


The Compact Muon Solenoid Experiment

# Conference Report

Mailing address: CMS CERN, CH-1211 GENEVA 23, Switzerland



22 October 2014

## IPbus A flexible Ethernet-based control system for xTCA hardware

Tom Williams for the CMS Collaboration

### Abstract

The ATCA and uTCA standards include industry-standard data pathway technologies such as Gigabit Ethernet which can be used for control communication, but no specific hardware control protocol is defined. The IPbus suite of software and firmware implements a reliable high-performance control link for particle physics electronics, and has successfully replaced VME control in several large projects. In this paper, we outline the IPbus system architecture, and describe recent developments in the reliability, scalability and performance of IPbus systems, carried out in preparation for deployment of uTCA-based CMS upgrades before the LHC 2015 run. We also discuss plans for future development of the IPbus suite. SUMMARY IPbus will be used for controlling the uTCA electronics in the CMS HCAL, TCDS, Pixel and Level-1 trigger upgrades. IPbus control has already been extensively used in the work of these upgrade projects so far, and final uTCA systems will be deployed in the experiment starting from Autumn 2014. IPbus is also being evaluated for use in the ATLAS and ALICE upgrades, as well as other particle physics experiments. A tightly-integrated suite of software and firmware components has been developed to implement the IPbus protocol the firmware core, a reference VHDL implementation of an IPbus server over UDP, decoding IPbus read/write requests within end-user hardware; uHAL, the C++/Python library providing an end-user API for IPbus reads and writes; and the ControlHub, a software application which arbitrates hardware access to each board from multiple clients. Over the past two years we have developed a new reliable, higher-throughput version of the IPbus protocol, firmware and software. We have set up an IPbus test system with realistic network topology in the CMS electronics integration centre, in order to validate the reliability and performance of the IPbus control system. The software has been optimised to increase the block write/read throughput towards the Gigabit Ethernet bandwidth, and to improve the scalability with the number of targets handled by each ControlHub instance. For 1 client and 1 target, the latency is about 250us for sequences of up to tens of transactions and the maximum block read/write throughput is 0.54Gbit/s; the throughput increases to 0.8Gbit/s for 3 or more targets. We have accumulated weeks of continuous high-throughput random writes and reads over IPbus, without any errors. We also investigated scenarios with network congestion in the MCH Ethernet switch for a full uTCA crate, and found that with appropriate configuration this congestion only has a small effect on the IPbus throughput (12pct reduction). Plans for future work include improving the monitoring of IPbus dataflows in large systems of hundreds of targets, and investigating further ideas for usability and performance improvements.

## 2 IPbus: A flexible Ethernet-based control system for 3 xTCA hardware

---

**C. Ghabrous Larrea<sup>a</sup>, K. Harder<sup>b</sup>, D. Newbold<sup>c</sup>, D. Sankey<sup>b</sup>, A. Rose<sup>d</sup>, A. Thea<sup>b</sup>, and  
T. Williams<sup>b\*</sup>**

*<sup>a</sup>Dept. of Physics, University of Wisconsin-Madison,  
1150 University Ave., Madison, WI 53706, U.S.A.*

*<sup>b</sup>Rutherford Appleton Laboratory, STFC,  
Harwell, OX11 0QX, U.K.*

*<sup>c</sup>H.H. Wills Physics Laboratory, University of Bristol,  
Tyndall Avenue, Bristol, BS8 1TL, U.K.*

*<sup>d</sup>Blackett Laboratory, Imperial College,  
Prince Consort Road, London, SW7 2BW, U.K.*

4 *E-mail: t.williams@cern.ch*

5 **ABSTRACT:** The ATCA and  $\mu$ TCA standards include industry-standard data pathway technologies such as Gigabit Ethernet which can be used for control communication, but no specific hardware control protocol is defined. The IPbus suite of software and firmware implements a reliable high-performance control link for particle physics electronics, and has successfully replaced VME control in several large projects. In this paper, we outline the IPbus control system architecture, and describe recent developments in the reliability, scalability and performance of IPbus systems, carried out in preparation for deployment of  $\mu$ TCA-based CMS upgrades before the LHC 2015 run. We also discuss plans for future development of the IPbus suite.

