the port associated with the entry itself must be cleared (to prevent loops).
58:54	REACH_PORT	Only valid for the first five entries in the table. Defines which ports can be reached by traffic received on the port associated with the entry. Each port is assigned a bit in this field with the LSB mapping to port 0. If a frame is received on the port associated with the entry and its destination MAC address is known (i.e. is contained in the L2 Address Lookup table), the frame is forwarded to the destination port only if the flag of the destination port is set in this field.
53:49	FL_DOMAIN	Only valid for the first five entries in the table. Defines the destination ports of unknown traffic at the port associated with this entry. Each port is assigned a bit in this field with the LSB mapping to port 0. If an Ethernet frame (that is not a broadcast frame) is received on the port associated with the entry and its destination MAC address is not known (i.e. is not contained in the L2 Address Lookup table), the frame is forwarded to those ports that have their respective flag set in this field. The flag of the port associated with the entry itself must be cleared (to avoid loops).
48:46	VLAN_PMAP[7]	For the first five entries in the table, this value defines the ingress VLAN priority remapping. The source port associated with the incoming frame is used as an index into the table, allowing ingress VLAN priority to egress VLAN priority mapping for each port. The result of the mapping is embedded in the transmitted frame on all ports included in the tagged set of the VLAN associated with the frame. For indices 5 to 12, this field contains the mapping of egress VLAN priority (determined by the first 12 entries in the table) to physical priority queues. In this case, the destination port is used as index into VLAN_PMAP.
:	:	
:	:	
:	:	
27:25	VLAN_PMAP[0]	
10:0	not used	

4.2.10 MAC Configuration table

Table 14 shows the layout of an entry in the MAC Configuration table. This table is used to define the configuration parameters for each switch port. The table contains five entries. However, all entries will not need to be loaded if some of the ports are not used. If N is the largest port number used in a specific configuration, N + 1 entries must be provided for this table. If any of the ports with a port number less than N are not used, dummy values for these unused ports must be provided.

Table 14. MAC Configuration table (block 09h)

Bit	Symbol	Description
223:215	TOP[7]	The fields TOP, BASE and ENABLED are used to define the maximum number of frames of the respective priority that may be waiting in the output queue of the associated Ethernet port. If the respective priority is enabled at the port as indicated by ENABLED being set, then the value of TOP must be at least as large as the value configured for BASE for the priority at the relevant port. The maximum number of frames of the respective priority at the port is then TOP minus BASE plus one. No two enabled priorities at the same port may have overlapping intervals set for the TOP and BASE parameters. If ENABLED is not set for a priority, the values configured for the TOP and BASE parameters are arbitrary. The maximum total value allowed for these parameters is 511.
214:206	BASE[7]	
205	ENABLED[7]	
:	:	
90:82	TOP[0]	
81:73	BASE[0]	
72	ENABLED[0]	
71:67	IFG	This parameter allows the standard Ethernet IFG of 12 bytes to be extended for output traffic on this port.
66:65	SPEED	Sets the port speed. 11 sets it to 10 Mbit/s; 10 sets it to 100 Mbit/s; 01 sets it to 1 Gbit/s; 00 allows the host to set the speed dynamically for this port.
64:49	TP_DELIN	Used to set a correction for updating the transparent clock of IEEE 1588v2 one-step event messages at the input port in multiples of 8 ns.
48:33	TP_DELOUT	Used to set a correction for updating the transparent clock of IEEE 1588v2 one-step event messages at the output port in multiples of 8 ns.
32:25	MAXAGE	Defines the maximum allowed age of critical traffic frames on the output in multiples of 4000 μs. The clock used to measure the age has a resolution of the same granularity. The minimum reasonable setting for this parameter is thus one (otherwise frames may get dropped depending on whether the clock wraps between reception and transmission). A value of 255 turns off the age check.
24:22	VLANPRIO	Defines the IEEE 802.1Q priority value used to prioritize an untagged frame on this port. The value is in the range of 0 to 7. This value is used as an index to VLAN_PMAP[7:0] in Table 13 to resolve the ingress priority to egress priority and ultimately physical priority queue mapping. This value is also used to calculate the index for the rate-policing entry to which untagged frames are assigned to (see Table 11).
21:10	VLANID	Defines the VLAN ID used to tag untagged incoming frames on this port. Values can be chosen arbitrarily in the range 0 to 4095. The respective entry in the VLAN Lookup table must be loaded and the flag of the port set in VMEMB_PORT. Otherwise all untagged frames received on the port will trigger WRONGPORTS or VNOTFOUNDS status errors (see Table 32) and cause N_VLANERR to be increased (see Table 50).
9	ING_MIRR	If this flag is set, all traffic received on this port is forwarded to the mirror port as defined by the MIRR_PORT field of the General Parameters configuration block.
8	EGR_MIRR	If this flag is set, all traffic forwarded to this port except for locally generated PCFs is forwarded to the mirror port as defined by the MIRR_PORT field of the General Parameters configuration block.
7	DRPNONA664	If this flag is set, frames carrying an EtherType other than 800h are dropped on input at this port. This includes VLAN-tagged frames. Only non-VLAN IP frames are accepted at the port.

Table 14. MAC Configuration table (block 09h) *…continued*

Bit	Symbol	Description
6	DRPDTAG	When this flag is set, double-tagged ingress traffic is dropped at the respective port (i.e. traffic that has a TPID defined in the General Parameters configuration block for either an outer or inner tag as well as traffic containing TPID2 in the outer tag - whether an inner tag exists or not). Flag affects L2 traffic only. Management traffic flows to the port regardless of the state of the INGRESS flag.
5	DRPUNTAG	If this flag is set, untagged ingress traffic is dropped at the respective port.
4	RETAG	When set, this flag enables retagging (using VLANID configured for the respective port but maintaining the priority value) of priority-tagged input on the respective port.
3	DYN_LEARN	This flag enables address learning at the respective port when set. Note that learning is independent of whether input traffic is enabled.
2	EGRESS	This flag enables output on the respective port when set.
1	INGRESS	This flag enables input on the respective port when set.
0	not used	-

4.2.11 Schedule Parameters

Table 15 shows the layout of an entry in the Schedule Parameters configuration block. This configuration block is compulsory if the user provides values for the Schedule Entry Points table (Table 4).

Table 15. Schedule Parameters table (block 0Ah)

Bit	Symbol	Description
95:86	SUBSCHEIND(7)	These fields define the last entry of the respective subschedule in the Schedule table. If the schedule is used at all, subschedule zero will always be active. For all other subschedules, the field must be set to the value of the active subschedule with the largest index that has an index smaller than the respective inactive subschedule when the subschedule is not active. For example, if subschedules 0 and 3 are used while all other subschedules are unused, and subschedule 0 contains 5 schedule entries while subschedule 3 contains 8, then SUBSCHEIND(0), SUBSCHEIND(1) and SUBSCHEIND(2) would be set to 4 while all other entries would be set to 12. The array must contain a valid entry for each active subschedule.
:	:	
:	:	
:	:	
:	:	
25:16	SUBSCHEIND(0)	
15:0	not used	

4.2.12 Schedule Entry Points Parameters

Table 16 shows the layout of an entry in the Schedule Entry Points Parameters configuration block. This configuration block is compulsory if the user provides values for the Schedule Entry Points table (Table 4).

Table 16. Schedule Entry Points Parameters table (block 0Bh)

Bit	Symbol	Description
31:30	CLKSRC	This field defines the clock source driving the schedule: 00 deactivates schedule execution 01 configures the switch to run in standalone mode, i.e., immediately integrating into the schedule without applying any clock correction 10 selects SAE AS6802 as clock source 11 selects the internal PTP clock which is under the control of the host as clock source.
29:27	ACTSUBSCH	This field defines the number of active subschedules. Subschedule zero must always be defined if the Schedule Entry Points table is loaded.
26:0	not used	

4.2.13 VL Forwarding Parameters

Table 17 shows the layout of an entry in the VL Forwarding configuration block.

Table 17. VL Forwarding Parameters table (block 0Ch)

Bit	Symbol	Description
95:86	PARTSPC(7)	These fields define the maximum amount of frame memory that a VL memory partition can use. A VL memory partition is a set of VLs that store their frames in shared memory. When a frame is received and passes all policing checks, it draws from the VL memory partition as identified by the PARTITION field of the respective entry of the VL Forwarding Table (Table 9). Once the frame has completed transmission to all ports, the memory needed to store the frame is credited to the respective VL memory partition. The parameter specifies the number of 128-byte memory blocks contained in the memory partition. A frame requires as many blocks as are needed to ensure that the sum of the bytes in the block is greater than or equal to the number of bytes contained in the frame, including Ethernet header and checksum but excluding a VLAN tag (if any). A block cannot be shared between frames.
:	:	
:	:	
:	:	
:	:	
:	:	
:	:	
25:16	PARTSPC(0)	
15	DEBUGEN	Mirroring and retagging are available for critical traffic only if this flag is set.
14:0	not used	

4.2.14 L2 Lookup Parameters

Table 18 shows the layout of the L2 Lookup Parameters block. Parameters that control the address learning process are loaded into this block. It specifies how long dynamically learned entries are valid. It also defines the maximum number of entries in the address lookup table that are available for the dynamic address learning process (in order to reserve space for entries used by higher layer protocols like MMRP, SRP or IGMP). It also specifies if the MAC addresses learned are shared among all VLANs or are distinct for every VLAN.

Table 18. L2 Lookup Parameters table (block 0Dh)

Bit	Symbol	Description
31:17	MAXAGE	This parameter defines the time-out for dynamically learned entries in multiples of 10 ms. An entry in the address hash table that reaches this age is forgotten. The timer is started every time a new entry is learned. The timer is restarted when the reception of another frame with an identical source MAC address confirms an existing entry. The aging mechanism does not affect entries in the L2 Address Lookup table that have been loaded during configuration. If the parameter is set to 0, aging is deactivated and learned addresses are not forgotten until reset or when changed by the host.
16:14	DYN_TBSZ	The value specified in this field limits the number of entries in the L2 Address Lookup table available for dynamic address learning to DYN_TBSZ $\times 2^8$. This parameter must be set to 0 to deactivate dynamic address learning. To make all entries in the L2 Address Lookup table eligible for dynamic address learning, this parameter must be set to four. The number of entries available for dynamic address learning can be limited to ensure that the application has $(4 - \text{DYN\_TBSZ}) \times 2^8$ entries available for higher layer protocols such as Multiple MAC reservation (MMRP) or IGMP (snooping).
13:6	POLY	This parameter defines the CRC polynomial used to compute the hash value from a MAC/VLAN pair. The polynomial is expected in Koopman notation and is provided in coefficients of degrees 1 to 8 with lower bit positions containing the coefficients of smaller degrees. The coefficient of degree 0 is hard-wired to 1. MAC addresses are fed MSB-first (i.e. first bit on the wire is fed first) to the linear-feedback chain constructed from the polynomial. The polynomial 0x97 resembles the following polynomial: $2^8 + 2^5 + 2^3 + 2^2 + 2^1 + 1$.
5	SHARED_LEARN	A value of 0 specifies that the VLAN ID is included in the hash computation. If this parameter is set to 1, the hash computation uses 0 instead of the actual VLAN ID, whether the frame is tagged or not.
4	NO_ENF_HOSTPRT	This parameter, when asserted, turns off port enforcement for management traffic received at the host port. All traffic producing a match with the MAC_FLT[i] and MAC_FLTRES[i] parameters of the General Parameters configuration block (see Table 22) is considered management traffic. The HOST_PORT parameter in the General Parameters configuration block identifies the host port. This flag is ignored if HOST_PORT does not contain a valid port number. Port enforcement is enabled for a MAC address by setting the ENFPORT flag of the respective entry in the L2 Address Lookup table. The NO_ENF_HOSTPRT flag overrules the ENFPORT flag for management traffic received at the host port.
3	NO_MGMT_LEARN	This parameter, when asserted, turns off address learning for management traffic received at the host port. Address learning includes learning a new address as well as updating a previously learned address (i.e. resetting its age and set receive port value). All traffic producing a match with the MAC_FLT[i] and MAC_FLTRES[i] parameters of the General Parameters configuration block (see Table 22) is considered management traffic. The HOST_PORT parameter in the General Parameters configuration block identifies the host port. This flag is only used when HOST_PORT contains a valid port number.
2:0	not used	

4.2.15 L2 Forwarding Parameters

Table 19 shows the layout of the L2 Forwarding Parameters block. This block defines the memory space available for traffic mapped to any of the available memory partitions through the configuration in the L2 Policing table (Table 11). This block also allows for the dynamic reconfiguration of priority queue mapping in order to ensure that low-priority traffic is not assigned to high-priority queues.

Table 19. L2 Forwarding Parameters table (block 0Eh)

Bit	Symbol	Description
95:93	MAX_DYNP	This field defines the maximum VLAN_PMAP[7:0] values (see Table 13) that will be accepted for dynamic updates. Note that this parameter only applies to dynamic updates. Larger values are accepted during configuration load entries. This parameter also only affects the mapping of egress priority values to physical priority queues (the latter 8 entries in the L2 Forwarding table; Table 13); the mapping of ingress priority values to egress priority values is not restricted.
92:83	PART_SPC[7]	These fields define the maximum amount of frame memory that a memory partition can use. A memory partition is used by a set of ports that store their frames in shared memory. When a frame is received and passes all policing checks, it draws from the memory partition as identified by the PARTITION field of the respective entry in the L2 Policing table (Table 11). Once the frame has completed transmission to all ports, the memory needed to store the frame is credited to the respective memory partition. The parameter specifies the number of 128-byte memory blocks contained in the memory partition. A frame requires as many blocks as needed to ensure that the sum of the bytes in the block is greater than or equal to the number of bytes contained in the frame, including Ethernet header and checksum but excluding any VLAN tag. A block cannot be shared between frames. The total number of assigned partitions must never exceed 910 if retagging is used or 929 if retagging is not used.
:	:	
:	:	
:	:	
:	:	
:	:	
:	:	
22:13	PART_SPC[0]	
12:0	not used	

4.2.16 Clock Synchronization Parameters

Table 20 shows the layout of an entry in the Clock Synchronization Parameters Configuration block.

Table 20. Clock Synchronization Parameters table (block 0Fh)

Bit	Symbol	Description
415:368	ETSSRCPCF	This field provides the Ethernet MAC address that is used as source address for protocol control frames generated by the switch when the SWMASTER flag is set.
367:364	WAITTHSYNC	If an out-of-schedule input protocol control frame has this number, or more, membership flags set, the synchronization engine re-sets its time base to the frame.
363:344	WFINTMOUT	This parameter defines the maximum length of time the synchronization engine remains in WAIT_FOR_IN state, in multiples of 8 ns.
343:340	UNSYTOTSYTH	The value provided for this field defines the minimum number of in-schedule clock masters needed to convince the synchronization engine to switch from UNSYNC state to TENTATIVE SYNC state.
339:336	UNSYTOSYTH	This parameter defines the minimum number of flags that must be set in the membership field of an input integration frame to cause the synchronization algorithm to switch from UNSYNC state to SYNC state.
335:332	TSYTOSYTH	This value defines the minimum number of in-schedule clock masters needed to switch the synchronization engine from TENTATIVE SYNC state to SYNC state.
331:328	TSYTH	This parameter defines the minimum number of in-schedule clock masters needed to keep the synchronization engine in TENTATIVE SYNC state.
327:324	TSYTOUSYTH	This field defines the minimum number of out-of-schedule clock masters needed to convince the synchronization engine to switch from TENTATIVE SYNC state UNSYNC state.
323:320	SYTH	The value provided for this field defines the minimum number of in-schedule clock masters needed to keep the synchronization engine in SYNC state.
319:316	SYTOUSYTH	This parameter defines the minimum number of out-of-schedule clock masters needed to convince the synchronization engine to switch from SYNC state to INTEGRATE state or UNSYNC state, respectively.
315:314	SYPRIORITY	Used to check the sync priority of input protocol control frames and to set the sync priority field of output protocol control frames.
313:310	SYDOMAIN	Used to check the sync domain of input protocol control frames and to set the sync domain field of output protocol control frames.
309:306	STTH	This parameter defines the minimum number of in-schedule clock masters needed to keep the synchronization engine in STABLE state; if the number of in-schedule clock masters remains below this value for as many cycles as defined by the NUMUNSTBCY parameter, the synchronization engine switches to INTEGRATE state.
305:302	STTOINTTH	This field defines the minimum number of out-of-schedule clock masters needed to convince the synchronization engine to switch from STABLE state to INTEGRATE state.
301:291	PCFSZE	This parameter determines the size, in bytes, of a PCF both on input as on output, including: 6-byte destination MAC address, 6-byte source MAC address, 2-byte EtherType field and 4-byte frame checksum. This parameter is used for to assemble output PCFs by appending as many trailing zeroes as needed to achieve the configured frame size. On input, the length of a PCF is expected to match the value of this parameter. If not, the frame is dropped. The clock synchronization block executes this check. To have a length limitation active for PCFs that pass through the switch, the MAXLEN field of the respective VL Policing table entries must be used. The minimum required value for this field is 64, the maximum allowed value is 2043.

