

COM-5403SOFT IP/TCP CLIENT/UDP/ARP/PING STACK for GbE VHDL SOURCE CODE OVERVIEW

Overview

Gigabit-speed IP protocols like TCP/IP can demand a high level of computation on processors. The trend has been to move the implementation of these fast but highly repetitive tasks to a TCP offload engine (TOE) to free the application processor from frequent interrupts.

The COM-5403SOFT is a generic Internet protocol stack (including the [VHDL source code](#)) designed to support 1Gbps speed on low-cost FPGAs. It is designed to achieve 950+ Mbps (UDP) or 450+ Mbps (per TCP client) throughputs on Gigabit Ethernet medium.

The following protocols are implemented in modular VHDL components: TCP client¹, UDP frames, ARP, PING, IP to MAC address routing table, IGMP (multicast) and DHCP client. Ancillary components are also included for streaming, test signal generation and bit error rate measurement.

These components can be instantiated as needed for the application. For a UDP receive-only application, one must instantiate *packet_parsing.vhd*, *arp.vhd*, *udp_rx.vhd*. The maximum number of concurrent TCP connections can be adjusted prior to VHDL synthesis depending on the available FPGA resources.

Wireshark Libpcap network capture files can be used as receiver stimulus for simulation purposes².

The code is written specifically for IEEE 802.3 Ethernet packet encapsulation (RFC 894), IPv4 protocols.

The code interfaces seamlessly with the COM-5401SOFT Tri-mode 10/100/1000 Mbps Ethernet MAC for the MAC / PHY layers

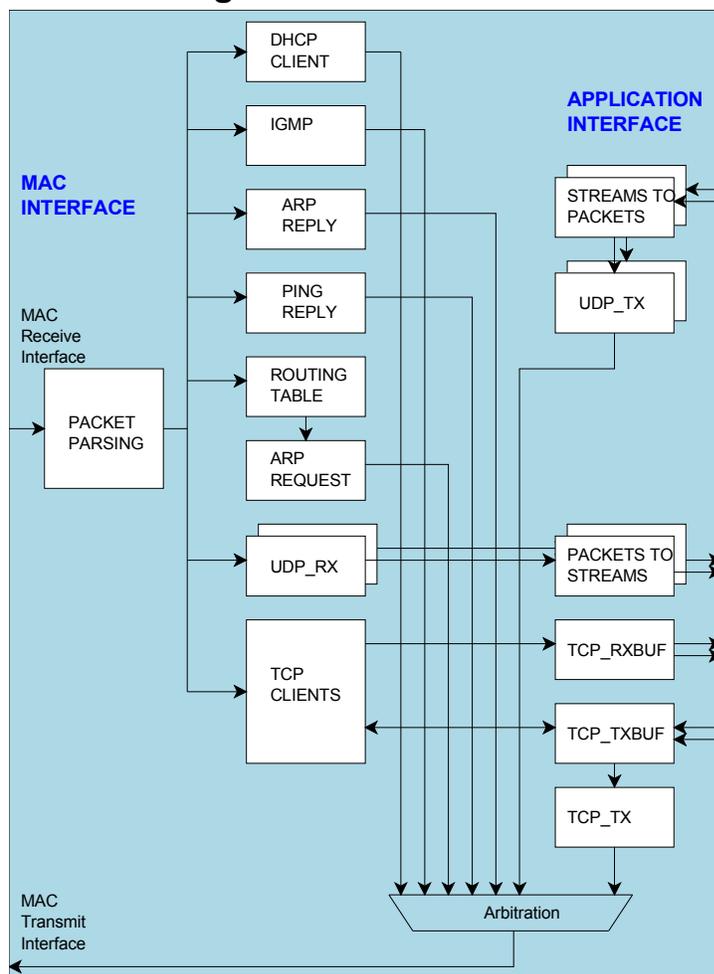
¹ See COM-5402SOFT for TCP server.

² Excludes TCP simulation

implementation. Code to interface with the Xilinx Tri-Mode Ethernet MAC (TEMAC) is also included. More generally, the MAC interface is generic and simple enough to interface with any Ethernet MAC component with minimum glue logic.

The component's very efficient implementation makes it suitable for multiple concurrent TCP and UDP streams instantiations within a small FPGA.

Block Diagram



Portable VHDL code

The code is written in generic standard VHDL and is thus portable to a variety of FPGAs. The code was developed and tested on a Xilinx 7-series FPGA but is expected to work similarly on other targets. No manufacturer-specific primitive is used.