6 **KEYWORDS:** Control system; IPbus;  $\mu$ TCA; MicroTCA; ATCA.

---

\*Corresponding author.

8 **Contents**

9	<b>1. Introduction</b>	<b>1</b>
10	<b>2. IPbus protocol</b>	<b>2</b>
11	<b>3. Firmware and software suite</b>	<b>2</b>
12	3.1 IPbus firmware	3
13	3.2 ControlHub	3
14	3.3 $\mu$ HAL library	4
15	<b>4. Control system topology</b>	<b>4</b>
16	<b>5. System reliability</b>	<b>6</b>
17	<b>6. Performance</b>	<b>6</b>
18	<b>7. Conclusions</b>	<b>8</b>

20 **1. Introduction**

21 New electronics systems within many particle physics experiments are based on the ATCA and  
22  $\mu$ TCA standards (henceforth collectively referred to as xTCA). The xTCA specifications incor-  
23 porate industry-standard serial communication technologies such as Gigabit Ethernet; however,  
24 unlike the VMEbus standard, they do not specify a hardware access protocol for controlling xTCA  
25 boards from external software applications.

26 Several important requirements must be considered when designing the architecture and im-  
27 plementation of a hardware control system. Control systems must have reliable and predictable  
28 behaviour under all conditions, since they form the main link by which hardware is configured,  
29 monitored, and debugged in case of problems. The control system architecture for large exper-  
30 iments should be highly scalable, ideally with the same ease of setup and use from the simple  
31 ‘board on benchtop’ scenario to the final system with hundreds of boards. In modern particle  
32 physics experiments, the same electronics setup is often used for decades before being replaced,  
33 and the associated control infrastructure must have the same maintainable lifetime. Hence, it is  
34 typically beneficial to use widespread industry-standard technologies, in order to avoid the risk of  
35 reliance on a single vendor. Experience from the CMS experiment’s online systems in LHC Run  
36 1 also shows that for monitoring and debugging issues in complex scenarios, in general it is help-  
37 ful to move complexity away from hardware/firmware into software running on commercial PC  
38 hardware.

39 The IPbus protocol — first developed by J. Mans et al. in 2009 — is a simple control pro-  
40 tocol for reading and modifying registers within IP-aware hardware. A tightly-integrated suite of  
41 IPbus software and firmware components which can be used to construct reliable, scalable, high-  
42 performance control systems has previously been presented in Ref. [1]. This IPbus suite will be  
43 used to control the xTCA off-detector electronics in the Phase-1 upgrades to the CMS experi-  
44 ment [2], as well as the ATLAS experiment’s Phase-0 and Phase-1 upgrades [3]. In this paper, we  
45 present recent improvements in the reliability, scalability and performance of the IPbus suite, based  
46 on a new version of the protocol.

## 47 2. IPbus protocol

48 The IPbus protocol is a simple protocol for controlling IP-aware hardware devices which have a  
49 virtual A32/D32 bus. It defines the following operations:

50 **Read** A read of user-definable depth. Two types of read are defined: incrementing (for multiple  
51 continuous registers in the IPbus address space) and non-incrementing (for a port or FIFO).

52 **Write** A write of user-definable depth. As with reads, two types of write are defined: incrementing  
53 and non-incrementing.

54 **Read-Modify-Write bits (RMWbits)** An atomic bit-masked write, defined as  $X := (X \& A) | B$ .  
55 This allows one to efficiently set/clear a subset of bits within a 32-bit register.