Table 20. Clock Synchronization Parameters table (block 0Fh) *…continued*

Bit	Symbol	Description
290:287	PCFPRIORITY	This parameter determines the priority of a PCF sourced by a switch configured as compression master. It does not affect PCFs that are routed by the switch (even if the switch is configured as a compression master). The value of this parameter must be larger than the PRIORITY field of entries of the VL Forwarding Table or the VLANPMAP values of the L2 Forwarding configuration block (or their respective dynamically changed values) for a PCF to win the priority selection.
286:272	OBVWINSZ	This parameter defines the value of the observation window used by the compression master to collect and compress input protocol control frames. It is specified in multiples of 8 ns.
271:265	NUMUNSTBCY	This value defines the number of cycles of low clique support (i.e. the number of in-schedule clock masters that remain below the value of the STTH parameter) that can pass before the synchronization engine switches to INTEGRATE state. If it is zero, the synchronization engine will transit to INTEGRATE state at the end of the first acceptance window in STABLE state (if support is not provided for the clique).
264:258	NUMSTBCY	This value defines the number of cycles of clique support (i.e. the number of in-schedule clock masters meets the value of the SYTH parameter) needed to cause the synchronization engine to transit from SYNC state to STABLE state. If it is zero, the synchronization engine will transit to STABLE state at the end of the first acceptance window in SYNC state (if there is support for the clique). The parameter is only used when the SYTOSTBEN flag is set.
257:231	MAXTRANSPCLK	Determines the age at which a frame is processed by the compression master or the sync engine. It is provided in multiples of 8 ns and must account for the maximum PCF latency from synchronization masters or compression masters to this switch.
230:223	MAXINTEGCY	Determines the maximum value of the integration cycle that the synchronization engine will accept in input PCFs. If the IPCFRAMESY flag is set, this field determines the maximum value of output PCFs generated by the switch.
222:193	LISTENTMOUT	This parameter defines the maximum number of clock cycles that the synchronization engine remains in INTEGRATE state before switching to UNSYNC state when the SWMASTER flag is set. It is specified in multiples of 8 ns.
192:166	INTCYDUR	This field defines the duration of an integration cycle in multiples of 8 ns. The actual duration of the cycle will be by one greater than the value given in this field.
165:162	INTTOTENTTH	This value defines the minimum number of clock masters supporting an integration frame needed to perform a transition from INTEGRATE state to WAIT 4 CYCLE START state.
161:158	INTTOSYNCTH	This value defines the minimum number of clock masters needed to convince the synchronization engine to switch from INTEGRATE state to SYNC state.
157:142	VLIDOUT	This field defines the VL IDs to be used for protocol control frames generated by the switch when the SWMASTER flag is set. This value is always used for the VL ID field of Coldstart frames. If the VLIDSELECT flag is cleared, the ID of coldstart acknowledgement frames will be one greater than this value; if the VLIDSELECT flag is set, this value will be the VL ID. If the VLIDSELECT flag is cleared, integration frames will have an ID that is two greater than this value; if the VLIDSELECT flag is set, this value will be the VL ID. The value of the field must be less than 65534 if the VLIDSELECT flag is cleared.
141:126	VLIDIMNMIN	This field defines the minimum VL ID that used for input protocol control frames. All critical traffic input frames having a VL ID greater than or equal to this value and less than or equal to the value of VLIDINMAX are processed by the clock synchronization block of the IP. The value of the field must be smaller than the value of the VLIDINMAX field.

Table 20. Clock Synchronization Parameters table (block 0Fh) *…continued*

Bit	Symbol	Description
125:110	VLIDINMAX	This field defines the maximum VL ID that is used for input protocol control frames. All critical traffic input frames having a VL ID less than or equal to this value and greater than or equal to the value of VLIDIMNMIN are processed by the clock synchronization block of the IP. The value of the field must be larger than the value of the VLIDIMNMIN field. If the SWMASTER flag is set, the value of this parameter must be less than VLIDIMNMIN plus 24. Otherwise it must be less than VLIDIMNMIN plus three.
109:90	CAENTMOUT	This parameter defines how long the synchronization engine remains in CA_ENABLED state following the reception of a coldstart or coldstart acknowledge protocol control frame in UNSYNC state. It is specified in multiples of 8 ns.
89:75	ACCDEVWIN	This parameter is provided in multiples of 8 ns and determines the maximum deviation of an integration protocol control frame from the expected arrival time to be considered in schedule.
74	VLIDSELECT	This parameter determines if all protocol control frames use the same VL ID (when set) or dedicated VL IDs for each type (when cleared).
73	TENTSYRELEN	Determines if the switch relays time-triggered traffic while the clock synchronization engine is in TENTATIVE_SYNC state.
72	ASYTENSYEN	If this flag is set, the asynchronous clique detection mechanisms are enabled in TENTATIVE_SYNC state.
71	SYTOSTBEN	This parameter enables the transition of the clock synchronization algorithm from SYNC state to STABLE state if enough rounds (as specified by NUMUNSTBCY) have been passed in SYNC state.
70	SYRELEN	Determines if the switch relays time-triggered traffic at times the clock synchronization engine resides in SYNC state.
69	SYSYEN	If this flag is set, the synchronous clique detection mechanisms are enabled in SYNC state.
68	SYASYEN	If this flag is set, the asynchronous clique detection mechanisms are enabled in SYNC state.
67	IPCFRAMESY	When this flag is set, the switch periodically transmits integration protocol control frames with a period defined by the INTCYDUR field. When set, it requires flags VLIDSELECT, SYTOSTBEN, ACCDEVWIN, VLIDOUT, MAXINTEGCY, SYDOMAIN, SYPRIORITY and ETSSRCPCF and either OUTPRTSLAVE or OUTPRTMASTER to be configured. The remaining Clock Synchronization Parameters are not used if this flag is set.
66	STABASYEN	If this flag is set, the asynchronous clique detection mechanisms are enabled in STABLE state.
65	SWMASTER	If this flag is set, the switch acts as a compression master for the clock synchronization algorithm.
64	FULLCBG	If this flag is set, a switch with SWMASTER set exhibits guardian functionality on input protocol control frames. It is intended to be cleared in high-integrity environments (i.e. where sync master end systems are set up as high-integrity components).
63:61	SRCPORT(7)	These fields establish a mapping of protocol control frame membership flags to switch source ports. They are used to prevent source port spoofing. In case of multi-hop networks, a single port may host several membership flags (at the compression master). If specific membership flags are not used, their respective entries would be set to values greater than 4. If the SWMASTER flag is cleared, SRCPORT(0) must be set to the index of the port connecting to the compression master and all other ports to a value greater than 4.
:	:	
:	:	
:	:	
42:40	SRCPORT(0)	

Table 20. Clock Synchronization Parameters table (block 0Fh) *…continued*

Bit	Symbol	Description
39:35	OUTPRTSLAVE	This vector contains a flag for each port defining if integration protocol control frames sourced at the switch are output to the respective Ethernet port. It has an effect only when the SWMASTER flag or the IPCFRAMESY flag is set. Any protocol control frames not sourced at the switch are routed through the switch like non-protocol control frames. Policing and forwarding tables must be set up for the respective VLs.
34:30	OUTPRTMASTER	This vector contains a flag for each port defining if protocol control frames sourced at the switch are output to the respective Ethernet port. In contrast to a port having its flag set in OUTPRTSLAVE, ports having their flag set in this field output coldstart protocol control frames as well as coldstart acknowledge protocol control frames, in addition to integration protocol control frames. It has an effect only when the SWMASTER flag or the IPCFRAMESY flag is set. Any protocol control frame not sourced at the switch will be routed through the switch in the same way as non-protocol control frames. Policing and forwarding tables must be set up for the respective VLs.
29:0	not used	

4.2.17 AVB Parameters

Table 21 shows the layout of the AVB Parameters block. The AVB Parameters table specifies the source and destination MAC addresses of the meta frame created by the switch for every frame trapped by filtering rules for which an ingress timestamp is captured. The source port information and the switch ID can be used to distinguish the meta frames from different switches when multiple switches are used in a cascaded architecture. If INCL_SRCPT is set, the switch embeds the device ID and the source port in bytes one and two of the destination MAC address of the original frame. The payload format of the meta frame is depicted in Figure 9. The meta frame is sent immediately after the trapped frame that triggered the action.

Table 21. AVB Parameters table (block 10h)

Bit	Symbol	Description
95:48	DESTMETA	This field defines the destination MAC address used for metadata follow-up frames (see Section 4.2.18)
47:0	SRCMETA	This field defines the source MAC address used for metadata follow-up frames (see Section 4.2.18)

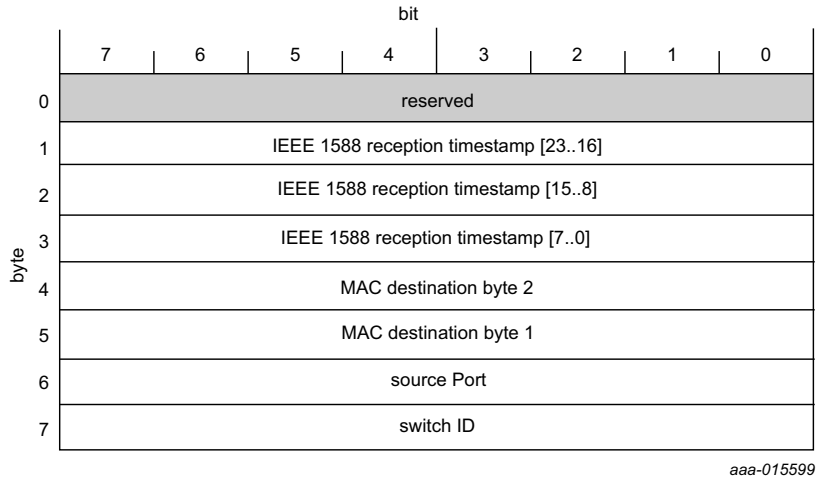

Fig 9. Structure of the meta-data follow-up frame

4.2.18 General Parameters

Table 22 shows the layout of the General Parameters table. This table contains general parameters used to configure basic properties of the switch.

Table 22. General Parameters table (block 11h)

Bit	Symbol	Description
319	VLLUPFORMAT	This field specifies the addressing scheme to be used to identify time-triggered and rate-constrained traffic. If this parameter is set to 1, time-triggered and rate-constrained traffic is identified by (1) the upper 32 bits of the frame destination MAC address filtered by VIMASK is equal to VIMARKER and (2) the lower 16 bits are configured together with the source port in the VL lookup table (see Table 6). If the parameter is set to 0, the frame is considered as rate-constrained or time-triggered if the combination of destination MAC address, VLAN ID, VLAN priority and source port is configured in the VL lookup table (see Table 5)
318	MIRR_PTACU	If this flag is set, the host can dynamically change the value of MIRR_PORT. If the flag is not set, changes of MIRR_PORT are prohibited.
317:315	SWITCHID	This field contains the configured switch ID used to identify the source of trapped frames forwarded to the host CPU in case the switch is composed from multiple cascaded devices.
314:312	HOSTPRIO	This field contains the priority value identifying the priority queue on HOST_PORT when a trapped frame is forwarded to the port connected to the host processor.
311:264	MAC_FLTRES[1]	This field contains a bit mask identifying a bridge level or MAC level management frame which shall be forwarded only to HOST_PORT. A received L2 frame produces a match if DEST_MAC and MAC_FLT[i] = MAC_FLTRES[i] holds. For example, to forward all groupcast traffic for the reserved OUI 01-80-C2-... to the host port, this field must be set to 01-80-C2-00-00-00.
263:216	MAC_FLTRES[0]	
215:168	MAC_FLT[1]	This field contains a bit mask identifying a bridge level or MAC level management frame which shall be forwarded only to HOST_PORT. A received L2 frame produces a match if DEST_MAC and MAC_FLT[i] = MAC_FLTRES[i] holds. For example, to forward all groupcast traffic for the reserved OUI 01-80-C2-... to the host port, this field must be set to FF-FF-FF-00-00-00. If the INCL_SRCPT[i] flag is set, MAC_FLT[i] must have set bytes 1 and 2 to zero (i.e.must be set to xx-xx-xx-00-00-xx, where an ‘x’ denotes an arbitrary user-defined value).
167:120	MAC_FLT[0]	
119	INCL_SRCPT[1]	If this field is set, the switch embeds for any frame where the destination MAC address matches the filter MAC_FLT[i] / MAC_FLTRES[i] the source port ID in byte 2 and the device ID in byte 1 of the MAC address, where byte 0 is the least significant byte. If the flag is set, MAC_FLT[i] must have set bytes 1 and 2 to zero (i.e. must be set to xx-xx-xx-00-00-xx, where an ‘x’ denotes an arbitrary user-defined value).
118	INCL_SRCPT[0]	
117	SEND_META[1]	If this field is set, the switch generates a meta frame containing the timestamp, source port ID and configured device ID for any frame where the destination MAC address matches MAC_FLT[i] / MAC_FLTRES[i]. The meta frame gets sent immediately after the filtered frame which produced the match.
116	SEND_META[0]	
115:113	CASC_PORT	If this field contains a valid port number, MAC bridge filtered MAC group traffic and standard MAC group traffic received on this port is automatically forwarded to HOST_PORT without including the source port and device ID information in the destination MAC address.
112:110	HOST_PORT	If this field contains a valid port number, MAC bridge filtered MAC group traffic and standard MAC group traffic is forwarded to this port. Otherwise, this traffic is dropped.

Table 22. General Parameters table (block 11h) *…continued*

Bit	Symbol	Description
109:107	MIRR_PORT	Traffic to be mirrored flows to this port if this field contains a valid port number. Traffic to be mirrored is identified by the EGR_MIRR and ING_MIRR flags of the entry in the MAC Configuration table (Table 14) for this port and by the VEGR_MIRR and VING_MIRR vectors of a VLANs entry in the VLAN Lookup table (Table 12). Note that mirroring will be faithful only if the switch operates lossless. The decision to mirror a frame is taken when the frame is dispatched to the destination port. The frame will still be mirrored if the destination port does not accept it because the respective priority queue is full or it exceeds the maximum permitted age. On the other hand, the mirror port may refuse to mirror frames because the respective priority queues are full or because the frames exceed the maximum configured age. A frame is dispatched to a port at most once; so if the port configured as mirror port by this field is also the intended forwarding destination for the frame, it is only dispatched once. Also, if several destination ports need to mirror a frame and, potentially, the source port, the frame will still be dispatched only once to the mirror port. In the case of time-triggered frames, the frame will be dispatched to the mirror port with its regular trigger in case the mirror port is on the route of the frame anyway (in case of loss of sync the frame may not get mirrored). If the mirror port is not on the route of a time-triggered frame, the frame will be dispatched to the mirror port with the first trigger that has the SETVALID flag asserted if ingress mirroring is enabled. If the mirror port is not on the route of a time-triggered frame and ingress mirroring is not selected for the frame, the frame will be dispatched to the mirror port with the first trigger that dispatches the frame to any port for which egress mirroring is selected.
106:75	VIMARKER	Provides the upper 32 bits of the Ethernet destination MAC address of frames considered to be critical traffic. More significant bytes of the MAC address are placed at lower bit positions in VIMARKER). The value of this parameter will be used for PCFs generated by the switch if it is configured as a compression master (without previously applying VIMASK).
74:43	VIMASK	Provides a bit-mask used for telling critical traffic from non-critical traffic. For the check, a bit-wise AND operation is performed both on VIMARKER and on the upper 32 bits of the Ethernet destination MAC address of the frame in question before the comparison. The parameter has the same byte order as VIMARKER so lower byte positions of VIMASK are used to mask the more significant bytes of the MAC address. When applying VIMASK to VIMARKER, bits at identical positions are logically ANDed.
42:27	TPID	This field contains the Ethernet Type Identifier used to identify tagged VLAN traffic.
26	IGNORE2STF	When set, this flag specifies that the ‘twoStepFlag’ of the ‘flagField’ of an IEEE 1588v2 event message shall be ignored. In this case, the ‘correctionField’ of an IEEE 1588v2 event message will always get updated with the residence time, even if the ‘twoStepFlag’ is asserted.
25:10	TPID2	This field contains the Ethernet Type Identifier used to identify double-tagged VLAN traffic.
9:0	not used	

4.2.19 Retagging table

Table 23 shows the layout of an entry in the Retagging table. This table is used to create copies of tagged Ethernet frames that will receive new VLAN IDs and will then be forwarded as if they were received on the original source port. The table has 32 entries and is optional.