Device Utilization Summary

Device: Xilinx Spartan-6

	TCP client (1) and ancillary protocols
Flip Flops	2666
LUTs	4198
RAMB16s	10
DSP48A1s	0
GCLKs	3
DCMs/PLLs	0

	TCP clients (2) and ancillary protocols
Flip Flops	3125
LUTs	5162
RAMB16s	14
DSP48A1s	0
GCLKs	3
DCMs/PLLs	0

	TCP client (1) and ancillary protocols, UDP tx (1), UDP rx (1)
Flip Flops	3115
LUTs	5163
RAMB16s	11
DSP48A1s	0
GCLKs	3
DCMs/PLLs	0

Throughput Performance Examples

Test setup1:

1 bidirectional connection between TCP server and TCP client over Gigabit Ethernet. 120 MHz FPGA processing clock. Measured sustained throughput: 452 Mbits/s concurrently in each direction.

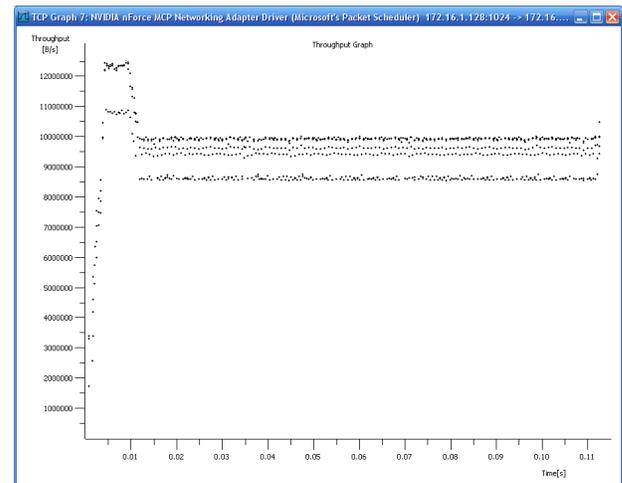
Test setup2:

512-byte UDP packets sent point-to-point over a LAN cable. Xilinx Spartan-6 FPGA –2 speed grade.

Measured: 0 bit error, payload throughput 878.5 Mbits/s. This matches the theoretical throughput (accounting for Ethernet, IP and UDP overhead and slower (120 MHz) user clock). The maximum throughput for this UDP frame size is 915 Mbits/s when user clock is 125 MHz or above.

Test setup 3:

TCP server transmit throughput on 100 Mbps LAN Wireshark measurement on receive PC. Average throughput 93 Mbps.



TCP Latency Performance Examples

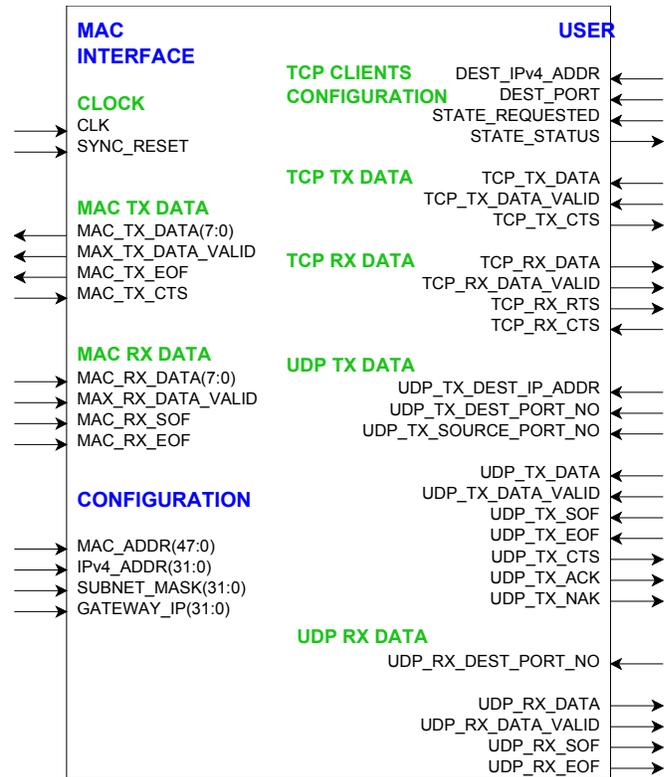
The transmit and receive latency depend on the frame length. For a maximum frame length of 1460 bytes, FPGA 125 MHz processing clock:

- Transmit latency (from the 1st byte of payload data input to the 1st byte of payload data output to the Ethernet MAC): 23.9µs
- Receive latency (from the 1st byte of Ethernet MAC input to the 1st byte of payload data output): 12.2µs

If latency is more important than throughput, the transmit segmentation threshold can be reduced to X payload bytes. In this more general case,

- Transmit latency (from the 1st byte of payload data input to the 1st byte of payload data output to the Ethernet MAC): $0.5 + 2X/125$ µs
- The receive latency (from the 1st byte of Ethernet MAC input to the 1st byte of payload data output): $0.5 + X/125$ µs

- Interfaces



User Interface

This interface comprises three primary signal groups: MAC interface (direct connection to COM-5401SOFT MAC core or equivalent), TCP streams, UDP frames or UDP streams to/from the user application.

All signals are clock synchronous with a user-selected clock CLK (it does not have to be the same as the PHY clock). To guarantee a 1 Gbps throughput, a minimum 125 MHz clock speed is required.

The user interface is buffered by internal elastic buffers in both tx/rx directions.

Configuration

The key configuration parameters are brought to the interface so that the user can change them dynamically at run-time. Other, more arcane, parameters are fixed at the time of VHDL synthesis.

Pre-synthesis configuration parameters

The following configuration parameters are set prior to synthesis in the *com5402pkg.vhd* package or at the top level component *com5403.vhd*.

<i>Configuration parameters in com5402pkg.vhd</i>	<i>Description</i>
Maximum number of concurrent TCP streams	NTCPSTREAMS_MAX
<i>Configuration parameters in com5403.vhd</i>	<i>Description</i>
DHCP client enable	DHCP_CLIENT_EN '1' to instantiate a DHCP client within. DHCP is a protocol used to dynamically assign IP addresses at power up from remote DHCP servers, like a gateways. '0' when a fixed (static) IP address is defined by the user.
Transmit UDP enabled	NUDPTX '1' to enable, '0' to disable
Receive UDP enabled	NUDPRX '1' to enable, '0' to disable
IGMP enabled	IGMP_EN '1' to enable UDP multicast (which requires IGMP)
TCP streams	NTCPSTREAMS. Number of concurrent TCP streams instantiated for this component. Each additional TCP stream requires additional resources (RAM block, logic). Must be less than or equal to NTCPSTREAMS_MAX
Inactive input stream timeout	TX_IDLE_TIMEOUT When segmenting a transmit stream, a packet will be sent out with pending data if no new data was received within the specified timeout. Expressed as integer multiple of 4µs.

TCP keep-alive period	TCP_KEEPALIVE_PERIOD period in seconds for sending no data keepalive frames. "Typically TCP Keepalives are sent every 45 or 60 seconds on an idle TCP connection, and the connection is dropped after 3 sequential ACKs are missed"
Elastic buffer size	Customized I/O elastic buffer sizes for various TCP and UDP components. Expressed number of address bits in Byte-wide RAM blocks. For example ADDR_WIDTH => 12 defines a 4KB buffer. Also expressed as an integer number NBUFS of 16Kbits RAM blocks. NBUFS is typically restricted to 1,2,4 or 8 (see code comments). Defined as generic parameters in the following components: <i>tcp_rxbufndemux2.vhd, tcp_txbuf.vhd, udp_tx, stream_2_packets.vhd.</i>

<i>Configuration parameters in arp_cache2.vhd</i>	<i>Description</i>
Routing table refresh period	Refresh period for this routing table. Expressed as an integer multiple of 100ms. Default value is 3000 (5 minutes).
REFRESH_PERIOD(19:0)	
<i>Configuration parameters in stream_2_packets.vhd</i>	<i>Description</i>
Maximum packet size when segmenting a stream to packets	When segmenting a transmit stream, a packet will be sent out as soon as MAX_PACKET_SIZE bytes are collected. The recommended size is 512 bytes for a low overhead.
MAX_PACKET_SIZE	
Retransmission timer	A re-transmission attempt will be made periodically until routing information is available and the transmit path to the MAC is available. The retry period is expressed as an integer multiple of 4µs.
TX_RETRY_TIMEOUT	

Run-time configuration parameters

The user can set and modify the following controls at run-time. All controls are synchronous with the user-supplied global CLK.