56 **Read-Modify-Write sum (RMWsum)** An atomic increment operation, defined as  $X := X + A$ ,  
57 which is useful for adding values to a register (or subtracting, using two’s complement).

58 The IPbus protocol lies in the application layer of the networking model and is transport protocol  
59 agnostic. Each IPbus host device (typically hardware in a remote electronics crate) has an IP  
60 address and port number on which it accepts IPbus control packets. The protocol is transactional  
61 — for each read, write or RMW operation, the IPbus client (typically software) sends a request to  
62 the IPbus device; the device then sends back a response message containing an error code (equal  
63 to 0 for a successful transaction), followed by return data in case of reads. In order to minimise  
64 latency, multiple transactions can be concatenated into a single IPbus packet.

65 Version 2.0 of the IPbus protocol [4] (finalised in early 2013) includes a reliability mechanism,  
66 through which the IPbus client can correct for any packet loss, duplication or re-ordering, if using  
67 an unreliable transport such as UDP. This mechanism is based on the client setting sequential packet  
68 ID values. In systems with multiple control applications, IPbus traffic must be routed via a network  
69 element that understands the IPbus protocol and thus can buffer the incoming request packets and  
70 reset their IDs (in practice, this is the role of the ControlHub).

## 71 3. Firmware and software suite

72 The IPbus software and firmware suite consists of the following components:

73 **IPbus firmware** A module that implements the IPbus protocol within end-user hardware

**Table 1.** Resource usage of IPbus firmware core.

Resource	Usage	
	Minimal configuration	Fully-featured
Flip flops	2000	3500
Slices	1000	2900
Block RAMs	5	17

74 **ControlHub** Software application that mediates simultaneous hardware access from multiple  $\mu$ HAL  
75 clients, and implements the IPbus reliability mechanism over UDP

76  **$\mu$ HAL** C++ and Python end-user programming interface for writes, reads and RMW operations

77 End-user instructions and source code for these components are available through the CERN CAC-  
78 TUS (Code Archive for CMS Trigger UpgradeS) website and SVN repository [5]. The software is  
79 packaged as RPMs for Scientific Linux versions 5 and 6, and available through a YUM repository.

### 80 3.1 IPbus firmware

81 The IPbus 2.0 firmware module is a reference system-on-a-chip implementation of an IPbus 2.0  
82 UDP server in VHDL; it interprets IPbus transactions on an FPGA. It has been designed as a  
83 common module to run alongside a device’s main processing logic (e.g. trigger algorithms) on the  
84 same FPGA, only using resources from within the FPGA. As a result of this, the IPbus firmware  
85 core must have a low resource usage, which is an important consideration in the choice of transport  
86 protocol. The TCP protocol exhibits various highly-desirable features of a transport protocol, such  
87 as reliable, ordered data transmission and congestion avoidance; however, the underlying algorithm  
88 is significantly more complex than for the other ubiquitous transport layer protocol, UDP. Hence,  
89 UDP has been chosen as the transport protocol; any loss, re-ordering or duplication of the IPbus  
90 UDP packets is automatically corrected by the ControlHub using the IPbus reliability mechanism.

91 The IPbus firmware module has been designed to be simple to integrate into variety of plat-  
92 forms, and there are example designs for several development boards and standard platforms. The  
93 source code is currently Xilinx-specific, but has been successfully adapted for Altera devices. In  
94 addition to UDP, the IPbus firmware module also implements: the echo request/reply semantics  
95 from ICMP (RFC 792, used in the Unix `ping` command); ARP (RFC 826, used for resolving  
96 IP addresses into MAC addresses); and RARP (RFC 903, used for requesting an IP address on  
97 startup). Several parameters are configurable at build time, including: the Ethernet frame MTU;  
98 the number of buffers for incoming/outgoing IPbus packets which determines the maximum pos-  
99 sible control throughput; and the method used for IP address assignment — RARP, IPMI, or fixed  
100 IP address. The resource usage of the IPbus firmware core under ‘minimal’ and ‘fully-featured’  
101 configurations is shown in table 1.