One of the applications of this table is to replicate and retag frames for debugging and monitoring purposes. Consider a deeply embedded network composed of multiple switch devices and different VLANs configurations. To access the traffic from a debugging port, all frames tagged for debug monitoring must be forwarded to this port. In order to do so, the device is re-reconfigured during runtime to create a copy of every frame selected for monitoring/mirroring and route it to a VLAN that is configured to carry monitoring/debugging traffic. The forwarding rules for this VLAN then make sure that the retagged traffic finds its way through the network to be accessed by the monitoring device.

Table 23. Retagging table (block 12h)

Bit	Symbol	Description
63:59	EGR_PORT	If a frame with VLAN ID VLAN_ING is forwarded to any port having its flag set in this field, a copy with VLAN ID VLAN_EGR is generated. Only one copy is generated, even if the frame is forwarded to multiple ports having their respective flags set in this field or if the frame is received on any of the ports having their respective flags set in ING_PORT. The LSB of this field is assigned to port 0.
58:54	ING_PORT	If a frame with VLAN ID VLAN_ING is received on any port having its flag set in this field, a copy with VLAN ID VLAN_EGR is generated. Only one copy is generated, even if the frame is forwarded to any of the ports having their respective flags set in EGR_PORT as well. The LSB of this field is assigned to port 0.
53:42	VLAN_ING	The VLAN ID of the ingress frame.
41:30	VLAN_EGR	This VLAN ID replaces the VLAN ID of the original frame at egress. The priority code point of the VLAN tag is not changed.
29	DO_NOT_LEARN	If this flag is asserted, address learning is disabled for all frames carrying a VLAN ID that matches VLAN_EGR of the entry. This applies to frames being received on any of the Ethernet ports. Source addresses of frames generated by the retagging function are never be learned.
28	USE_DEST_PORTS	If this flag is asserted, DESTPORTS of the entry is used to route all frames carrying a VLAN ID that matches VLAN_EGR of the entry. This applies to both frames generated by the retagging function and to frames being received on any of the Ethernet ports. The configured route will bypass all other forwarding decisions. A frame may still be dropped at the egress port (if the respective transmit priority queue is filled to capacity) or by rate limitation and memory partition constraints. If several entries produce a match, the DESTPORTS field of the one with the smallest index is used.
27:23	DESTPORTS	This field provides a dedicated route for all frames carrying a VLAN ID that matches VLAN_EGR of the entry. This applies to both frames generated by the retagging function and to frames being received on any of the Ethernet ports. The configured route will bypass all other forwarding decisions. A frame may still be dropped at the egress port (if the respective transmit priority queue is filled to capacity) or by rate limitation and memory partition constraints.
22:0	not used	

4.2.20 xMII Mode Parameters

Table 24 shows the layout of the xMII Mode Parameters block. This block is used to set the xMII mode of operation. When PHY mode is selected, the port on this switch interface behaves as a PHY and the partner should behave as a MAC. When MAC mode is selected, the switch interface port behaves as a MAC and the partner should be a PHY. In order to set up the clocking scheme, the CGU must also be configured (see Section 5.3).

Table 24. xMII Mode Parameters table (block 4Eh)

Bit	Symbol	Description
31	PHY_MAC[4]	This parameter is used select the interface mode for port 4: 1 = PHY mode; 0 = MAC mode
30:29	xMII_MODE[4]	This parameter is used set the xMII mode for port 4: 00 = MII; 01 = RMII; 10 = RGMII, 11 = not used
28	PHY_MAC[3]	This parameter is used select the interface mode for port 3: 1 = PHY mode; 0 = MAC mode
27:26	xMII_MODE[3]	This parameter is used set the xMII mode for port 3: 00 = MII; 01 = RMII; 10 = RGMII, 11 = not used
25	PHY_MAC[2]	This parameter is used select the interface mode for port 2: 1 = PHY mode; 0 = MAC mode
24:23	xMII_MODE[2]	This parameter is used set the xMII mode for port 2: 00 = MII; 01 = RMII; 10 = RGMII, 11 = not used
22	PHY_MAC[1]	This parameter is used select the interface mode for port 1: 1 = PHY mode; 0 = MAC mode
21:20	xMII_MODE[1]	This parameter is used set the xMII mode for port 1: 00 = MII; 01 = RMII; 10 = RGMII, 11 = not used
19	PHY_MAC[0]	This parameter is used select the interface mode for port 0: 1 = PHY mode; 0 = MAC mode
18:17	xMII_MODE[0]	This parameter is used set the xMII mode for port 0: 00 = MII; 01 = RMII; 10 = RGMII; 11 = not used
16:0	not used	

5. Programming interface

This section describes the host computer interface to the SJA1105TEL. The host communicates with the device via the SPI interface (see Section 3). Register addresses are relative to a base address 00000000h.

5.1 Status area

The status area is divided into five sections: general status information, clock synchronization, memory partitioning, Ethernet port status information and virtual link status information. Reserved bits return 0 when read.

5.1.1 General status information

5.1.1.1 Device ID

Configuration data loaded via the configuration interface must start with the device ID.

Table 25. Device ID register (address 00h)

Bit	Symbol	Access	Value	Description
31:0	ID	R	9E00030Eh	device identification code

5.1.1.2 Configuration status information

The host should check the CONFIGS flag after loading the configuration file as described in Section 4.2. If CONFIGS = 0, the configuration load process should be reset and the configuration process restarted (with a valid configuration file). If CONFIGS = 1, the configuration is locked and configuration reset is no longer possible.

If CRCCHKL is set, a local CRC check failure occurred while the configuration file was being loaded. The host should check this flag after each block is loaded and reset the configuration process if it is set. If CRCCHKG is set, a global CRC check failure occurred during the configuration load process. IDS is set if the configuration file loaded did not contain a matching identifier.

NSLOT is a free-running 0 to 9 counter used for internal processing. It is intended to be used for debugging.

Table 26. Initial device configuration flag register (address 01h)

Bit	Symbol	Access	Value	Description
31	CONFIGS	R	[1]	device configuration status:
			0	configuration is invalid
			1	configuration is valid
30	CRCCHKL	R	[2]	local CRC check:
			0	local CRC check OK
			1	local CRC check failed
29	IDS	R	[2]	device identifier flag:
			0	matching device identifier found
			1	matching device identifier not found

Table 26. Initial device configuration flag register (address 01h) *...continued*

Bit	Symbol	Access	Value	Description
28	CRCCHKG	R	[2]	global CRC check:
			0	global CRC check OK
			1	global CRC check failed
27:4	reserved	R	all 0s	
3:0	NSLOT	R	xxxx	0 to 9 counter; overflows after 9

[1] Flag cleared at power-on/reset.

[2] Flag cleared at power-on/reset and when host resets the configuration load process.

5.1.1.3 Configured VL route and partition status information

VLROUTES is set when a critical frame is dropped because the configured route does not contain a port index. This information could be used, for example, to drop PCF traffic processed by a switch configured as a compression master for which the user wants to avoid receiving VLNOTFOUND errors on the respective input ports.

VLPARTS is set when a critical frame is dropped because the associated memory partition did not have enough space to store it. When this flag is set, the corresponding flag in the memory partition section is also set (see Section 5.1.3).

Table 27. Route/partition status register (address 02h)

Bit	Symbol	Access	Value	Description
31:16	VLIND	R	xxxxh	index in VL Forwarding Table that caused the VLROUTES or VLPARTS flag to be set
15:8	VLPARIND	R	xxxxh	index of partition that causes the VLPARTS flag to be set
7:2	reserved	R	all 0s	
1	VLROUTES	R	[1]	configured route status:
			0	route ok; no dropped frames
			1	critical traffic frame dropped because route configuration does not include a valid port
0	VLPARTS	R	[1]	partition overflow error status:
			0	memory partition ok; no dropped frames
			1	critical traffic frame dropped due to memory partition overflow

[1] Flag cleared at power-on/reset and on a read access by the host.

5.1.1.4 General status registers

L2BUSYFDS (see Table 28) is set if a frame received at PORTENF was dropped because it was received while the L2 Address Lookup table was being initialized. If the host respects the rule not to load the L2 Address Lookup table before the L2BUSYS flag is cleared, this condition will only occur in a setup that does not load the L2 Address Lookup table. To avoid setting this flag, wait for the L2BUSYS flag to be cleared before writing the last configuration word.

L2BUSYS = 1 indicates that the L2 Address Lookup table is being initialized. This flag will be set after a configuration reset condition: the device has not yet received a valid configuration but has received the first word of the configuration stream. The flag is cleared once initialization is complete and will remain cleared until a new power cycle or reset occurs. The L2 Address Lookup table cannot be loaded before this flag has been cleared.

Table 28. General status register 1 (address 03h)

Bit	Symbol	Access	Value	Description
31:16	MACADDL	R	xxxxh	lower (15 to 0) 16 bits of the source MAC address that triggered ENFFDS
15:8	PORTENF	R	xxh	number to the port that triggered ENFFDS
7:5	reserved	R	000	
4	FWDS	R	[1]	forwarding frame drop status:
			0	no dropped frames
			1	frame dropped because input port was not set to ‘forwarding’ when frame received
3	MACFDS	R	[1]	standard group MAC address frame drop status:
			0	no dropped frames
			1	frame dropped because it contained a filtered MAC address when the switch is not configured to forward such traffic
2	ENFFDS	R	[1]	enforced frame drop status:
			0	no dropped frames
			1	frame received at PORTENF dropped because it carried a source MAC address declared to be enforced on a different port in the L2 Address Lookup table
1	L2BUSYFDS	R	[2]	L2 Address Lookup table busy/frame drop status:
			0	not busy; no dropped frames
			1	frame received at PORTENF dropped because it was received while L2 Address Lookup table was being initialized
0	L2BUSYS	R		L2 Address Lookup table status:
			0	L2 Address Lookup table initialization complete
			1	L2 Address Lookup table is being initialized

[1] Flag cleared at power-on/reset and on a read access by the host.

[2] Flag cleared at power-on/reset.

Table 29. General status register 2 (address 04h)

Bit	Symbol	Access	Value	Description
31:0	MACADDU	R	xxxxxxxxh	upper (47 to 16) 32 bits of the source MAC address that triggered ENFFDS

Table 30. General status register 3 (address 05h)

Bit	Symbol	Access	Value	Description
31:16	MACADDHCL	R	xxxxh	lower (15 to 0) 16 bits of the source MAC address that triggered HASHCONFS
15:4	VLANIDHC	R	xxxh	VLAN ID that triggered HASHCONFS
3:1	reserved	R	000	
0	HASHCONFS	R		hash conflict status:
			0	no hash conflict encountered
			1	unresolved hash conflict in L2 Address Lookup table; means that all available ways for dynamic address learning are already occupied for the hash calculated for the pair MACADDHCL/MACADDHCU and VLANIDHC

Table 31. General status register 4 (address 06h)

Bit	Symbol	Access	Value	Description
31:0	MACADDHCU	R	xxxxxxxxh	upper (47 to 16) 32 bits of the source MAC address that triggered HASHCONFS

WPVLANID (see Table 32) contains the VLAN ID (either the VLAN ID contained in a tagged frame or the default VLAN ID of the port as specified by the VLANID field of the MAC Configuration table (Table 14) in the case of an untagged frame) that triggered WRONGPORTS or VNOTFOUNDS. The field only contains valid data when one of these flags is set.

VLANBUSYS = 1 indicates that the VLAN Lookup table is being initialized. This flag will be set after a configuration reset condition: the device has not yet received a valid configuration but has received the first word of the configuration stream. The flag is cleared once initialization is complete and will remain cleared until a new power cycle or reset occurs. The VLAN Lookup table cannot be loaded before this flag has been cleared.

Table 32. General status register 5 (address 07h)

Bit	Symbol	Access	Value	Description
31:16	WPVLANID	R	xxxxh	VLAN ID that triggered WRONGPORTS or VNOTFOUNDS
15:8	PORT	R	xxh	input port number that triggered WRONGPORTS or VNOTFOUNDS; contains valid data when one of these flags is set.
7:5	reserved	R	000	
4	VLANBUSYS	R		VLAN Lookup table status:
			0	VLAN Lookup table initialization complete
			1	VLAN Lookup table is being initialized

Table 32. General status register 5 (address 07h) *…continued*

Bit	Symbol	Access	Value	Description
3	WRONGPORTS	R	[1]	port status for VLAN frame:
			0	frame/port status OK
			1	frame received at PORT dropped because the port is not configured for the VLAN ID in the VLAN Lookup table
2	VNOTFOUNDS	R	[1]	port/VLAN ID status:
			0	frame OK
			1	frame received at port PORT containing the VLAN ID dropped because the VLAN ID is not configured in the VLAN Lookup table
1:0	reserved	R	00	

[1] Flag cleared at power-on/reset and on a read access by the host.

Table 33. General status register 6 (address 08h)

Bit	Symbol	Access	Value	Description
31:16	VLID	R	xxxxh	VLID that did not produce a hit in the VL Lookup table; contains valid data when VLNOTFOUND is set
15:8	PORTVL	R	xxh	index of port at which VLID (that did not produce a hit in the VL Lookup table) was received
7:1	reserved	R	00h	
0	VLNOTFOUND	R	[1]	port/VLAN ID status:
			0	frame OK
			1	PORTVL and VLID contain a tuple that did not generate a hit in the VL Lookup table

[1] Flag cleared at power-on/reset and on a read access by the host.

EMPTYS = 1 (see Table 34) indicates that dynamic memory management experienced an ‘out of memory’ condition between the most recent read access to this field and the current one. This means that the configuration of memory partition sizes does not comply with the configuration rules. The switch will not behave as expected when this error occurs.

Table 34. General status register 7 (address 09h)

Bit	Symbol	Access	Value	Description
31	EMPTYS	R	[1]	dynamic memory status:
			0	memory ok
			1	‘out of memory’ condition registered in dynamic memory management
30:0	BUFFERS	R	xxxxxxxxh	number of frame buffers available to dynamic memory management; it will be set to 1024 after power-on or reset, but the receive ports will immediately start to draw buffers (1 buffer per port); used for testing and debugging

[1] Flag cleared at power-on/reset and on a read access by the host.

Table 35. General status register 8 (address 0Ah)

Bit	Symbol	Access	Value	Description
31:16	reserved	R	all 0s	
15:8	PORT	R	xxh	port number where frame dropped if FWDS or PARTS is set
7:2	reserved	R	all 0s	
1	FWDS	R	[1]	port forwarding status:
			0	no dropped frames
			1	port identified by PORT sourced a frame that was dropped because the configured forwarding direction did not contain any ports
0	PARTS	R	[1]	port/VLAN ID status:
			0	no dropped frames
			1	frame received at PORT dropped because the respective memory partition did not have enough space to hold the frame

[1] Flag cleared at power-on/reset and on a read access by the host.

Each memory block is assigned a dedicated RAMPARERR flag (see Table 36 and Table 37). The associated flag is set when a parity error is detected in a memory block. These flags are provided for test and debug purposes. If any of these flags is set during normal operation, the host must reset the switch. The switch stops forwarding frames when a parity error is detected.

Table 36. General status register 9 (address 0Bh)[1]

Bit	Symbol	Access	Description
31:21	reserved	R	
20:0	RAMPARERRL[20:0]	R	If one of these flags is found set, a parity error has been detected in a memory block. Each memory block is assigned a dedicated flag in this vector. If any of these flags are found set during operation, the host must reset the switch.

[1] All flags cleared at power-on/reset.

Table 37. General status register 10 (address 0Ch)[1]

Bit	Symbol	Access	Description
31:5	reserved	R	
4:0	RAMPARERRU[25:21]	R	If one of these flags is found set, a parity error has been detected in a memory block. Each memory block is assigned a dedicated flag in this vector. If any of these flags are found set during operation, the host must reset the switch.

[1] All flags cleared at power-on/reset.

PTPEGR_TSn contains the PTP egress timestamp of the most recent management frame for which the user specified an egress timestamp to be captured.

Table 38. General status registers 11 to 20 (addresses C0h to C9h)

Bit	Symbol	Access	Value	Description
31:8	PTPEGR_TSn[1]	R	xxxxxxh[2]	PTP egress timestamp
7:1	reserved	R	00h	
0	UPDATEn[1]	R	x[2]	set if value PTPEGR_TSn has changed

[1] 'n' is an index from 0 (address 192) to 9 (address 201) and is calculated using the formula $n = 2 \times PORT + TSREG$, where PORT is the port on which the timestamp was taken and TSREG specifies which of the two timestamp registers of the given port is addressed. The dynamic reconfiguration entry used for routing the management frame (MGMTROUTE; see Table 64) determines which register contains the timestamp.