<i>Run-time configuration</i>	<i>Description</i>
MAC address <code>MAC_ADDR(47:0)</code>	This network node 48-bit MAC address. The user is responsible for selecting a unique 'hardware' address for each instantiation. Natural bit order: enter x0123456789ab for the MAC address 01:23:45:67:89:ab It is essential that this input matches the MAC address used by the MAC/PHY.
Dynamic vs static IP <code>DYNAMIC_IP</code>	'1' for dynamic addressing '0' for static IP address. The device IP address can be assigned dynamically by an external DHCP server, or defined as static address by the user. Dynamic addressing requires instantiating a DHCP client: set the generic parameter <code>DHCP_CLIENT_EN = '1'</code> .
IPv4 address <code>REQUESTED_IPv4_ADDR(31:0)</code>	Static address when <code>DYNAMIC_IP = '0'</code> Last dynamically assigned address when <code>DYNAMIC_IP = '1'</code> . Address 0.0.0.0 can also be used in conjunction with dynamic addressing if the user does not 'remember' the last dynamic IP address. 4 bytes for IPv4. Byte order: (MSB)192.68.1.30(LSB)
Multicast IP address <code>MULTICAST_IP_ADDR(31:0)</code>	to receive UDP multicast messages. One multicast address only 0.0.0.0 to signify that IP multicasting is not supported here.
Subnet Mask <code>IPv4_SUBNET_MASK(31:0)</code>	Subnet mask to assess whether an IP address is local (LAN) or remote (WAN) Byte order: (MSB)255.255.255.0(LSB)
Gateway IP address <code>IPv4_GATEWAY_ADDR(31:0)</code>	One gateway through which packets with a WAN destination are directed. Byte order: (MSB)192.68.1.1(LSB)

Transmit UDP destination IP address <code>UDP_TX_DEST_IP_ADDR</code>	The UDP destination IP address can be modified dynamically on a frame-by-frame basis.
Transmit UDP destination port number <code>UDP_TX_DEST_PORT_NO</code>	The UDP destination port number can be modified dynamically on a frame-by-frame basis.
Transmit UDP source port number <code>UDP_TX_SOURCE_PORT_NO</code>	The UDP source port number can be modified dynamically on a frame-by-frame basis.
Check UDP port number <code>CHECK_UDP_RX_DEST_PORT_NO</code>	check the received UDP frame destination port number matches <code>UDP_RX_DEST_PORT_NO</code> (1) or ignore it (0) In the latter case, the application is responsible for checking destination ports.
Receive UDP port number <code>UDP_RX_DEST_PORT_NO</code>	Local UDP port listening for incoming UDP frames. Receive and transmit UDP streams can use identical or different ports at the user's discretion.
Client TCP destination IP address <code>DEST_IPv4_ADDR(I)</code>	For each TCP client, define a remote server IP address to connect to. It should be set prior to triggering a connection.
Client TCP destination port <code>DEST_PORT(I)</code>	For each TCP client, define a remote server TCP port to connect to. It should be set prior to triggering a connection.
Client TCP connection control <code>STATE_REQUESTED(I)</code>	For each TCP client, set to 1 to request a connection to a remote server, 0 to clear any such connection.

Limitations

This software does not support the following:

- IEEE 802.3/802.2 encapsulation, RFC 1042, only the most common Ethernet encapsulation.

Only one gateway is supported at any given time.

Software Licensing

The COM-5403SOFT is supplied under the following key licensing terms:

1. A nonexclusive, nontransferable license to use the VHDL source code internally, and
2. An unlimited, royalty-free, nonexclusive transferable license to make and use products incorporating the licensed materials, solely in bit stream format, on a worldwide basis.

The complete VHDL/IP Software License Agreement can be downloaded from <http://www.comblock.com/download/softwarelicense.pdf>

Configuration Management

The current software revision is 002r.

Directory	Contents
/doc	Specifications, user manual, implementation documents when applicable.
/src	.vhd source code, .constraint file, .pkg packages. One component per file.
/sim	Testbenches. Please note that end-to-end TCP simulation requires both TCP client and TCP server code (such as COM-5402SOFT)
/project1	Xilinx Vivado 2017 project Xilinx Vivado 2019 project
/use_example	Examples of interface with the Ethernet MAC (COM-5401SOFT and Xilinx TEMAC)

VHDL development environment

The VHDL software was developed using the following development environment:

- (a) Xilinx Vivado 2019.2 for synthesis, place and route and VHDL simulation
- (b) Xilinx ISE 14.7 for synthesis, place and route

The entire project fits easily within a Xilinx Artix7-100T. Therefore, the ISE project can be processed using the free Xilinx WebPack tools.