### 102 3.2 ControlHub

103 The ControlHub is a software application that forms a single point of access for IPbus control of  
104 each device; specifically, it arbitrates simultaneous access from multiple control applications to  
105 one or more devices, and it implements the IPbus reliability mechanism for the ControlHub–device

106 UDP packets. Since the ControlHub is a software application, the  $\mu$ HAL–ControlHub communi-  
107 cation uses TCP, which has sophisticated congestion mitigation and flow-control algorithms.

108 **Design requirements and implementation.** The ControlHub must be at least as reliable and  
109 transparent as a VME crate controller, since failure or crash within the ControlHub could disrupt  
110 the communications of several upstream control or monitoring applications. Additionally its design  
111 must allow multiple clients to communicate with multiple targets reliably, efficiently and indepen-  
112 dently.

113 Erlang is a general-purpose, concurrent programming language, designed by Ericsson to build  
114 high-availability, fault-tolerant applications. The main structural unit in Erlang is the process.  
115 Erlang processes are lightweight compared to operating system processes; they share no state,  
116 instead communicating by message passing. These features are well-suited to the ControlHub’s  
117 requirements for high reliability, performance, and scalability in routing IPbus transactions, and  
118 therefore the ControlHub is implemented in Erlang. The ControlHub uses a separate Erlang process  
119 for each connected  $\mu$ HAL client and each IPbus device, ensuring workload can be spread across  
120 multiple CPU cores; its internal structure is described in more detail in Ref. [1].

### 121 3.3 $\mu$ HAL library

122  $\mu$ HAL is the Hardware Access Library (HAL) providing an end-user C++/Python API for IPbus  
123 reads, writes and RMW transactions. It is based on a delayed dispatch model in which multi-  
124 ple transactions are queued and concatenated within the transport layer payload buffers until the  
125 `dispatch` method is called.

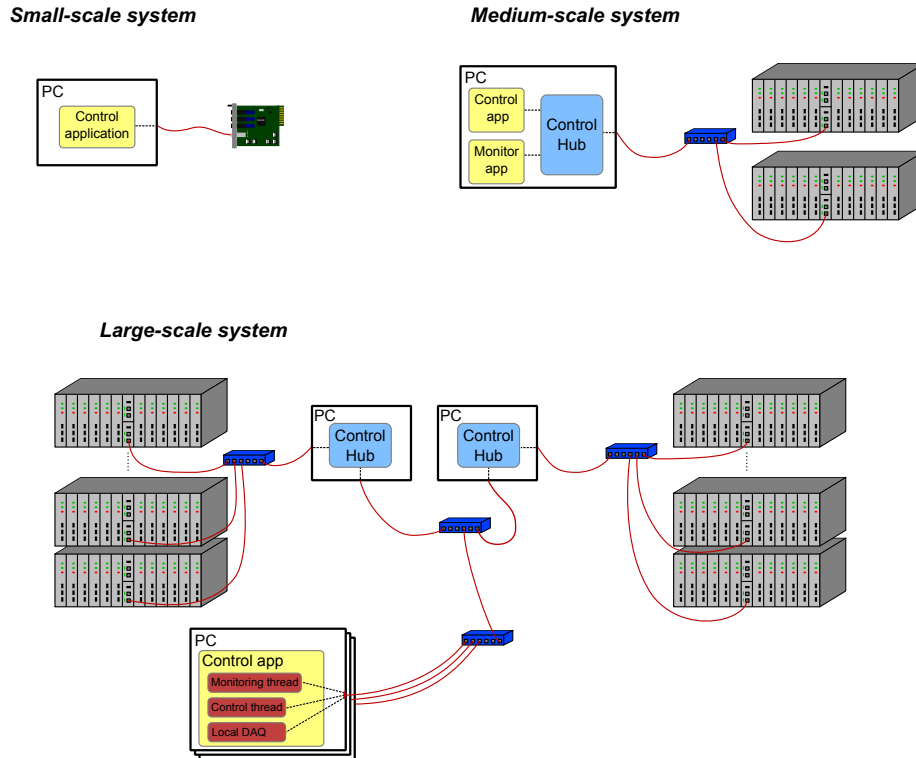
126 In  $\mu$ HAL each device’s register layout is specified by XML files. Each node of the XML tree  
127 represents either a single register, block RAM, FIFO, or a collection of these; the nodes in one  
128 file can reference other address files, such that the interfaces to repeated instances of a firmware  
129 module can be generated with minimal copy-paste of address file contents. This enables the user  
130 to write control software in a manner that intuitively mirrors the modular, hierarchical structure of  
131 large firmware designs.