[2] Flag cleared at power-on/reset and on a read access by the host.""""

5.1.2 Clock synchronization

Table 39. Synchronization engine status register (address 40h)

Bit	Symbol	Access	Value	Description
31:8	reserved	R	000h	
7:0	SYNCSTATE	R		state of synchronization engine
			0	INIT
			1	INTEGRATE
			2	UNSYNC
			3	WAIT_4_CYCLE_START
			4	CA_ENABLED
			5	WAIT_FOR_IN
			6	TENTATIVE_SYNC
			7	SYNC
			8	STABLE

Table 40. Synchronization status registers (address range 41h to 57h)

Address	Bits	Symbol	Access	Description
41h	31:0	INTEGCY	R	This field contains the *local_integration_cycle* state variable of the clock synchronization algorithm
42h	31:0	ACTCORR	R	This field displays the correction value most recently applied by the clock synchronization algorithm as a 32-bit two's complement number. It displays zero if no term has ever been applied. A negative value indicates that the local clock is running too fast with respect to the cluster clock.
43h	31:0	MAXCORR	R	This field displays the largest correction value ever applied by the clock synchronization algorithm as 32 bits two's complement. It displays zero if no term has ever been applied. A negative value indicates that the local clock is running too fast with respect to the cluster clock.
44h	31:0	MINCORR	R	This field displays the smallest correction value ever applied by the clock synchronization algorithm as 32 bits two's complement. A negative value indicates that the local clock is running too fast with respect to the cluster clock.

Table 40. Synchronization status registers (address range 41h to 57h) *…continued*

Address	Bits	Symbol	Access	Description
45h	31:16	reserved	R	-
	15:0	SYNCLOSS	R	This field displays the number of transitions of the clock synchronization state machine from SYNC/STABLE state to a state other than SYNC/STABLE state as a 16-bit integer value. The counter stops counting when it reaches its maximum value.
46h	31:8	reserved	R	-
	7:0	LOCMEM		This field displays the *local_membership_comp* state variable of the clock synchronization algorithm.
47h	31:8	reserved	R	-
	7:0	MINMEMVAR		This field displays the membership vector with the smallest number of flags set of all in-schedule integration frames ever processed by the synchronization state machine. The field maintains its value when synchronization is lost; it is never reset. The field does not reflect PCF frames dropped by low-level filtering or because failing on the BAG check. Field will contain all 1s after power-on or reset.
48h	31:8	reserved	R	-
	7:0	MAXMEMVAR		This field displays the membership vector with the largest number of flags set of all in-schedule integration frames ever processed by the synchronization state machine. The field maintains its value when synchronization is lost; it is never reset. The field does not reflect PCF frames dropped by low-level filtering or because failing on the BAG check. Field will contain all 0s after power-on or reset.
49h	31:17	reserved	R	-
	23:16	LOCASYNMEM0	R	This field displays the *local_async_membership_comp* vector state variable of the clock synchronization algorithm with index zero.
	15:8	reserved	R	-
	7:0	LOCASYNMEM1	R	This field displays the *local_async_membership_comp* vector state variable of the clock synchronization algorithm with index one.
4Ah	31:8	reserved	R	-
	7:0	MINASYNMEM	R	This field displays the bit vector of the async membership count variable of the clock synchronization algorithm with the smallest number of flags set or all-ones if no evaluation has yet taken place. The field maintains its value even if synchronization is lost. It is never reset. After power-on or reset, it contains all-ones.
4Bh	31:8	reserved	R	-
	7:0	MAXASYNMEM	R	This field displays the bit vector of the async membership count variable of the clock synchronization algorithm with the largest number of flags set ever evaluated or all zeros if no evaluation has yet taken place. The field maintains its value even if synchronization was lost. It is never reset. Field will contain all 1s after power-on or reset.
4Ch	31:8	reserved	R	-
	7:0	CAMEM	R	This field displays the membership vector contained in the coldstart frame or coldstart acknowledge frame that caused the most recent transition from UNSYNC state to CA_ENABLED state, or all-zeroes if no such transition has taken place. Note that this field retains all 0s for a switch that is not configured as a compression master.

Table 40. Synchronization status registers (address range 41h to 57h) *…continued*

Address	Bits	Symbol	Access	Description
4Dh	31:8	reserved	R	-
	7:0	SYSNMEM	R	This field displays the membership vector contained in the integration frame that caused the most recent transition from UNSYNC state to SYNC state or TENTATIVE_SYNC state or all-zeroes if no such transition has yet taken place. Note that this field retains all 0s if the switch never took part in starting the network (which is always the case for switches that are not configured as compression masters).
4Eh:55h	31:0	MOFFSETn	R	This field (in each of the eight registers) contains a timestamp taken whenever a frame of the respective synchronization master becomes permanent. The source of the timestamp is the local clock of the synchronization engine. The address assigned to a specific synchronization master equals its bit position within the membership vector. These fields always return zero on read for switches whose SWMASTER flag is not set; 'n' is an index from 0 (address 78) to 7 (address 85)
56h	31:30	reserved	R	-
	29:0	TIMER	R	This field displays the local timer state variable of the clock synchronization algorithm.
57h	31:27	reserved	R	-
	26:0	CLOCK	R	This field displays the local clock state variable of the clock synchronization algorithm.

5.1.3 Memory partitioning

Table 41. VL memory partition status registers (address range 80h to 87h)

Bit	Symbol	Access	Value	Description
31	PARTDRPVLn[1]	R	[2]	VL memory partition status:
			0	no memory error
			1	If this flag is set, the respective VL memory partition suffered at least one out-of-memory error after the previous read access to this field.
30:0	PARTSPCVLn[1]	R	xxxxxxxxh	Each of these fields contains the number of frames left for the respective VL memory partition at the time of the read access. After configuration (and before receiving the first frame drawing from a particular partition), each field will be set to the value specified by the respective PartitionSpace parameter of the VL Forwarding Parameters configuration block. A critical traffic frame will only be accepted if there is space left within the respective VL memory partition.

[1] 'n' is an index from 0 (address 80h) to 7 (address 87h).

[2] Flag cleared at power-on/reset and on a read access by the host.

Table 42. VL memory partition error counters (address range 800h to 807h)

Bit	Symbol	Access	Description
31:0	PARTDRPCNTVLn[1]	R	Each of these fields contains the number of frames dropped due to lack of VL memory partition space since power-on or reset. The counter wraps.

[1] 'n' is an index from 0 (address 800h) to 7 (address 807h).

Table 43. L2 memory partition status registers (address range 100h to 107h)

Bit	Symbol	Access	Value	Description
31	L2PARTSn[1]	R	[2]	switch memory partition status for Ethernet traffic:
			0	no memory error
			1	If this flag is set, the respective memory partition suffered at least one out-of-memory error after the previous read access to this field.
30:0	N_L2PSPCn[1]	R	xxxxxxxxh	Each of these fields contains the number of frames left for the respective L2 memory partition at the time of the read access. After configuration (and before receiving the first frame drawing from a particular partition) each field will be set to the value specified by the respective PART_SPC parameter of the L2 Forwarding Parameters configuration block. A frame will only be accepted if there is space left within the respective memory partition.

[1] 'n' is an index from 0 (address 100h) to 7 (address 107h).

[2] Flag cleared at power-on/reset and on a read access by the host.

Table 44. L2 memory partition error counters (address range 1000h to 1007h)

Bit	Symbol	Access	Description
31:0	PARTDROPn[1]	R	Each of these fields contains the number of frames dropped due to lack of best-effort memory partition space since power-on or reset. The counter wraps.

[1] 'n' is an index from 0 (address 1000h) to 7 (address 1007h)

5.1.4 Ethernet port status

MAC-level diagnostics counters and flags are provided for each port as detailed in Table 48 and Table 49. Addresses in these tables are relative to the base address of the respective port, as listed in Table 45. High-level diagnostics counters are provided for each port as detailed in Table 50. Addresses in this table are relative to the base address of the respective port, as listed in Table 46.

Table 45. Ethernet MAC-level port status base addresses

Port number	Base address
Ethernet port 4	208h
Ethernet port 3	206h
Ethernet port 2	204h
Ethernet port 1	202h
Ethernet port 0	200h

Table 46. Ethernet high-level port status part 1 base addresses

Port number	Base address
Ethernet port 4	440h
Ethernet port 3	430h
Ethernet port 2	420h
Ethernet port 1	410h
Ethernet port 0	400h

Table 47. Ethernet high-level port status part 2 base addresses

Port number	Base address
Ethernet port 4	640h
Ethernet port 3	630h
Ethernet port 2	620h
Ethernet port 1	610h
Ethernet port 0	600h

5.1.4.1 MAC-level status

Table 48. Ethernet port status - MAC-level diagnostic counters (relative address 0h)

Bit	Symbol	Access	Description
31:24	N_RUNT[1]	R	This field counts the number of frames that do not have a SOF, alignment or MII error, but are shorter than 64 bytes. The counter does not wrap.
23:16	N_SOFERR[1]	R	This field counts the number of frames that started less than 16 clock cycles after the most recent frame that at least had a correct SOF pattern with a byte other than 55h or D5h, have a byte other than D5h being the first byte that is different from 55h (if the frame starts with a preamble), that have the MII error input being asserted prior to or up to the SOF delimiter byte or that terminated before the SOF delimiter byte or immediately after the SOF delimiter byte. The field does not wrap.
15:8	N_ALIGNERR[1]	R	This field counts the number of frames that started with a valid start sequence (preamble plus SOF delimiter byte) but whose length is not a multiple of 8 bits (at the given line speed). The field does not wrap.
7:0	N_MIIERR	R	This field counts the number of frames that started with a valid start sequence (preamble plus SOF delimiter byte) but terminated with the MII error input being asserted. The field does not wrap.

[1] Flag cleared at power-on/reset and on a read access by the host.

Table 49. Ethernet port status - diagnostic flags (relative address 1h)

Bit	Symbol	Access	Description
31:28	reserved	R	-
27	TYPEERR[1]	R	This flag is set if a protocol control frame was received on the respective port with a type value that is not defined.
26	SIZEERR[1]	R	This flag is set if a protocol control frame was received on the respective port with a size other than was specified by the PCFSZE parameter in the Clock Synchronization Parameters block of the configuration file (Table 20).
25	TCTIMEOUT[1]	R	This flag is set if a protocol control frame was received on the respective port that had a transparent clock value that does not allow processing of the frame at its permanence point (i.e. the value contained in the transparent clock field does not have enough margin with respect to the value of the MAXTRANSPCLK parameter of the Clock Synchronization Parameters block of the configuration file).
24	PRIORERR[1]	R	This flag is set if a protocol control frame was received on the respective port that had a priority value other than the one set in the SYPRIORITY parameter of the Clock Synchronization Parameters block of the configuration file.
23	NOMASTER[1]	R	This flag is set if a coldstart protocol control frame or a coldstart acknowledgement protocol control frame has been received on the respective port although the switch is not configured as a compression master.
22	MEMOV[1]	R	This flag is set if a protocol control frame was received on the respective port that had flags of the membership vector set that are not supported (only bit positions 0 to 7 are supported).
21	MEMERR[1]	R	This flag is set if the switch is configured as compression master (the SWMASTER flag of the Clock Synchronization Parameters configuration block is set) and a protocol control frame was received on the respective port whose membership vector does not match the VL ID of the protocol control frame.
20	reserved	R	-
19	INVTYP[1]	R	This flag is set if a protocol control frame (as defined by the VL ID) was received on the respective port that had an EtherType value other than h'891D.

Table 49. Ethernet port status - diagnostic flags (relative address 1h) *...continued*

Bit	Symbol	Access	Description
18	INTCYOV[1]	R	This flag is set if a protocol control frame was received on the respective port that had an integration cycle value larger than MAXINTEGCY of the Clock Synchronization Parameters configuration block.
17	DOMERR[1]	R	This flag is set if a protocol control frame was received on the respective port that had a domain value that does not match the value set for the SYDOMAIN parameter of the Clock Synchronization Parameters configuration block.
16	PCFBAGDROP[1]	R	This flag is set if a protocol control frame was received on the respective port that was to become permanently more than MAXTRANSPCLK past the preceding protocol control frame of the same master. This check is performed at compression master switches only.
15:12	SPCPRIOR	R	see flag SPCERR
11:8	AGEPRIOR	R	see flag AGEDRP
7	reserved	R	
6	PORTDROP[1]	R	This flag is set to indicate that a frame was dropped at the respective port because the port has not been enabled for traffic in the L2 Policing table (Table 11).
5	LENDROP[1]	R	This flag is set to indicate that a frame was dropped at the respective port because the frame was longer than defined in the L2 Policing table.
4	BAGDROP[1]	R	This flag is set to indicate that a frame was dropped at the respective port because there was no bandwidth left on the port as defined in the L2 Policing table.
3	POLIECERR[1]	R	This flag is set to indicate that a critical traffic frame failed on either the VL-specific length check or on the timeliness check (BAG check, TT window check, or unreleased check for input RC frames that are routed as TT) at the respective port. The flag is reset on power-on/reset and whenever the host reads the status word containing the flag.
2	DRPNONA664ERR[1]	R	This flag is set to indicate that a frame was dropped at the respective port because its EtherType field contained a value other than 800h while the DRPNONA664 flag of the MAC Configuration block is set (see Section 4.2.10).
1	SPCERR[1]	R	This flag is set to indicate that a frame was dropped at the respective port because the respective priority queue as defined by the BASE and TOP parameters (in the MAC configuration block) did not have any space left or is deactivated (as defined by the ENABLED array of flags within the MAC Configuration table). If the flag is set, SPCPRIOR will contain the index of the priority queue that hosted the dropped frame.
0	AGEDRP[1]	R	This flag is set to indicate that a critical traffic frame was dropped at the respective port because it hit the maximum age as specified by the MAXAGE parameter of the port (in the MAC Configuration table). If the flag is set, AGEPRIOR will contain the index of the priority queue that hosted the dropped frame.

[1] Flag cleared at power-on/reset and on a read access by the host.

5.1.4.2 High-level status

Table 50. Ethernet high-level port status diagnostic counters part 1

Relative address	Bit	Symbol	Access	Description
Fh	31:0	N_N664ERR	R	This field counts the number of frames dropped since power-on or reset because they had an EtherType field other than 800h while the DRPNONA664 flag was set for the respective port in the MAC Configuration table (Table 14), they were not tagged while untagged traffic was not allowed (DRPUNTAG = 1; see Table 14), or that were not routed to any destination (because destination ports were down because flag EGRESS = 0, destination ports were not reachable for traffic sourced at the respective ingress port as per REACH_PORT of the respective ingress port, or destination ports were not members of the VLAN broadcast domain as per VLAN_BC of the respective VLAN). The counter wraps.
Eh	31:0	N_VLANERR	R	This field counts the number of frames that were dropped since power-on or reset because the VLAN ID was either not found in the VLAN Lookup table, the respective port is not listed in the VMEMB_PORT vector of the configured VLANID, or a legal or illegal double-tagged frame was received while double-tagged traffic was not allowed (DRPDTAG = 1; see Table 14). The counter wraps.
Dh	31:0	N_UNRELEASED	R	This counter is incremented when a frame is received on the respective port that meets the conditions for increasing the Unreleased counter of any Virtual Link Status entry. This means that this counter is also incremented when the respective Virtual Link Status entry is not increased because it has already reached its maximum value (since the Unreleased counter of the respective Virtual Link Status entry does not wrap). The counter wraps.
Ch	31:0	N_SIZEERR	R	This field counts the number of frames received since power-on or reset with an invalid length (2 kB or more or the length contained in the Type/Length field of the frame did not match the actual length) as well as frames received while ingress traffic was disabled (INGRESS = 0; see Table 14) on this port. The counter wraps.
Bh	31:0	N_CRCERR	R	This field counts the number of frames that had a receive-side CRC error on this port since power-on or reset. The counter wraps.
Ah:9h	-	reserved	R	-
8h	31:0	N_POLERR	R	This field counts the number of frames that were dropped based on the L2 policing operation (rate limit exceeded, length limit exceeded, source address spoofing, the port is not configured, a frame received from the host port produced a match with MAC_FLT/MAC_FLTRES but the host did not provide routing information for the respective destination MAC address) since power-on or reset. The counter wraps.
7h	31:0	N_RXFRMSH	R	This field contains the upper bits of the most recently read N_TXBYTE, N_TXFRM, N_RXBYTE or N_RXFRM counter, no matter what the port was. The intended use is to read the N_RXFRM field and the RSFRAMESSH field of a specific port as an atomic action, i.e. without reading any of the N_TXBYTE, N_TXFRM, N_RXBYTE or N_RXFRM fields of the same port or of other ports in between. Only in this way can a consistent counter value be received. The counter wraps.
6h	31:0	N_RXFRM	R	This field contains the lower 32 bits of the number of MAC-level correct frames received on the respective port since power-on or reset. The counter wraps. When reading from this address, the upper bits of the counter are stored to a shadow register accessible on relative address seven.