Ready-to-use Hardware

Use examples are available to run on the following Comblock hardware modules:

- COM-1800 FPGA + GbE LAN + DDR3 SODIMM SOCKET + ARM + NAND development platform
<http://www.comblock.com/com1800.html>

All hardware schematics are available at <https://comblock.com/download.html>

Top-Level VHDL hierarchy

```
COM5403(Behavioral) (com5403.vhd) (16)
├── Inst_TIMER_4US : TIMER_4US(Behavioral) (tim
├── Inst_PACKET_PARSING : PACKET_PARSING(B
├── Inst_ARP : ARP(Behavioral) (arp.vhd)
├── Inst_PING : PING(Behavioral) (ping.vhd) (1)
│   ├── BRAM_DP2_001 : BRAM_DP2(Behavioral) (
│   ├── WHOIS2_X.WHOIS2_001 : WHOIS2(Behavioral)
│   ├── ARP_CACHE2_X.ARP_CACHE2_001 : ARP_CA
│   ├── DHCP_CLIENT_002.Inst_DHCP_CLIENT : DHC
│       ├── UDP_RX_001 : UDP_RX(Behavioral) (udp_r
│       > ── UDP_TX_001 : UDP_TX(Behavioral) (udp_tx
│   ├── IGMP_QUERY_001 : IGMP_QUERY(Behavioral)
│   ├── IGMP_REPORT_001 : IGMP_REPORT(Behavior
│   ├── Inst_UDP2SERIAL : UDP2SERIAL(Behavioral) (
│   ├── UDP_RX_X.UDP_RX_001 : UDP_RX(Behavioral)
│   ├── UDP_TX_NZ.UDP_TX_001 : UDP_TX(Behavioral)
│   ├── TCP_CLIENTS_X.TCP_CLIENTS_001 : TCP_CI
│   ├── TCP_CLIENTS_X.Inst_TCP_TX : TCP_TX(Behav
│   ├── TCP_CLIENTS_X.Inst_TCP_RXBUFNDMUX2 :
│       ├── BRAM_DP2_X[0].BRAM_DP2_001 : BRAM_I
│   ├── TCP_CLIENTS_X.Inst_TCP_TXBUF : TCP_TXBL
│       ├── OB_001 : BRAM_DP2(Behavioral) (bram_dp
│       ├── BRAM_DP2_X[0].BRAM_DP2_001 : BRAM_I
```

The code is stored with one, and only one, component per file.

The root entity (highlighted above) is *COM5403.vhd*. It contains instantiations of the IP protocols and a transmit arbitration mechanism to select the next packet to send to the MAC/PHY.

The root also includes the following components:

- The *PACKET_PARSING.vhd* component parses the received packets from the MAC and efficiently extracts key information relevant for multiple protocols. Parsing is done on the fly without storing data. Instantiated once.
- The *ARP.vhd* component detects ARP requests and assembles an ARP response Ethernet packet for transmission to the MAC. Instantiated once.
- The *PING.vhd* component detects ICMP echo (ping) requests and assembles a ping echo Ethernet packet for transmission to the MAC. Instantiated once.
- The *WHOIS2.vhd* component generates an ARP request (broadcast) packet requesting that the target identified by its IP address responds with its MAC address. Instantiated once. Required only if *ARP_CACHE2.vhd* is instantiated.
- The *ARP_CACHE2.vhd* component is a shared routing table that stores up to 128 IP addresses with their associated 48-bit MAC addresses and a 'freshness' timestamp. An arbitration circuit is used to arbitrate the routing request from multiple transmit instances. Instantiated once. Required if either *TCP_CLIENTS.vhd* or *UDP_TX.vhd* components are enabled.
- The *DHCP_CLIENT.vhd* component asks DHCP servers for an IP address and network information such as gateway, subnet mask and DNS address. Limitations: Client and server are on the same subnet (no relay)
- The *IGMP_QUERY.vhd* component detects a valid IGMP membership query and triggers a response when applicable.
- The *IGMP_REPORT.vhd* component sends an IGMP membership report out to whom it may concern.
- The flexible *UDP_TX.vhd* component encapsulates a data packet into a UDP frame addressed from any port to any port/IP destination. Instantiated once, irrespective of the number of source or destination UDP ports.
- The *UDP_RX.vhd* component validates received UDP frames and extracts the data packet within. As the validation is performed on the fly (no storage) while received data is passing through, the validity confirmation is made available at the end of the packet. The calling application should therefore be able to 'backtrack' upon receiving an invalid packet. Instantiated once, irrespective of the number of UDP ports being listened to. Although this component is written for one port, it can very easily be modified to accommodate several ports (follow the PORT_NO signal). Therefore, there is never any need to instantiate more than one component.
- The *TCP_CLIENTS.vhd* component is the heart of the TCP protocol. It is written parametrically so as to support NTCPSTREAMS concurrent TCP connections. It essentially handles the TCP state machines of several TCP clients: user-initiated connection request to user-specified remote server IP/Port address, establishing and tearing down the connections and managing flow control while the bi-directional connections are established.
- The *TCP_TX.vhd* component formats TCP tx frames, including all layers: TCP, IP, MAC/Ethernet. It is common to all concurrent streams.
- The *TCP_TXBUF.vhd* component stores TCP tx payload data in individual elastic buffers, one for each transmit stream. The buffer size is configurable prior to synthesis as NBUFS*16Kbits RAM blocks.
- The *TCP_RXBUFNDEMUX2.vhd* component demultiplexes several TCP rx streams. It does not include any elastic buffer. It is expected that the user will instantiate elastic buffers if the application requires it. Data bytes are received in sequence without gaps or backtracking.