132 The  $\mu$ HAL interface to each device (based on the methods of the `HwInterface` and `Node`  
133 classes) can run in one of two modes of operation. In the local-client mode, the  $\mu$ HAL library  
134 communicates directly with device over UDP. In the remote-client mode, the  $\mu$ HAL library com-  
135 municates with hardware exclusively via a ControlHub. These differing modes of operation are  
136 implemented through the inheritance of of common interface, such that users can switch be-  
137 tween the modes of operation by simply changing the prefix of a single string when creating a  
138 `HwInterface` instance.

139  $\mu$ HAL is also packaged with an example GUI that is useful for monitoring the values of a  
140 subset of registers on a device during hardware development.

## 141 4. Control system topology

142 The topologies of an IPbus control system in some common scenarios are shown in figure 1. The  
143 simplest system (*upper left*) is a single target running the IPbus firmware, directly connected by a



**Figure 1.** Example topologies of IPbus control systems involving  $\mu$ TCA hardware, from small to large scale.

144 single Ethernet cable to a computer running a C++/Python control application based on the  $\mu$ HAL  
 145 library. This is the typical layout during early hardware development.

146 In a more complex scenario such as a beam test or integration tests, there will typically be  
 147 several devices, with multiple control, monitoring and DAQ applications, as shown in figure 1 (*up-*  
 148 *per right*). Due to multiple applications simultaneously communicating with the devices, the IPbus  
 149 traffic would be routed via a ControlHub, which would also recover any lost packets making the  
 150 IPbus communication 100 % reliable.

151 For a full-scale IPbus system at a large experiment (such as ATLAS or CMS) there would  
 152 be hundreds of IPbus devices spread across many crates, and the control/monitoring applications  
 153 would be spread across many computers, as shown in figure 1 (*lower*). In this case the use of an  
 154 Ethernet network naturally allows scalability with the ease of extending the network using multiple  
 155 switches and routers. Additionally the recovery from computer failure is simplified with the pos-  
 156 sibility of having spare computers already connected to the network. Notably, the network could  
 157 be divided into a separate subnet for each subdetector so that the network's logical segmentation  
 158 matches the typical IPbus dataflow. The exact number of devices per ControlHub would be adapted  
 159 based on performance requirements.

160 **IPbus test system.** A test system was set up in the CMS electronics integration centre at CERN,  
 161 in order to investigate the reliability and performance of the IPbus suite using very similar network  
 162 layout and hardware to that planned for final deployment in the CMS experiment. The test sys-

163 tem consists of network infrastructure, two computers, and one  $\mu$ TCA shelf containing 12  $\mu$ TCA  
164 boards (AMCs), each running the IPbus 2.0 firmware core. The computers are Dell PowerEdge  
165 R300 rack PCs; three of the AMCs are GLIBs [6] and the other nine are Mini-T5s [7].

## 166 5. System reliability

167 The reliability and robustness of the IPbus suite has been ensured by extensive testing of both the  
168 software and firmware in a range of scenarios.