Table 50. Ethernet high-level port status diagnostic counters part 1 *…continued*

Relative address	Bit	Symbol	Access	Description
5h	31:0	N_RXBYTESH	R	This field contains the upper bits of the most recently read N_TXBYTE, N_TXFRM, N_RXBYTE or N_RXFRM counter, no matter what the port was. The intended use is to read the N_RXBYTE field and the N_RXBYTESH field of a specific port as an atomic action, i.e. without reading any of the N_TXBYTE, N_TXFRM, N_RXBYTE or N_RXFRM fields of the same port or of other ports in between. This ensures that a consistent counter value is received. The counter wraps.
4h	31:0	N_RXBYTE	R	This field contains the lower 32 bits of the number of bytes (all data bytes of an Ethernet frame from the first byte of the Ethernet destination MAC address to the last byte of the checksum but not including preamble bytes nor SOF delimiters) received on the respective port in MAC-level correct frames since power-on or reset. The counter wraps. When reading from this address, the upper bits of the counter are stored to a shadow register accessible on relative address 5h.
3h	31:0	N_TXFRMSH	R	This field contains the upper bits of the most recently read N_TXBYTE, N_TXFRM, N_RXBYTE or N_RXFRM counter, no matter what the port was. The intended use is to read the N_TXFRM field and the N_TXFRMSH field of a specific port as an atomic action, i.e. without reading any of the N_TXBYTE, N_TXFRM, N_RXBYTE or N_RXFRM fields of the same port or of other ports in between. This ensures that a consistent counter value is received. The counter wraps.
2h	31:0	N_TXFRM	R	This field contains the lower 32 bits of the number of frames transmitted to the respective port since power-on or reset. The counter wraps. When reading from this address, the upper bits of the counter are stored to a shadow register accessible on relative address 3h.
1h	31:0	N_TXBYTESH	R	This field contains the upper bits of the most recently read N_TXBYTE, N_TXFRM, N_RXBYTE or N_RXFRM counter, no matter what the port was. The intended use is to read the N_TXBYTE field and the N_TXBYTESH field of a specific port as an atomic action, i.e. without reading any of the N_TXBYTE, N_TXFRM, N_RXBYTE or N_RXFRM fields of the same port or of other ports in between. This ensures that a consistent counter value is received. The counter wraps.
0h	31:0	N_TXBYTE	R	This field contains the lower 32 bits of the number of bytes (all data bytes of an Ethernet frame from the first byte of the Ethernet destination MAC address to the last byte of the checksum but not including preamble bytes nor SOF delimiters) transmitted to the respective port since power-on or reset. The counter wraps. When reading from this address, the upper bits of the counter are stored to a shadow register accessible on relative address 1h.

Table 51. Ethernet high-level port status diagnostic counters part 2

Relative address	Bit	Symbol	Access	Description
3h	31:0	N_QFULL	R	This field counts the number of frames that were dropped on egress because the respective priority queue of the destination port (as defined per VLAN_PMAP of the L2 Forwarding table) or of a critical traffic frame (as defined per PRIORITY of the VL Forwarding table received at this port) did not have any space left since power-on or reset. The counter wraps.
2h	31:0	N_PART_DROP	R	This field counts the number of frames that were dropped on ingress because the respective memory partition of the port (as defined per PARTITION of the L2 Policing table) or of a critical traffic frame (as defined per PARTITION of the VL Forwarding table received at this port) had no space left after power-on or reset. The counter wraps.
1h	31:0	N_ERG_DISABLED	R	This field counts the number of frames that were not routed to the port this counter is assigned to, since power on or reset, because the port was down (EGRESS = 0;see Table 14). The counter wraps.
0h	31:0	N_NOT_REACH	R	This field counts the number of frames that produced a match in the L2 Lookup table since power-on or reset, but were not routed to the port this counter is assigned to because the port is not reachable for the respective ingress port as per REACH_PORT in the L2 Forwarding table. The counter wraps.

5.1.5 Virtual link status

Virtual link status is provided by two double words per VL that contain all status information for the respective VL. The status words at addresses 65536 and 65537 are assigned to the input VL with the lowest ID (as contained in the VL Lookup table). The status words at addresses 65538 and 65539 are assigned to the second lowest ID, and so on. All the counters are reset at power-on and whenever the host submits a request to clear the status words by asserting the CLEARVLS flag in the Control area (see Table 57). All of the counters wrap. Reading counters at an offset that are not defined in the VL Lookup table of the configuration data returns arbitrary values.

Table 52. Virtual link status (address 10000h)

Bit	Symbol	Access	Description
31:16	TIMINGERRn	R	This counter is incremented when a VL configured to be RC-policed on input fails on the BAG check, if a VL configured to be TT-policed on input fails on the window check (either because the window is closed, synchronization is not established, or flushing the TT buffers after loss of synchronization has not been completed yet), or if a VL configured to be TT-dispatched is received at a time synchronization is not established or flushing the TT buffers after loss of synchronization has not been completed yet. The counter is incremented by 1 at most for each frame (e.g. when receiving a frame of a VL that is configured to be RC-policed and TT-dispatched at a time synchronization is not established and if this frame also violates the BAG, the counter is incremented by 1).
15:0	UNRELEASEDn	R	This counter is incremented whenever a frame is received for a VL that is dispatched in time-triggered fashion at a time the most recently received frame has not yet been dispatched to all destination ports. In a correct setup, this can only be the case for a VL that is policed in RC mode but dispatched in TT mode. For VLs being policed in TT fashion the receive window check is supposed to fail at times the VL has ports left in the dispatch schedule. The respective check is only performed for frames that passed all policing checks (BAG or window) as well as the memory partition check.

Table 53. Virtual link status (address 10001h)

Bit	Symbol	Access	Description
31:16	reserved	R	
15:0	LENGTHERRn	R	This counter is incremented whenever a frame of the VL is received that contains more bytes than specified by the MAXLENGTH parameter of the respective VL Policing table entry.

5.2 Control area

The controls area manages some of the switch functionality during run time. Write access to an address not listed in this section is ignored.

5.2.1 General control

Bits RPARINITL[20:0] and RPARINITU[25:21] are used to configure the RAM parity check. After power-on or reset, these bits will contain all 0s, resulting in even parity. These bits are intended for self-testing: the host could load a configuration, apply stimulus, change the parity bit of a specific block and check if the respective RAMPARERR flag is set after applying enough stimulus to trigger a parity check. When read, these fields will return the value most recently written by the host or all 0s after power-on or reset.

Table 54. RAM parity check configuration register 2 (address 0Dh)

Bit	Symbol	Access	Description
31:21	reserved	R	
20:0	RPARINITL[20:0]	R/W	This field is used to change the parity check. For each bit, 0 selects even parity and 1 selects odd parity. The parity is set to all zeros after reset.

Table 55. RAM parity check configuration register 2 (address 0Eh)

Bit	Symbol	Access	Description
31:5	reserved	R	
4:0	RPARINITU[25:21]	R/W	This field is used to change the parity check. For each bit, 0 selects even parity and 1 selects odd parity. The parity is set to all zeros after reset.

Table 56. Ethernet port status control register (address 0Fh)

Bit	Symbol	Access	Description
31:5	reserved	R	
4:0	CLEARPORT	W	This field is used to reset the MAC-level diagnostics counters and flags for each port. Setting a bit to 1 resets the relevant MAC-level port status information, as described in Section 5.1.4. High-level diagnostic counters belonging to a port cannot be reset. This field returns all 0s on read.

Table 57. VL status control register (address 10h)

Bit	Symbol	Access	Description
31:1	reserved	R	
0	CLEARVLS	R/W	Setting this bit to 1 triggers a reset of the status information for all VLs. It remains set during the reset process and is cleared once the operation is complete. Setting this bit again while a reset is in progress has no effect. The switch will not update any counters while the reset is in progress. All counters will show zero immediately after this flag is cleared.

Table 58. Ethernet port status control register (address 11h)

Bit	Symbol	Access	Description
31:5	reserved	R	
4:0	INHIBITTX	W	This vector represents the set of ports on which the transmission is inhibited. A port inhibits transmission if and only if the respective bit is set to 1. Transmission to the respective port resumes when the bit is cleared. Only frame output is stopped, while frame queue processing continues. In effect, resource utilization in the switch (e.g. frame memory occupied by frames being forwarded to an inhibited port) will not change in response to this flag. This also means that frames that would have been transmitted to a port at a time the port had its flag in this vector set will never be transmitted to this port. Changing the flag will have an effect once the port is in an IFG, a change cannot generate malformed packets. This field returns all 0s on read.

Table 59. PTP control register 1 (address 17h)

Bit	Symbol	Access	Description
31	VALID	W	Setting this flag in combination with PTPSTRTSCH or PTPSTOPSCH dynamically changes the behavior of the external PTP_CLK pin. This flag always returns 0 on read.
30	PTPSTRTSCH	R/W	Setting this flag in combination with VALID to true triggers the switch to integrate all active subschedules into the schedule at the first entry configured in the Schedule Entry Points table (Table 4). This action is performed at a time when the PTP clock synchronized by the host exceeds the value of PTPSCHTM and only if the value of CLKSRC in the Schedule Entry Points Parameter table (Table 16) is set to 11 (PTP clock) and the schedule is not already running. On read, this flag is found true if CLKSRC is set to PTP and the schedule is running correctly running.
29	PTPSTOPSCH	R/W	Setting this flag in combination with VALID to true triggers the switch to stop schedule execution immediately if CLKSRC (in Table 16) is set to 11 (PTP clock). This flag has precedence over PTPSTRTSCH. On read, this flag will be found set if CLKSRC is set to PTP and the schedule is not running.
28	STARTPTPCP	W	Setting this flag in combination with VALID triggers the switch to begin toggling the external PTP_CLK pin at a rate of PTPPINDUR when the PTP clock synchronized by the host exceeds the value of PTPPINST.
27	STOPPTPCP	R/W	Setting this flag in combination with VALID triggers the switch to stop toggling the external PTP_CLK pin.
26:3	reserved	R	
2	RESPTP	W	Asserting this flag in combination with VALID causes PTPCLK, PTPTSCLK and PTPCLKRATE to be reset to their power-on defaults: the clocks are set to 0 and the rate is set to 1. Note that this may corrupt ingress timestamps when it happens in a time frame between frame start at the source port at 2 μs past the frame end at the source port. This flag always returns 0 on read.
1	CORRCLK4TS	W	Asserting this flag in combination with VALID causes subsequent timestamps on ingress and egress management frames to be taken based on PTPCLK. If the flag is de-asserted, timestamps are taken based on PTPTSCLK. The latter is also the default after power-on or reset. Note that taking ingress timestamps is only safe when they are taken based on PTPTSCLK. Using PTPCLKVAL for timestamping may deliver corrupted ingress timestamps if (a) PTPCLKRATE is larger than 1 : 2 or (b) the user writes to PTPCLK. Also changing the value of this field may corrupt ingress timestamps that occur in a time frame between frame start at the source port at 2 μs past the frame end at the source port. This flag always returns 0 on read.
0	PTPCLKADD	W	Asserting this flag in combination with VALID causes subsequent writes to PTPCLKVAL to be added to the clock rather than setting a new value. After power-on or reset the switch will be in set mode (i.e. writes to PTPCLKVAL will set a new value rather than adding an offset to the current value). This flag always returns 0 on read.

Table 60. PTP control registers 2 to 6 (address 16h to 12h)

Address	Bits	Symbol	Access	Description
16h	31:0	PTPPINDUR	W	This field specifies the interval between two edges of the external clock on pin PTP_CLK (in multiples of 8 ns). This field returns all 0s on read.
15h	31:0	PTPPINST[63:32]	W	This field specifies THE value of PTPTSCLK at which the switch starts toggling the external PTP_CLK pin. This field returns all 0s on read.
14h	31:0	PTPPINST[31:0]	W	
13h		PTPSCHTM[63:32]		This field specifies the time on the synchronized PTP clock at which the switch integrates into the schedule if the parameter CLKSRC in the Schedule Entry Points Parameter table (Table 16) is set to select the PTP clock to drive the schedule.
12h		PTPSCHTM[31:0]		

Table 61. PTP control registers 7 to 12 (address 1Dh to 18h)

Address	Bits	Symbol	Access	Description
1Dh	31:0	PTPCLKCORP	W	On write, this field defines the time between consecutive clock corrections applied to the schedule module in multiples of 8 ns (i.e. the first clock correction gets applied to the schedule execution module at time PTPSCHTM + PTPCLKCORP × 8 ns. The second clock correction is applied after a period of PTPCLKCORP × 8 ns). The field is ignored on read.
1Ch	31:0	PTPTSCLK[63:32]	R	This field is read only, write access is ignored. On read, this field contains the current value of the PTP timestamp clock that is used to timestamp MAC management frames on ingress and egress. Upon reading the least significant 32 bits, the most significant 32 bits are latched to a shadow register to provide a consistent snapshot of this 64-bit value. The field represents the time elapsed since power-on or reset in multiples of 8 ns measured on the free running clock.
1Bh	31:0	PTPTSCLK[31:0]		
1Ah	31:0	PTPCLKRATE	W	This field determines the speed of PTPCLKVAL. It implements a fixed-point clock rate value with a single-bit integer part and a 31-bit fractional part allowing for sub-ppb rate corrections. PTPCLKVAL ticks at the rate of PTPTSCLK multiplied by this field. So any value having the integer part set to 0 (i.e. bit 31 set to 0) will cause PTPCLKVAL to be slower than PTPTSCLK. Any value having the integer part set to one will cause PTPCLKVAL to be at least as fast as PTPTSCLK. E.g. a value of h'90000000 will cause PTPCLKVAL to tick 1.125 = ($2^0 + 2^{-3}$) faster than PTPTSCLK. This field returns all 0s on read.
19h	31:0	PTPCLKVAL[63:32]	R/W	Depending on the value of PTPCLKADD, a write to this field will cause the internal PTP clock counter to be set to the value provided by the host (in case PTPCLKADD is de-asserted) or to add the value provided by the host to the current value of the internal PTP clock counter. A read access to this field returns the current value of the internal (rate-corrected) PTP clock counter. On reading the least significant 32 bits, the most significant 32 bits are latched to a shadow register to provide a consistent snapshot of this 64-bit value.
18h	31:0	PTPCLKVAL[31:0]		

5.2.1.1 Control of the credit-based shaping blocks

Table 62. Credit-based shaping block register 1 (address 30h)

Bit	Symbol	Access	Description
31	VALID	W	When this flag is set, the host triggers a dynamic change of the entry with index SHAPER_ID. This flag always returns 0 on read.
30:20	reserved	R	
19:16	SHAPER_ID	W	On write, this field specifies the index of the credit-based shaper which is subject to dynamic reconfiguration; ignored on read. This field returns all 0s on read.
5:3	CBS_PORT	W	On write, this field specifies the port to which the credit-based shaper is assigned; ignored on read. This field returns all 0s on read.
2:0	CBS_PRIO	W	On write, this field specifies the priority queue to which the credit-based shaper is assigned; is ignored on read. This field returns all 0s on read.

Table 63. Credit-based shaping block registers 2 to 5 (address 2Fh to 2Ch)

Address	Bit	Symbol	Access	Description
2Fh	31:0	CREDIT_LO	W	On write, this field specifies the value at which the credit counter negatively saturates upon transmission of a frame. This can be used to reduce the gap between multiple burst high priority frames from the same queue, if shaping is enabled on this queue. This parameter defines the upper 32 bits of the credit counter, the lower 16 bits are set to 0. This field returns all 0s on read.
2Eh	31:0	CREDIT_HI	W	On write, this field specifies the value at which the credit counter positively saturates upon transmission of a frame. This can be used to limit the burst length of frames from a queue to which a shaper is applied to. The parameter defines the upper 32 bits of the credit counter, the lower 16 bits get set to zero. This field returns all 0s on read.
2Dh	31:0	SEND_SLOPE	W	On write, this field specifies the value at which the credit counter gets decreased at a rate of bytes per second times link speed. The credit counter gets decreased whenever the currently transmitted frame is sourced from the priority queue to which the shaper is applied. This field returns all 0s on read.
2Ch	31:0	IDLE_SLOPE	W	On write, this field specifies the value at which the credit counter gets increased at a rate of bytes per second times link speed. The counter gets increased whenever it is negative or the priority queue to which the shaper is applied to holds a frame ready for transmission but the media is occupied by a frame sourced from a different queue. This field returns all 0s on read.