Additional components are also provided for use during system integration or tests.

- *STREAM_2_PACKETS.vhd* segments a continuous data stream into packets. The

transmission is triggered by either the maximum packet size or a timeout waiting for fresh stream bytes. Useful when sending a continuous data stream over UDP.

- *PACKETS_2_STREAM.vhd* reassembles a data stream from received valid packets while discarding invalid packets. The packet's validity is assessed at the end of packet. It is designed to connect seamlessly with the *TCP_RX.vhd* component. Useful when receiving a continuous data stream over UDP.
- *LFSR11P.vhd* generates a pseudo-random binary stream PRBS11 for use during throughput and bit error rate tests. It is capable of generating 1 Gbps (8 bit per clock @ 125 MHz).
- *BER2.vhd* synchronizes with a received data stream and counts bit errors. It is also capable of working at 1 Gbps.

VHDL simulation

A testbench (tb*.vhd), located in the /sim directory, can be used to validate the source code through VHDL simulation. However, because of the interactive nature of the TCP and DHCP protocols, other IP cores are needed for such simulations:

- TCP end-to-end simulation is only possible when a TCP server (COM-5402SOFT or equivalent) are present.
- Simulation of DHCP dynamic assignment of IP addresses is only possible when a DHCP server (COM-5404SOFT or COM-5402SOFT or equivalent) are present.

Use examples

An example of COM-5403SOFT interface with the Xilinx tri-mode Ethernet MAC is in */use_example/COM1800_TEMAC_IF.vhd*.

Clock / Timing

The software uses one synchronous clock CLK. The clock should be at least 125 MHz in order to take full advantage of the Gbit Ethernet speed. The code can operate properly at less than 125 MHz, albeit at reduced throughput.

The code is written to run at 125 MHz on the following targets:

- Xilinx Spartan-6 –2 speed grade with 2 concurrent TCP streams instantiated.
- Any Xilinx 7-series FPGA/SoC, even at the lowest speed grade.

ComBlock Compatibility List

FPGA development platform
COM-1800 FPGA (Xilinx Artix7-100T) + DDR3 SODIMM + GbE LAN development platform
Software
COM-5401SOFT Tri-mode 10/100/1000 Mbps Ethernet MAC. VHDL source code.
Xilinx TEMAC, rev 9.0 or above
COM-5402SOFT IP/UDP/TCP SERVER/ARP/PING stack. VHDL source code.

Acronyms

Acronym	Definition
BER	Bit Error Rate
BRAM	Dual-port Block RAM
CTS	Clear To Send, a flow-control signal allowing the data source to send data.
IP	Internet Protocol
LSb	Least Significant bit
MSb	Most Significant bit
PRBS-11	Pseudo-Random Binary Sequence, 2047-bit period
Rx	Receive
TCP	Transmission Control Protocol
Tx	Transmit
UDP	User Datagram Protocol

ComBlock Ordering Information

COM-5403SOFT IP/TCP CLIENT/
UDP/ARP/PING PROTOCOL STACK,
VHDL SOURCE CODE

ECCN: EAR99

MSS • 845 Quince Orchard Boulevard Ste N •
Gaithersburg, Maryland 20878-1676 • U.S.A.
Telephone: (240) 631-1111
Facsimile: (240) 631-1676
E-mail: sales@comblock.com