169 The software is tested by itself (independent of the hardware) each night using a *dummy hard-*  
170 *ware* executable which emulates the response of an IPbus device. A suite of unit test executables  
171 are run in order to test  $\mu$ HAL and the ControlHub with basic read/write/RMW operations to the  
172 dummy hardware running on the same machine. By configuring the operating system to randomly  
173 drop IP packets, these executables are also used to test the ControlHub's reliability mechanism.

174 The full IPbus control link ( $\mu$ HAL–ControlHub–firmware) has been tested with a variety of  
175  $\mu$ TCA boards, using a  $\mu$ HAL-based C++ executable. This executable issues random sequences of  
176 reads, writes and RMW transactions to a device using random addresses, random depths for the  
177 reads and writes, and random values for the data written and the RMW parameters. The executable  
178 checks that all of the returned error codes indicate success, and checks that the values returned  
179 by the reads and RMW transactions are always correct. The released version of the firmware  
180 core was validated by running the executable for over 20 hours (corresponding to over 10 billion  
181 transactions) against the IPbus firmware core loaded on each of the Mini-T5, GLIB and MP7  
182 boards. No errors were observed during this final testing.

## 183 6. Performance

184 The latency and block transfer throughput are two important parameters of a control system:


185 **Latency** is defined as the total round-trip time taken to perform an IPbus transaction, as measured  
186 in the  $\mu$ HAL client application.

187 **Throughput** is defined as the amount of user data transferred or received per unit of time.

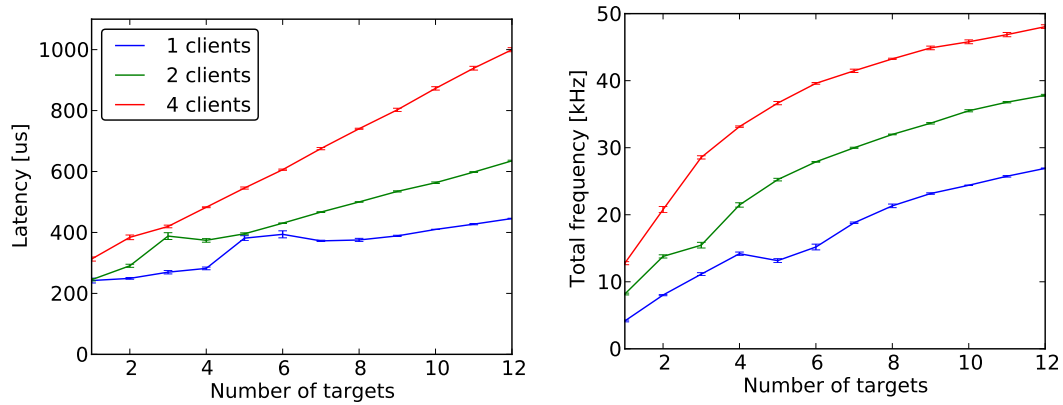
188 In order to predict the performance of the future CMS IPbus control system, and verify the design of  
189 the IPbus components and their planned layout, the system performance was measured in several  
190 benchmark scenarios. These measurements were carried out in the IPbus test system, with the  
191  $\mu$ HAL clients running on one computer and the ControlHub on the other computer.

192 **1-to-1 block transfers.** The block read/write latency and throughput for one  $\mu$ HAL client con-  
193 trolling one device via the ControlHub, is shown in figure 2. The median single-word write/read  
194 latency is approximately 250  $\mu$ s. Although this single-word latency is significantly larger than  
195 with VME/PCIe-based control, for multiple transactions or large block transfers this is compen-  
196 sated by concatenating multiple transactions into each packet, and by having multiple packets in  
197 flight around the system at any given time. Hence, the block read/write throughput for payloads  
198 larger than 1 MByte is above 0.5 Gbit/s.