5.2.2 Dynamic reconfiguration

Dynamic reconfiguration of the switch refers to those features of the programming interface's control area that allow specific parameter values of the loaded configuration to be changed at run time. The following sections provide details of the dynamic reconfiguration of specific parts of the loaded configuration.

5.2.2.1 Dynamic reconfiguration of the L2 Address Lookup table

The register entries in this section are used to dynamically reconfigure the L2 Address Lookup table (Table 10).

Table 64. L2 Address Lookup table reconfiguration register 1 (address 23h)

Bit	Symbol	Access	Description
31	VALID	R/W	The host sets this flag to trigger a dynamic change in the contents of the L2 Address Lookup table (if RDWRSET is set) or a read access (when RDWRSET is cleared). A write access from the host is only accepted when this flag is cleared. The flag remains set until the switch has completed the access and is cleared automatically afterwards. The address of the access is extracted from the INDEX field of ENTRY both for reads and writes.
30	RDWRSET	R/W	Determines whether access is a read access (flag is cleared) or a write access (flag is set). On read this flag displays the value most recently written by the host.
29	ERRORS	R	This flag return 0 on read when the VALID flag is set. If it is found to be set while the VALID flag is set, the most recent access resulted in an error. A write access is ignored.
28	LOCKEDS	R	This flag should return 0 on read when the VALID flag is set. The flag will also be cleared when the MGMTROUTE flag is read set. If a read operation finds the MGMTROUTE flag cleared, this flag is set if the most recent access operated on an entry that was either loaded at configuration or through dynamic reconfiguration (as opposed to automatically learned entries). A write access is ignored.
27	VALIDENT	R/W	In the case of a write access with the MGMTROUTE flag cleared, this flag determines if the respective entry should be marked valid. Marking an entry as invalid (i.e., clearing VALIDENT) has the effect that the entry at the respective position will be available again for address learning. For a read operation with the MGMTROUTE flag cleared, this flag will be set if the most recent access operated on a valid entry (i.e. an entry that contains either a programmed route or a dynamically learned one). This flag is ignored during a write access with the MGMTROUTE flag set. It will always be found cleared during a read access with MGMTROUTE set.
26	MGMTROUTE	R/W	On write, the host sets this flag to indicate that the request is targeted for a management route entry. In this case, the INDEX field of the ENTRY must point to one of the four supported management route entries. All management frames received from the port as indexed by the HOST_PORT field of the General Parameters block are checked for a match in the management forwarding entries and forwarded accordingly. The respective management frame is dropped if no matching entry is found. A management route entry is only valid if the ENFPORT flag is set and it is only valid for a single frame. The ENFPORT flag of the respective entry is cleared when a match is found. The host can use this flag as an acknowledgement. If the host provides several management route entries with identical values for the MACADDR, the one at the lowest index is used first. On read, the flag displays the value most recently written by the host. To specify if a PTP egress timestamp shall be captured on each port upon transmission of the frame, the LSB of VLANID in the ENTRY field provided by the host must be set. Bit 1 of VLANID then specifies the register where the timestamp for this port is stored in (see Table 38: if the respective management frame is sent on port n, the timestamp can be received at address $192 \times n$ where $n = 2 \times PORT + TSREG$. TSREG is the LSB of VLANID and indicates which of the two timestamp registers is to be used.
25:0	reserved	R	

Table 65. L2 Address Lookup table reconfiguration registers 2 to 4 (addresses 22h to 20h)

Address	Bits	Symbol	Access	Description
22h	31:0	ENTRY[75:44]	R/W	On write this field contains the new value for the entry in the L2 Address Lookup table to be updated if VALIDENT is set. If VALIDENT is not set on write, this field is ignored. The format to be used matches that specified for Table 10. On read this field displays the value most recently written by the host if this most recent access had the RDWRSET set. If the most recent host access requested a read (RDWRSET not set), the field displays the table entry data once the access completes (as indicated by the VALID flag being cleared).
21h	31:0	ENTRY[43:12]		
20h	31:20	ENTRY[11:0]		
	19:0	reserved	R	

5.2.2.2 Dynamic reconfiguration of the L2 Forwarding table

The register entries in this section are used to dynamically reconfigure the VLAN priority mapping and the port reachability limitations defined in the L2 Forwarding table (Table 13).

Table 66. L2 Forwarding table reconfiguration register 1 (address 26h)

Bit	Symbol	Access	Description
31	VALID	R/W	The host sets this flag to trigger a dynamic change in the entry indicated by INDEX. If this flag is found set on read, the switch is still busy processing the most recent update request (which may need up to 10 clock cycles to complete)
30	ERRORS	R	Write access is ignored. This value indicates whether a dynamic reconfiguration attempt was successful. An attempt may fail if at least one value in VLAN_PMAP (Table 13) exceeds the value configured for MAX_DYNP (Table 19). A write access is ignored.
29:5	reserved	R	
4:0	INDEX	W	Contains the index of the entry being dynamically reconfigured. Read access is ignored. This field returns all 0's on read.

Table 67. L2 Forwarding table reconfiguration registers 2 to 4 (addresses 25h to 24h)

Address	Bits	Symbol	Access	Description
25h	31:0	ENTRY[38:7]	W	Contains the entry indicated by INDEX in the same format as described in Table 13. This field returns all 0s on read.
24h	31:25	ENTRY[6:0]		

5.2.2.3 Dynamic reconfiguration of the VLAN Lookup table

The register entries in this section are used to dynamically reconfigure the VLAN lookup table (Table 12).

Table 68. VLAN lookup table reconfiguration register 1 (address 2Ah)

Bit	Symbol	Access	Description
31	VALID	W	The host sets this flag to trigger a dynamic change in the entry associated with the VLANID field of ENTRY (see Table 69 and Table 12). A write access is only accepted when this flag is cleared. Finding this flag set on read indicates that a dynamic reconfiguration is in progress.
30:28	reserved	R	
27	VALIDENT	W	For a write access, this flag determines if the respective entry should be marked valid. Marking an entry as invalid (i.e. clearing VALIDENT) causes the VLAN with the VLANID field specified in the ENTRY field to be deactivated on the switch. The flag returns 0 on read.
26:0	reserved	R	

Table 69. VLAN lookup table reconfiguration registers 2 to 4 (addresses 28h to 27h)

Address	Bit	Symbol	Access	Description
28h	31:0	ENTRY[36:5]	W	Contains the value for the entry in the VLAN Lookup table to be updated. The format to be used matches that specified in Table 12. Returns all 0's on read.
27h	31:25	ENTRY[4:0]		

5.2.2.4 Dynamic reconfiguration of MAC configuration table

The register entries in this section are used to dynamically reconfigure the parameters defined in the MAC configuration table.

Table 70. MAC configuration table reconfiguration register 1 (address 37h)

Bit	Symbol	Access	Description
31	VALID	W	The host sets this flag to trigger a dynamic change to the contents of the MAC Configuration table. The flag returns 0 on read.
30:29	SPEED	W	Used to set the port speed. 11 sets the speed to 10 Mbit/s; 10 sets the speed to 100 Mbit/s; 01 sets the speed to 1 Gbit/s; 00 disables the port. This field is only evaluated for ports whose SPEED parameter in the MAC Configuration table was set to 00 in the configuration initially loaded. If the loaded configuration contains the value 00 for a port, the host may change the port speed by setting this field to any value at any time. The field returns all 0's when read.
28:27	reserved	-	
26:24	PORT	W	Specifies the port affected by this dynamic reconfiguration. This field returns all 0s on read.
23	DRPDTAG	W	If this flag is set, double-tagged ingress traffic is dropped at the respective port (i.e. traffic that has a TPID defined in the General Parameters configuration block for either an outer or inner tag as well as traffic containing TPID2 in the outer tag - whether an inner tag exists or not). Flag affects L2 traffic only. Management traffic flows to the port regardless of the state of the INGRESS flag.
22	DRPUNTAG	W	If this flag is set, untagged ingress traffic is dropped at the respective port.
21	RETAG	W	When set, this flag enables retagging (using VLANID configured for the respective port but maintaining the priority value) of priority-tagged input on the respective port.
20	DYN_LEARN	W	This flag enables address learning at the respective port when set. Note that learning is independent of whether input traffic is enabled.
19	EGRESS	W	This flag enables output on the respective port when set.
18	INGRESS	W	This flag enables input on the respective port when set.
17	INGMIRR	W	If this flag is set, all traffic received on this port is forwarded to the mirror port as defined by the MIRR_PORT field in the General Parameters configuration block (provided it does not fail on any of the filtering rules). The field returns 0 on read.
16	EGRMIRR	W	If this flag is set, all traffic forwarded to this port except for locally generated PCFs is forwarded to the mirror port as defined by the MIRR_PORT field in the General Parameters configuration block. This field returns 0 on read.
15	reserved	R	
14:12	VLANPRIO	W	Defines the IEEE 802.1Q VLAN priority level that is used for tagging untagged incoming frames on this port. The field returns 0 on read.
11:0	VLANID	W	Contains the IEEE 802.1Q VLAN ID that is used for tagging untagged incoming frames on this port. This value defines the key for lookups in the VLAN Lookup table. The respective entry of the VLAN Lookup table must be defined and have the flag of port set in VMEMB_PORT (which may require dynamic reconfiguration of the entry as well), otherwise all untagged frames received on the port will trigger WRONGPORTS or VNOTFOUNDS status errors (see Table 32) and cause N_VLANERR to be increased (see Table 50). This field returns all 0s on read.

Table 71. MAC configuration table reconfiguration register 2 (address 36h)

Bit	Symbol	Access	Description
31:16	TPDELIN	W	Used to set a correction for updating the transparent clock of IEEE 1588v2 one-step event messages at the input port in multiples of 8 ns. This field is only evaluated for ports that have their respective SPEED parameter of MAC reconfiguration register 1 set to 00 in the configuration initially loaded. If the loaded configuration contains the value 00 for a port, however, the host may change the delay by setting this field accordingly at any time to any value. The field returns all 0s on read.
15:0	TPDELOUT	W	Used to set a correction for updating the transparent clock of IEEE 1588v2 one-step event messages at the output port in multiples of 8 ns. This field is only evaluated for ports that have their respective SPEED parameter of MAC reconfiguration register 1 set to 00 in the configuration initially loaded. If the loaded configuration contains the value 00 for a port, however, the host may change the delay by setting this field accordingly at any time to any value. The field returns all 0s on read.

5.2.2.5 Dynamic reconfiguration of the Retagging table

The register entries in this section are used to dynamically reconfigure the parameters defined in the Retagging table (Table 23).

Table 72. Retagging table reconfiguration register 1 (address 33h)

Bit	Symbol	Access	Description
31	VALID	W	The host sets this flag to trigger a dynamic change to the entry with index INDEX. If this flag is found set on read, the switch is still busy processing the most recent update request. The flag returns 0 on read.
30	ERRORS	R	Write access is ignored. If found set on read, the most recent access resulted in an error because it was issued prior to completing the configuration load procedure.
29	VALIDENT	W	Indicates that the entry at position INDEX should be enabled (in this case the host is supposed to have provided the data for the entry at ENTRY (Table 73) when asserted; the entry at position INDEX is disabled if this flag is found de-asserted). The flag returns 0 on read.
28:6	reserved	R	
5:0	INDEX	W	Contains the index of the entry being dynamically reconfigured. Read access is ignored. returns all 0s on read.

Table 73. Retagging table reconfiguration registers 2 to 4 (addresses 32h to 31h)

Address	Bits	Symbol	Access	Description
32h	31:0	ENTRY[63:32]	W	New retagging entry to be applied to INDEX. ENTRY is in the same format as described for the Retagging table.
31h	31:23	ENTRY[31:23]		
	22:0	reserved	R	

5.2.2.6 Dynamic reconfiguration of the general parameters configuration block

The register entries in this section are used to dynamically reconfigure the parameters defined in the General Parameters table (Table 22).

Table 74. General Parameters block reconfiguration register (address 34h)

Bit	Symbol	Access	Description
31	VALID	W	The host sets this flag to trigger a reconfiguration of the general parameters block. The flag always returns 0 on read.
30	ERRORS	R	Write access is ignored. If found set on read, the most recent access resulted in an error because dynamic reconfiguration was not enabled (as indicated by MIRR_PTACU flag being de-asserted).
29:3	reserved	R	
2:0	MIRRORP	W	Contains the current mirror port setting. A value greater than four turns off mirroring. Returns all 0s on read.

5.2.2.7 Dynamic reconfiguration of the VL Lookup table

The register entries in this section are used to dynamically reconfigure the parameters defined in the VL Lookup tables (Table 10 and Table 6).

Table 75. VL Lookup table reconfiguration register (address 35h)

Bit	Symbol	Access	Description
31	VALID	R/W	The host sets this flag to trigger an update of the VL Lookup table. If this flag is found set on read, the switch is still busy processing the most recent update request.
30	ERROR	R	Write access is ignored. If found set on read, the most recent access resulted in an error (and did not update any entry of the VL Lookup table) because either it was issued prior to completing the configuration load procedure, or VLLUPFORMAT in Table 22 is set to 0, or the VL Lookup table has not been loaded (i.e. there is no critical traffic enabled), or the host provided an index in INDEX that has not been loaded.
29:18	reserved	R	
21:17	EGMIRR	W	Contains the value to be written to the EGRMIRR field in VL Lookup table as indexed by INDEX. Returns all zeros on read.
16	INGMIRR	W	Contains the value to be written to the INGRMIRR field in VL Lookup table as indexed by INDEX. Returns zero on read.
15:10	reserved	R	
9:0	INDEX	W	Contains the index of the entry to be reconfigured. Returns all zeros on read.

5.2.2.8 Dynamic reconfiguration of the L2 Lookup Parameters table

The register entries in this section are used to dynamically reconfigure the parameters defined in the L2 Lookup table (Table 18). The SJA1105 allows the hash polynomial used for index generation in the forwarding process to be changed. Note that the hardware does not reorganize changes to the polynomial made in software.

Table 76. L2 Lookup Parameters table reconfiguration register (address 38h)

Bit	Symbol	Access	Description
31	VALID	R/W	The host sets this flag to trigger an update of the L2 Lookup Parameters table. If this flag is found set on read, the switch is still busy processing the most recent update request. The flag is only evaluated if CONFIGS in the Initial device configuration flag register (Table 26) is asserted. An update request issued before CONFIGS is asserted is processed (provided VALID remains asserted) once CONFIGS = 1.
30:8	reserved	R	
7:0	POLY	R/W	On write, contains the new value for the CRC polynomial. Returns the polynomial currently in use if VALID = 0 (undefined if VALID = 1).

5.3 Clock Generation Unit (CGU)

The CGU generates multiple internal clocks to drive the internal core and the xMII ports. Depending on the operating mode (see Table 24), selected clocks are used to drive the internal xMII interface and transmit clock pins (configured as TX_CLK, REF_CLK or TXC). The CGU clocking scheme is shown in Figure 10. PLL0 generates a 125 MHz clock used for the switch core and, optionally, RGMII. PLL1 generates a 50 MHz clock for RMII. The dividers, IDIV0 to IDIV4, are configurable per port and can divide the input clock down to 2.5 MHz. After reset, PLL0 is automatically set to provide a 125 MHz clock for the switch core and PLL1 is disabled. PLL1 must be manually enabled when a port is configured for RMII.