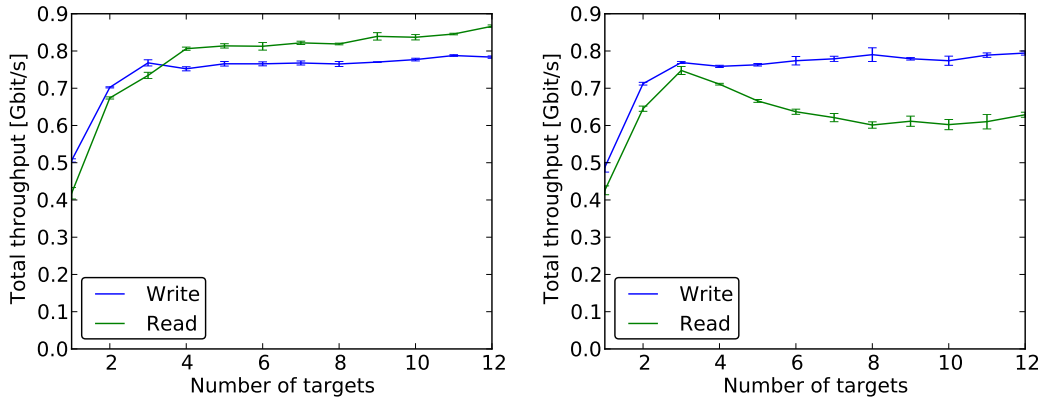
**Figure 2.** The median write/read and throughput as a function of depth, for one software client controlling one IPbus device, via the ControlHub.



**Figure 3.** The latency and total system polling frequency for  $n$  clients each simultaneously polling a register in one of the  $m$  targets.

199 **n-to-m polling.** The system performance for multiple  $\mu$ HAL clients polling a single-word reg-  
 200 ister in multiple devices is shown in via one ControlHub was also measured. The mean polling  
 201 latency, and total system polling frequency, for 1, 2 or 4 clients per device are shown in figure 3  
 202 as a function of the number of devices. The latency experienced by each client gradually increases  
 203 with the number of clients or devices, due to the the increasing load of network interrupts on the  
 204 computers. However, the total polling frequency increases with the number of clients or devices in  
 205 the system, as the ControlHub spreads its increasing workload over the four CPU cores.

206 **n-to-n block transfers.** The performance for continuous block reads and writes of all 12 boards  
 207 in the  $\mu$ TCA crate was also measured. The Ethernet connection to a  $\mu$ TCA crate is via a Gigabit  
 208 Ethernet socket on the front panel of the crate management module, the MCH (MicroTCA Carrier  
 209 Hub). Each individual AMC in a  $\mu$ TCA crate is connected to the MCH’s Ethernet switch by  
 210 a separate bidirectional 1 Gbit/s link. In theory, this network topology could lead to congestion  
 211 in the MCH switch during simultaneous block reads from multiple AMCs. For block reads, the  
 212 reply packets are significantly larger than the request packets, and so the total instantaneous return  
 213 bandwidth from the 12 AMCs into the MCH could exceed the 1 Gbit/s capacity of the link from



**Figure 4.** The total system throughput for  $n$  IPbus clients each simultaneously writing to / reading from one of  $n$  devices, via one ControlHub, using a NAT MCH (*left*) or a Vadatech MCH (*right*).

214 the MCH to the local network. However, within the IPbus protocol only a limited number of  
 215 requests are in flight to each target at any given time, which imposes an upper limit on the total  
 216 size of packets that would have to be buffered within the MCH switch. In practice whether or not  
 217 such congestion leads to reduced performance depends on various factors, including the number of  
 218 packets in flight to each AMC, and the design of the MCH switch. Within the CMS collaboration,  
 219 MCH modules are currently being purchased from two vendors: NAT and Vadatech.

220 The IPbus system throughput for multi-client block reads and writes with multiple targets are  
 221 shown in figure 4 for both the NAT and Vadatech MCHs. For the NAT MCH (V3.4), the read and  
 222 write throughputs are similar; over 75 % of the Gigabit Ethernet bandwidth is utilised with three  
 223 or more devices. However, using the Vadatech MCH (model UTC002-210-440-010), the system  
 224 throughput degrades for simultaneous block reads from four or more devices due to congestion in  
 225 the MCH switch, with read throughput approximately 20 % lower than write throughput for 8 or  
 226 more targets. In order to reduce congestion, the system performance was re-measured with fewer  
 227 packets in flight to each device; this can be achieved by editing one line in the ControlHub config-  
 228 uration file. With 11 packets in flight to each device (default value is 16), there is less congestion-  
 229 induced packet loss, and so the simultaneous read throughput is above 0.75 Gbit/s for three or more  
 230 devices; however, the maximum 1-client-to-1-target throughput decreases by approximately 12 %.

## 231 7. Conclusions

232 A new reliable, high-performance version of the IPbus protocol has been developed along with the  
 233 associated suite of software and firmware, in order to control xTCA hardware via Gigabit Ethernet.  
 234 An IPbus test system with realistic network topology was set up in the CMS electronics integration  
 235 centre in order to verify the control system's reliability, and investigate its performance. For one  
 236 software client controlling one device, the single-word read/write latency is approximately 250  $\mu$ s  
 237 and the block read/write throughput is above 0.5 Gbit/s for payloads larger than 1 MByte; the total  
 238 block read/write throughput is above 0.75 Gbit/s for three or more boards in a single  $\mu$ TCA shelf.

239 The first large-scale IPbus system in the CMS experiment was deployed in August 2014, in  
 240 preparation for the start of LHC Run 2 in 2015. Hence, development is now focused on simplifying

241 the monitoring of IPbus dataflows in large systems of hundreds of devices. The IPbus software  
242 and firmware suite will be optimised in order to improve performance with 10 Gigabit Ethernet.  
243 Additionally, an IPbus locking mechanism is being considered in order to provide exclusive access  
244 to IPbus devices from a single client for extended configuration sequences.

## 245 **Acknowledgments**

246 Acknowledgments.

## 247 **References**

- 248 [1] R. Frazier, G. Iles, D. Newbold, and A. Rose, *Software and firmware for controlling CMS trigger and*  
249 *readout hardware via gigabit Ethernet*, *Physics Procedia* **37** (2012) 1892
- 250 [2] The CMS collaboration, *Technical proposal for the upgrade of the CMS detector through 2020*, CERN,  
251 Geneva 2011. CERN-LHCC-2011-006.
- 252 [3] The ATLAS collaboration, *Letter of Intent for the Phase-I Upgrade of the ATLAS Experiment*, CERN,  
253 Geneva 2011. CERN-LHCC-2011-012.
- 254 [4] R. Frazier, G. Iles, M. Magrans de Abril, D. Newbold, A. Rose, D. Sankey, and T. Williams, *The IPbus*  
255 *protocol: version 2.0*, [https://svnweb.cern.ch/trac/cactus/browser/trunk/doc/ipbus\\_protocol\\_v2\\_0.pdf](https://svnweb.cern.ch/trac/cactus/browser/trunk/doc/ipbus_protocol_v2_0.pdf)
- 256 [5] The CMS Level-1 trigger project, *The CACTUS (Code Archive for CMS Trigger UpgradeS) SVN*  
257 *repository*, <http://cactus.web.cern.ch/>
- 258 [6] P. Vichoudis et al, *The Gigabit Link Interface Board (GLIB) ecosystem*, 2013 *JINST* **8** C03012.
- 259 [7] C. Foudas, R. Frazier, G. Hall, G. Iles, J. Jones, J. Marrouche, D. Newbold, and A. Rose, *A*  
260 *demonstrator for a level-1 trigger system based on MicroTCA technology and 5Gb/s optical links*,  
261 2010 *JINST* **5** C11015.