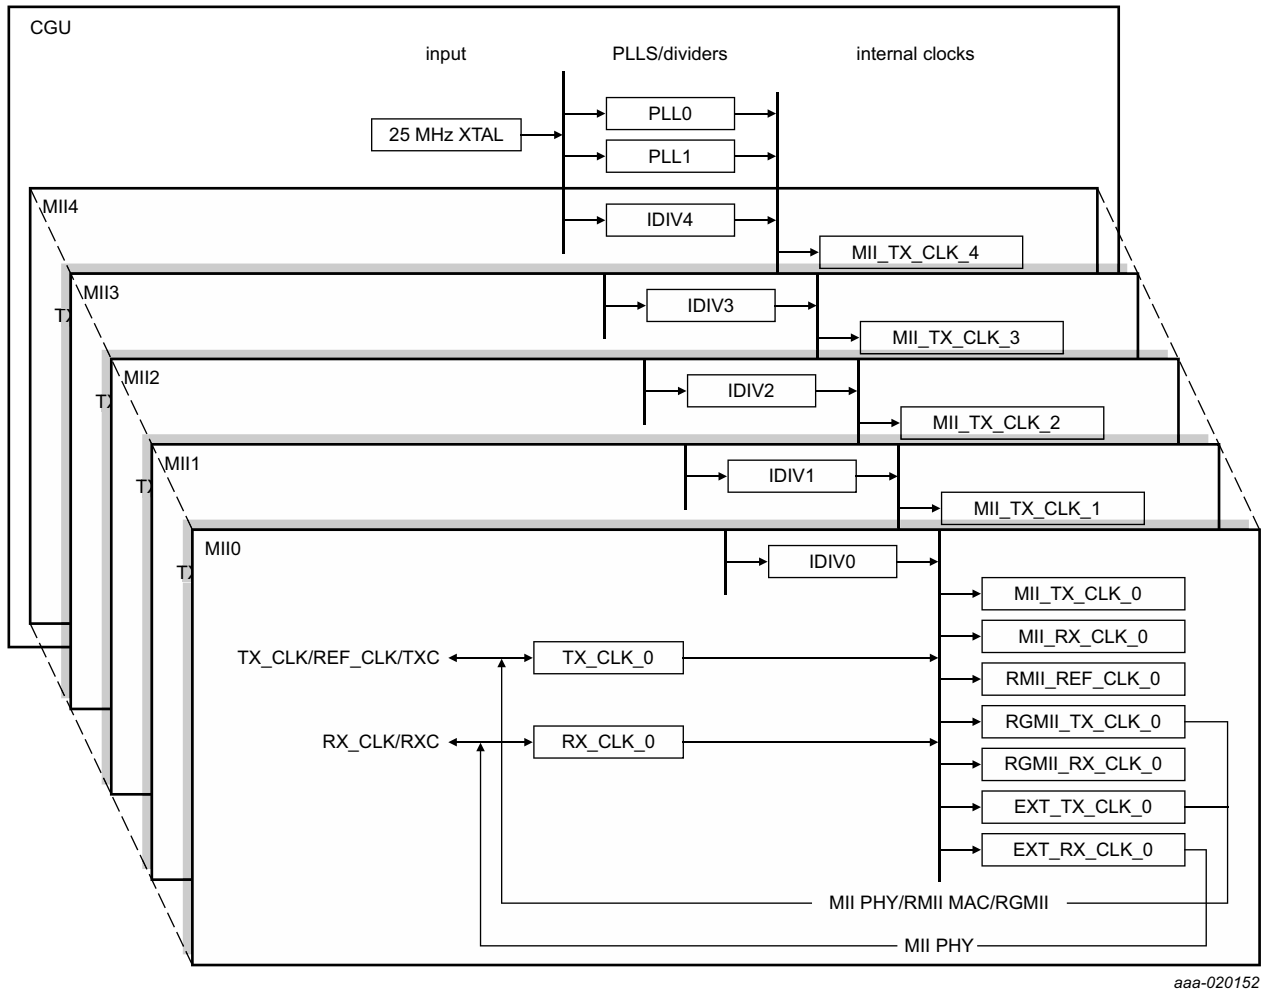

Fig 10. CGU clocking scheme

Clocks must be configured according to the xMII configuration selected for the associated port. Only required clocks are enabled at any time. If PORT 1 is configured for MII MAC via the static configuration interface, for example, only clocks MII0_MII_TX_CLK and MII0_MII_RX_CLK are needed. Table 77 provides an overview of which clocks are enabled for each port configuration.

Table 77. Clocks and associated xMII configuration

Mode	Internal clocks	Notes
MII-MAC	MIIx_MII_TX_CLK MIIx_MII_RX_CLK	used to clock the internal MII interface logic; the clocks are supplied externally by the PHY
MII-PHY	MIIx_MII_TX_CLK MIIx_MII_RX_CLK MIIx_EXT_TX_CLK MIIx_EXT_RX_CLK	MIIx_MII_TX_CLK and MIIx_MII_RX_CLK used to drive the internal MII interface MIIx_EXT_TX_CLK and MIIx_EXT_RX_CLK drive the TX_CLK and RX_CLK clock pins
RMII_MAC	MIIx_RMII_REF_CLK MIIx_EXT_TX_CLK	MIIx_RMII_REF_CLK drives the internal RMII interface MIIx_EXT_TX_CLK drives the REF_CLK output clock pin
RMII_PHY	MIIx_RMII_REF_CLK	MIIx_RMII_REF_CLK used to drive the internal RMII interface; the RMII interface reference clock is supplied externally in RMII_PHY mode
RGMII	MIIx_RGMII_RX_CLK	MIIx_RGMII_RX_CLK drives the internal RGMII interface (this clock is supplied externally via the RXC clock pin
	MIIx_RGMII_TX_CLK	MIIx_RGMII_TX_CLK drives the TXC output clock pin

Table 78. CGU register overview

Address	Name	Access	Reset value
100007h	PLL_0_S	R	00000000h
100009h	PLL_1_S	R	00000000h
10000Ah	PLL_1_C	R/W	0A000003h
10000Bh	IDIV_0_C	R/W	0A000000h
10000Ch	IDIV_1_C	R/W	0A000000h
10000Dh	IDIV_2_C	R/W	0A000000h
10000Eh	IDIV_3_C	R/W	0A000000h
10000Fh	IDIV_4_C	R/W	0A000000h
100013h	MII0_MII_TX_CLK	R/W	11000000h
100014h	MII0_MII_RX_CLK	R/W	11000000h
100015h	MII0_RMII_REF_CLK	R/W	0E000000h
100016h	MII0_RGMII_TX_CLK	R/W	11000000h
100018h	MII0_EXT_TX_CLK	R/W	11000000h
100019h	MII0_EXT_RX_CLK	R/W	11000000h
10001Ah	MII1_MII_TX_CLK	R/W	12000000h
10001Bh	MII1_MII_RX_CLK	R/W	12000000h
10001Ch	MII1_RMII_REF_CLK	R/W	0E000000h
10001Dh	MII1_RGMII_TX_CLK	R/W	12000000h
10001Fh	MII1_EXT_TX_CLK	R/W	12000000h
100020h	MII1_EXT_RX_CLK	R/W	12000000h
100021h	MII2_MII_TX_CLK	R/W	13000000h
100022h	MII2_MII_RX_CLK	R/W	13000000h
100023h	MII2_RMII_REF_CLK	R/W	0E000000h
100024h	MII2_RGMII_TX_CLK	R/W	13000000h
100026h	MII2_EXT_TX_CLK	R/W	13000000h
100027h	MII2_EXT_RX_CLK	R/W	13000000h
100028h	MII3_MII_TX_CLK	R/W	14000000h
100029h	MII3_MII_RX_CLK	R/W	14000000h
10002Ah	MII3_RMII_REF_CLK	R/W	0E000000h
10002Bh	MII3_RGMII_TX_CLK	R/W	14000000h
10002Dh	MII3_EXT_TX_CLK	R/W	14000000h
10002Eh	MII3_EXT_RX_CLK	R/W	14000000h
10002Fh	MII4_MII_TX_CLK	R/W	15000000h
100030h	MII4_MII_RX_CLK	R/W	15000000h
100031h	MII4_RMII_REF_CLK	R/W	0E000000h
100032h	MII4_RGMII_TX_CLK	R/W	15000000h
100034h	MII4_EXT_TX_CLK	R/W	15000000h
100035h	MII4_EXT_RX_CLK	R/W	15000000h

Table 79. Clock selection matrix

'd' indicates the reset value; 'a' indicates available clock sources; all other values are invalid.

Internal clock	Internal clock selection (via CLKSRC)																
	MII0_TX_CLK (CLKSRC = 00h)	MII0_RX_CLK (CLKSRC = 01h)	MII1_TX_CLK (CLKSRC = 02h)	MII1_RX_CLK (CLKSRC = 03h)	MII2_TX_CLK (CLKSRC = 04h)	MII2_RX_CLK (CLKSRC = 05h)	MII3_TX_CLK (CLKSRC = 06h)	MII3_RX_CLK (CLKSRC = 07h)	MII4_TX_CLK (CLKSRC = 08h)	MII4_RX_CLK (CLKSRC = 09h)	PLL0 (CLKSRC = 0Bh)	PLL1 (CLKSRC = 0Eh)	IDIV0 (CLKSRC = 11h)	IDIV1 (CLKSRC = 12h)	IDIV2 (CLKSRC = 13h)	IDIV3 (CLKSRC = 14h)	IDIV4 (CLKSRC = 15h)
MII0_MII_TX_CLK	a												d				
MII0_MII_RX_CLK		a											d				
MII0_RMII_REF_CLK	a											d					
MII0_RGMII_TX_CLK											a		d				
MII0_EXT_TX_CLK												a	d				
MII0_EXT_RX_CLK													d				
MII1_MII_TX_CLK			a											d			
MII1_MII_RX_CLK				a										d			
MII1_RMII_REF_CLK			a									d					
MII1_RGMII_TX_CLK											a			d			
MII1_EXT_TX_CLK												a		d			
MII1_EXT_RX_CLK														d			
MII2_MII_TX_CLK					a										d		
MII2_MII_RX_CLK						a									d		
MII2_RMII_REF_CLK					a							d					
MII2_RGMII_TX_CLK											a				d		
MII2_EXT_TX_CLK												a			d		
MII2_EXT_RX_CLK															d		
MII3_MII_TX_CLK							a									d	
MII3_MII_RX_CLK								a								d	
MII3_RMII_REF_CLK							a					d					
MII3_RGMII_TX_CLK											a					d	
MII3_EXT_TX_CLK												a				d	
MII3_EXT_RX_CLK																d	
MII4_MII_TX_CLK									a								d
MII4_MII_RX_CLK										a							d
MII4_RMII_REF_CLK									a		a	d					
MII4_RGMII_TX_CLK																	d
MII4_EXT_TX_CLK												a					d
MII4_EXT_RX_CLK																	d

Table 80. PLL_x_S clock status registers 0 and 1 - address 100007h and 100009h

*Legend: * reset value*

Bit	Symbol	Access	Value	Description
31:1	reserved	R	0h	
0	LOCK	R		PLL lock indicator:
			0*	PLL not locked
			1	PLL locked

Table 81. PLL_1_C control register (address 10000Ah; see Table 78)

*Legend: * reset value*

Bit	Symbol	Access	Value	Description
31:29	reserved	R/W	0h	
28:24	PLLCLKSRC	R/W	0Ah	input clock selection; must be set to 0Ah to select the 25 MHz reference clock
23:16	MSEL	R/W		M divider value:
			00h*	disabled
			01h	RMII clock generation
15:12	reserved	R/W	0h	
11	AUTOBLOCK	R/W		block clock automatically when settings are being changed to prevent glitches in the output clock:
			0*	disabled
			1	enabled
10	reserved	R/W	0	
9:8	PSEL	R/W		P divider value:
			00*	disabled
			01	50 MHz generation
7	DIRECT	R/W		direct clock output control:
			0	RMII setting: clock signal goes through post divider
			1*	clock signal goes directly to output
6	FBSEL	R/W		PLL feedback select:
			0*	disabled
			1	50 MHz generation
5:2	reserved	R/W	0h	
1	BYPASS	R/W		bypass:
			0*	PLL not bypassed; must be set to 0 to enable RMII clock generation
			1	PLL bypassed
0	PD	R/W		power down:
			0	PLL1 enabled
			1*	PLL1 disabled

Table 82. IDIV_0_C to IDIV_4_Ccontrol registers (addr. 10000Bh to 10000Fh; see Table 78)
*Legend: * reset value*

Bit	Symbol	Access	Value	Description
31:29	reserved	R/W	0h	
28:24	CLKSRC	R/W	0Ah	input clock selection; must be set to 0Ah to select the 25 MHz reference clock; all other values invalid
23:12	reserved	R/W	0h	
11	AUTOBLOCK	R/W		block 25 MHz reference clock automatically when configuration settings are being changed:
			0*	disabled
			1	enabled (recommended when configuration settings are being changed)
10:6	reserved	R/W	0h	
5:2	IDIV	R/W		integer IDIV divide by value:
			0000*	divide by 1
			0001	reserved
			“ “ “ “	“ “ “ “ “ “
			1000	reserved
			1001	divide by 10
1	reserved	R/W	0	
0	PD	R/W		IDIV power down:
			0*	IDIV enabled
			1	IDIV disabled

Table 83. MIIx clock control registers 1 to 30 (addresses 100013h to 100035h; see Table 78)
*Legend: * reset value*

Bit	Symbol	Access	Value	Description
31:29	reserved	R/W	0h	
28:24	CLKSRC	R/W	[1]	internal clock selection:
			00h	MII0_TX_CLK
			01h	MII0_RX_CLK
			02h	MII1_TX_CLK
			03h	MII1_RX_CLK
			04h	MII2_TX_CLK
			05h	MII2_RX_CLK
			06h	MII3_TX_CLK
			07h	MII3_RX_CLK
			08h	MII4_TX_CLK
			09h	MII4_RX_CLK
			0Bh	PLL0 (RGMII)
			0Eh	PLL1 (RMII)
			11h	IDIV0
			12h	IDIV1
			13h	IDIV2
			14h	IDIV3
			15h	IDIV4
23:12	reserved	R/W	0h	
11	AUTOBLOCK	R/W		block clock automatically when frequency is changing:
			0*	disabled
			1	enabled (recommended when clock source is being changed)
10:1	reserved	R/W	0h	
0	PD	R/W		MIIx clock power down:
			0*	MIIx clock enabled
			1	MIIx clock disabled

[1] Only the values shown are valid.

5.4 Reset Generation Unit (RGU)

The RGU provides reset sources across the device and can be used to trigger a software cold or warm reset.

Table 84. RGU register overview

Address	Name	Access	Reset value
100440h	RESET_CTRL	W	00000000h

Table 85. RESET_CTRL register - address 100440h

Bit	Symbol	Access	Value	Description
31:9	reserved	W	0h	
8:0	RESET	W		reset control:
			04h	trigger a cold reset
			08h	trigger a warm reset

5.5 Auxiliary Configuration Unit (ACU)

The auxiliary configuration unit controls the I/O characteristics and provides auxiliary functionality.

Table 86. ACU register overview

Address	Name	Access	Reset value	Description
100800h	CFG_PAD_MII0_TX	R/W	0x12121212	configuration register for TX pads of MII0
100801h	CFG_PAD_MII0_RX	R/W	0x02020212	configuration register for RX pads of MII0
100802h	CFG_PAD_MII1_TX	R/W	0x12121212	configuration register for TX pads of MII1
100803h	CFG_PAD_MII1_RX	R/W	0x02020212	configuration register for RX pads of MII1
100804h	CFG_PAD_MII2_TX	R/W	0x12121212	configuration register for TX pads of MII2
100805h	CFG_PAD_MII2_RX	R/W	0x02020212	configuration register for RX pads of MII2
100806h	CFG_PAD_MII3_TX	R/W	0x12121212	configuration register for TX pads of MII3
100807h	CFG_PAD_MII3_RX	R/W	0x02020212	configuration register for RX pads of MII3
100808h	CFG_PAD_MII4_TX	R/W	0x12121212	configuration register for TX pads of MII4
100809h	CFG_PAD_MII4_RX	R/W	0x02020212	configuration register for RX pads of MII4
100840h	CFG_PAD_MISC	R/W	0x00320412	configuration register for MISC pads
100880h	CFG_PAD_SPI	R/W	0x12040407	configuration register for SPI pads
100881h	CFG_PAD_JTAG	R/W	0x02000000	configuration register for JTAG pads
100900h	PORT_STATUS_MII0	R	0x0000001B	port configuration status register for MII0
100901h	PORT_STATUS_MII1	R	0x0000001B	port configuration status register for MII1
100902h	PORT_STATUS_MII2	R	0x0000001B	port configuration status register for MII2
100903h	PORT_STATUS_MII3	R	0x0000001B	port configuration status register for MII3
100904h	PORT_STATUS_MII4	R	0x0000001B	port configuration status register for MII4
100A00h	TS_CONFIG	R/W	0x00000065	temperature sensor configuration register
100A01h	TS_STATUS	R	0x00000000	temperature sensor status register
100A80h	RGMII_MEAS_SETUP	R/W	0x00000000	RGMII input timing measurement setup register
100BC0h	PROD_CFG	R	-	product configuration status register 1
100BC3h	PROD_ID	R	-	product configuration status register 2

Table 87. CFG_PAD_MIIx_TX registers 1 to 5 (addresses 100800/2/4/6/8h; see Table 86)
*Legend: * reset value*

Bit	Symbol	Access	Value	Description
31:29	not used	R	0*	
28:27	D32_OS	R/W		TXD3 and TXD2 pad output stage speed selection:
			00	very low noise/low speed
			01	low noise/medium speed
			10*	medium noise/fast speed
			11	high noise/high speed
26	not used	R	0*	
25:24	D32_IPUD	R/W		TXD3 and TXD2 pad input stage (weak) pull-up/pull-down selection (when pins configured as inputs):
			00	pull-up
			01	repeater
			10*	plain input
			11	pull-down
23:21	not used	R	0*	
20:19	D10_OS	R/W		TXD1 and TXD0 pad output stage speed selection:
			00	very low noise/low speed
			01	low noise/medium speed
			10*	medium noise/fast speed
			11	high noise/high speed
18	not used	R	0*	
17:16	D10_IPUD	R/W		TXD1 and TXD0 pad input stage (weak) pull-up/pull-down selection (when pins configured as inputs):
			00	pull-up
			01	repeater
			10*	plain input
			11	pull-down
15:13	not used	R	0*	
12:11	CTRL_OS	R/W		TX_EN/TX_CTL and TX_ER pad output stage speed selection:
			00	very low noise/low speed
			01	low noise/medium speed
			10*	medium noise/fast speed
			11	high noise/high speed
10	not used	R	0*	
9:8	CTRL_IPUD	R/W		TX_EN/TX_CTL and TX_ER pad input stage (weak) pull-up/pull-down selection:
			00	pull-up
			01	repeater
			10*	plain input
			11	pull-down
7:5	not used	R	0*	

Table 87. CFG_PAD_MIIx_TX registers 1 to 5 (addresses 100800/2/4/6/8h; see Table 86) *…continued*
*Legend: * reset value*

Bit	Symbol	Access	Value	Description
4:3	CLK_OS	R/W		TX_CLK/REF_CLK/TXC pad output stage speed selection:
			00	very low noise/low speed
			01	low noise/medium speed
			10*	medium noise/fast speed
			11	high noise/high speed
2	CLK_IH	R/W		TX_CLK/REF_CLK/TXC pad input stage hysteresis:
			0*	non-Schmitt
			1	Schmitt
1:0	CLK_IPUD	R/W		TX_CLK/REF_CLK/TXC pad input stage (weak) pull-up/pull-down selection:
			00	pull-up
			01	repeater
			10*	plain input
			11	pull-down

Table 88. CFG_PAD_MIIx_RX registers 1 to 5 (addresses 100801/3/5/7/9h; see Table 86)
*Legend: * reset value*

Bit	Symbol	Access	Value	Description
31:27	not used	R	0*	
26	D32_IH	R/W		RXD3 and RXD2 pad input stage hysteresis:
			0*	non-Schmitt
			1	Schmitt
25:24	D32_IPUD	R/W		RXD2 and RXD3 pad input stage (weak) pull-up/pull-down selection:
			00	pull-up
			01	repeater
			10*	plain input
			11	pull-down
23:19	not used	R	0*	
18	D10_IH	R/W		RXD1 and RXD0 pad input stage hysteresis:
			0*	non-Schmitt
			1	Schmitt
17:16	D10_IPUD	R/W		RXD1 and RXD0 pad input stage (weak) pull-up/pull-down selection:
			00	pull-up
			01	repeater
			10*	plain input
			11	pull-down
15:11	not used	R	0*	
10	CTRL_IH	R/W		RX_DV/CRS_DV/RX_CTL and RX_ER pad input stage hysteresis:
			0*	non-Schmitt
			1	Schmitt

Table 88. CFG_PAD_MIIx_RX registers 1 to 5 (addresses 100801/3/5/7/9h; see Table 86) *...continued*
*Legend: * reset value*

Bit	Symbol	Access	Value	Description
9:8	CTRL_IPUD	R/W		RX_DV/CRS_DV/RX_CTL and RX_ER pad input stage (weak) pull-up/pull-down selection:
			00	pull-up
			01	repeater
			10*	plain input
			11	pull-down
7:5	not used	R	0*	
4:3	CLK_OS	R/W		RX_CLK/RXC pad output stage speed selection:
			00	very low noise/low speed
			01	low noise/medium speed
			10*	medium noise/fast speed
			11	high noise/high speed
2	CLK_IH	R/W		RX_CLK/RXC pad input stage hysteresis:
			0*	non-Schmitt
			1	Schmitt
1:0	CLK_IPUD	R/W		RX_CLK/RXC pad input stage (weak) pull-up/pull-down selection:
			00	pull-up
			01	repeater
			10*	plain input
			11	pull-down

Table 89. CFG_PAD_MISC (address 100840h; see Table 86)
*Legend: * reset value*

Bit	Symbol	Access	Value	Description
31:22	not used	R	0*	
21	PTPCLK_EN	R/W		PTP_CLK pad output stage enable (active LOW):
			0	enabled
			1*	disabled
20:19	PTPCLK_OS	R/W		PTP_CLK pad output stage speed selection:
			00	very low noise/low speed
			01	low noise/medium speed
			10*	medium noise/fast speed
			11	high noise/high speed
18	not used	R	0*	
17:16	PTPCLK_IPUD	R/W		PTP_CLK pad input stage (weak) pull-up/pull-down:
			00	pull-up
			01	repeater
			10*	plain input
			11	pull-down
15:11	not used	R	0*	

Table 89. CFG_PAD_MISC (address 100840h; see Table 86) *…continued*
*Legend: * reset value*

Bit	Symbol	Access	Value	Description
10	RSTN_IH	R/W		RST_N pad input stage hysteresis:
			0	non-Schmitt
			1*	Schmitt
9:8	RSTN_IPUD	R/W		RST_N pad input stage (weak) pull-up/pull-down:
			00*	pull-up
			01	repeater
			10	plain input
			11	pull-down
7:6	not used	R	0*	
5	CLKOUT_EN	R/W		CLK_OUT pad output stage enable (active LOW):
			0*	enabled
			1	disabled
4:3	CLKOUT_OS	R/W		CLK_OUT pad output stage speed selection:
			00	very low noise/low speed
			01	low noise/medium speed
			10*	medium noise/fast speed
			11	high noise/high speed
2	not used	R	0*	
1:0	CLKOUT_IPUD	R/W		CLK_OUT pad input stage (weak) pull-up/pull-down:
			00	pull-up
			01	repeater
			10*	plain input
			11	pull-down

Table 90. CFG_PAD_SPI (address 100880h; see Table 86)
*Legend: * reset value*

Bit	Symbol	Access	Value	Description
31:29	not used	R	0*	
28:27	SDO_OS	R/W		SDO pad output stage speed selection:
			00	very low noise/low speed
			01	low noise/medium speed
			10*	medium noise/fast speed
			11	high noise/high speed
26	not used	R	0*	
25:24	SDO_IPUD	R/W		SDO pad input stage (weak) pull-up/pull-down:
			00	pull-up
			01	repeater
			10*	plain input
			11	pull-down
23:19	not used	R	0*	

Table 90. CFG_PAD_SPI (address 100880h; see Table 86) *…continued*
*Legend: * reset value*

Bit	Symbol	Access	Value	Description
18	SDI_IH	R/W		SDI pad input stage hysteresis:
			0	non-Schmitt
			1*	Schmitt
17:16	SDI_IPUD	R/W		pad input stage (weak) pull-up/pull-down:
			00*	pull-up
			01	repeater
			10	plain input
			11	pull-down
15:11	not used	R	0*	
10	SSN_IH	R/W		SS_N pad input stage hysteresis:
			0	non-Schmitt
			1*	Schmitt
9:8	SSN_IPUD	R/W		SS_N pad input stage (weak) pull-up/pull-down:
			00*	pull-up
			01	repeater
			10	plain input
			11	pull-down
7:3	not used	R	0*	
2	SCK_IH	R/W		SCK pad input stage hysteresis:
			0	non-Schmitt
			1*	Schmitt
1:0	SCK_IPUD	R/W		SCK pad input stage (weak) pull-up/pull-down:
			00	pull-up
			01	repeater
			10	plain input
			11*	pull-down

Table 91. CFG_PAD_JTAG (address 100881h; see Table 86)
*Legend: * reset value*

Bit	Symbol	Access	Value	Description
31:26	not used	R	0*	
25:24	TDO_IPUD	R/W		TDO pad input stage (weak) pull-up/pull-down:
			00	pull-up
			01	repeater
			10*	plain input
			11	pull-down
23:18	not used	R	0*	

Table 91. CFG_PAD_JTAG (address 100881h; see Table 86) *…continued*
*Legend: * reset value*

Bit	Symbol	Access	Value	Description
17:16	TDI_IPUD	R/W		TDI pad input stage (weak) pull-up/pull-down:
			00*	pull-up
			01	repeater
			10	plain input
			11	pull-down
15:10	not used	R	0*	
9:8	TRSTNTMS_IPUD	R/W		TRST_N/TMS pad input stage (weak) pull-up/pull-down:
			00*	pull-up
			01	repeater
			10	plain input
			11	pull-down
7:2	not used	R	0*	
1:0	TCK_IPUD	R/W		TCK pad input stage (weak) pull-up/pull-down:
			00*	pull-up
			01	repeater
			10	plain input
			11	pull-down

Table 92. PORT_STATUS_MII0 to MII4 registers (addresses 100900h to 100904h; see Table 86)
*Legend: * reset value*

Bit	Symbol	Access	Value	Description
31:5	not used	R	0h	
4:3	SPEED	R		port speed setting:
			00	10 Mbit/s
			01	100 Mbit/s
			10	1 Gbit/s
			11*	1 Gbit/s
2	PHY_MAC	R/W		port mode setting:
			0*	MAC mode
			1	PHY mode
1:0	xMII_MODE	R		xMII interface:
			00	MII interface
			01	RMII interface
			10	RGMII interface
			11*	MII interface

Table 93. TS_CONFIG (address 100A00h; see Table 86)

*Legend: * reset value*

Bit	Symbol	Access	Value	Description
31:7	not used	R	0h	
6	PD	R/W		temperature sensor power-down control:
			0	temperature sensor active
			1*	power-down temperature censor
5:0	THRESHOLD	R/W	xxxxxx	temperature threshold selection (valid values in range 01h to 39h); default value 25h; see Table 94 for values

Table 94. Temperature threshold selection (selected via bits THRESHOLD; see Table 93)

Bit value	Temp. (°C)	Bit value	Temp. (°C)	Bit value	Temp. (°C)	Bit value	Temp. (°C)	Bit value	Temp. (°C)
000000	invalid	001000	−11.4	010000	+25.6	011000	+63.3	100000	+102.5
000001	−45.7	001001	−6.1	010001	+30.9	011001*	+67.9	100001	+106.9
000010	−41.7	001010	−2.1	010010	+36.4	011010	+72.6	100010	+111.4
000011	−37.5	001011	+2.1	010011	+42.0	011011	+77.4	100011	+116.0
000100	−33.0	001100	+6.5	010100	+46.1	011100	+82.4	100100	+120.7
000101	−28.4	001101	+11.0	010101	+50.2	011101	+87.5	100101	+125.5
000110	−23.5	001110	+15.7	010110	+54.5	011110	+92.8	100110	+130.5
000111	−18.3	001111	+20.6	010111	+58.8	011111	+98.2	100111	+135.5

Table 95. TS_STATUS (address 100A01h; see Table 86)

Bits	Symbol	Access	Value	Description
31:1	not used	R	0h	
0	EXCEEDED	R		temperature detection status:
			0*	temperature below threshold
			1	temperature above threshold

Table 96. PROD_CFG (address 100BC0h; see Table 86)

Bits	Symbol	Access	Value	Description
31:1	not used	R	0h	
0	DISABLE_TTETH	R		TTEthernet features status:
			0	TTEthernet features of switch enabled
			1	TTEthernet features of switch disabled

Table 97. PROD_ID (address 100BC3h; see Table 86)

Bits	Symbol	Access	Value	Description
31:20	not used	R	0h	
19:4	PART_NR	R	9A83h	part number
3:0	VERSION	R	2h	version

6. Abbreviations

Table 98. Abbreviations

Acronym	Description
CRC	Cyclic Redundancy Check
LSB	Least Significant Bit
MAC	Media Access Control
MII	Media Independent Interface
MSB	Most Significant Bit
PHY	Physical Layer (of the interface)
SOF	Start Of Frame
VLAN	Virtual Local Area Network

7. Legal information

7.1 Definitions

Draft — The document is a draft version only. The content is still under internal review and subject to formal approval, which may result in modifications or additions. NXP Semiconductors does not give any representations or warranties as to the accuracy or completeness of information included herein and shall have no liability for the consequences of use of such information.

7.2 Disclaimers

Limited warranty and liability — Information in this document is believed to be accurate and reliable. However, NXP Semiconductors does not give any representations or warranties, expressed or implied, as to the accuracy or completeness of such information and shall have no liability for the consequences of use of such information. NXP Semiconductors takes no responsibility for the content in this document if provided by an information source outside of NXP Semiconductors.

In no event shall NXP Semiconductors be liable for any indirect, incidental, punitive, special or consequential damages (including - without limitation - lost profits, lost savings, business interruption, costs related to the removal or replacement of any products or rework charges) whether or not such damages are based on tort (including negligence), warranty, breach of contract or any other legal theory.

Notwithstanding any damages that customer might incur for any reason whatsoever, NXP Semiconductors' aggregate and cumulative liability towards customer for the products described herein shall be limited in accordance with the *Terms and conditions of commercial sale* of NXP Semiconductors.

Right to make changes — NXP Semiconductors reserves the right to make changes to information published in this document, including without limitation specifications and product descriptions, at any time and without notice. This document supersedes and replaces all information supplied prior to the publication hereof.

Suitability for use in automotive applications — This NXP Semiconductors product has been qualified for use in automotive applications. Unless otherwise agreed in writing, the product is not designed, authorized or warranted to be suitable for use in life support, life-critical or safety-critical systems or equipment, nor in applications where failure or malfunction of an NXP Semiconductors product can reasonably be expected to result in personal injury, death or severe property or environmental damage. NXP Semiconductors and its suppliers accept no liability for inclusion and/or use of NXP Semiconductors products in such equipment or applications and therefore such inclusion and/or use is at the customer's own risk.

Applications — Applications that are described herein for any of these products are for illustrative purposes only. NXP Semiconductors makes no representation or warranty that such applications will be suitable for the specified use without further testing or modification.

Customers are responsible for the design and operation of their applications and products using NXP Semiconductors products, and NXP Semiconductors accepts no liability for any assistance with applications or customer product design. It is customer's sole responsibility to determine whether the NXP Semiconductors product is suitable and fit for the customer's applications and products planned, as well as for the planned application and use of customer's third party customer(s). Customers should provide appropriate design and operating safeguards to minimize the risks associated with their applications and products.

NXP Semiconductors does not accept any liability related to any default, damage, costs or problem which is based on any weakness or default in the customer's applications or products, or the application or use by customer's third party customer(s). Customer is responsible for doing all necessary testing for the customer's applications and products using NXP Semiconductors products in order to avoid a default of the applications and the products or of the application or use by customer's third party customer(s). NXP does not accept any liability in this respect.

Export control — This document as well as the item(s) described herein may be subject to export control regulations. Export might require a prior authorization from competent authorities.

Evaluation products — This product is provided on an "as is" and "with all faults" basis for evaluation purposes only. NXP Semiconductors, its affiliates and their suppliers expressly disclaim all warranties, whether express, implied or statutory, including but not limited to the implied warranties of non-infringement, merchantability and fitness for a particular purpose. The entire risk as to the quality, or arising out of the use or performance, of this product remains with customer.

In no event shall NXP Semiconductors, its affiliates or their suppliers be liable to customer for any special, indirect, consequential, punitive or incidental damages (including without limitation damages for loss of business, business interruption, loss of use, loss of data or information, and the like) arising out the use of or inability to use the product, whether or not based on tort (including negligence), strict liability, breach of contract, breach of warranty or any other theory, even if advised of the possibility of such damages.

Notwithstanding any damages that customer might incur for any reason whatsoever (including without limitation, all damages referenced above and all direct or general damages), the entire liability of NXP Semiconductors, its affiliates and their suppliers and customer's exclusive remedy for all of the foregoing shall be limited to actual damages incurred by customer based on reasonable reliance up to the greater of the amount actually paid by customer for the product or five dollars (US$5.00). The foregoing limitations, exclusions and disclaimers shall apply to the maximum extent permitted by applicable law, even if any remedy fails of its essential purpose.

Translations — A non-English (translated) version of a document is for reference only. The English version shall prevail in case of any discrepancy between the translated and English versions.

7.3 Trademarks

Notice: All referenced brands, product names, service names and trademarks are the property of their respective owners.

8. Contents

For more information, please visit: http://www.nxp.com
For sales office addresses, please send an email to: salesaddresses@nxp.com

Date of release: 25 January 2017
Document identifier: UM10944