

OCP-IP News

Membership Announcements

Defense Science and Technology Organization (DSTO): part of Australia's Department of Defense, DSTO is the Australian Government's lead agency charged with applying science and technology to protect and defend Australia and its national interests



Analog Devices: A leader in data conversion and signal conditioning technology, Analog Devices serves over 60,000 customers, representing virtually all types of electronic equipment.



Celebrating over 40 years as a leading global manufacturer of high-performance

integrated circuits used in analog and digital signal processing applications, Analog Devices is headquartered in Norwood, Massachusetts, with design and manufacturing facilities throughout the world. Analog Devices' common stock is listed on the New York Stock Exchange under the ticker "ADI" and is included in the S&P 500 Index.

OCP-IP Sponsors EDAC's Annual CEO Forecast

•Once again this year OCP-IP had the pleasure of sponsoring the **2010 EDA Consortium CEO Forecast and Industry Vision**. The event was a great success with presentations from: Robert Gardner, EDAC Executive, Jay Vleeschhouwer, Senior Software Analyst, Ticonderoga Securities, and the CEO from each of four market-leading EDA companies

- Lip-Bu Tan, Cadence Design Systems
- Wally Rhines, Mentor Graphics
- John Kibarian, PDF Solutions
- Aart de Geus, Synopsys

The presentations and slides from the event are now available for download.

Watch the presentations: <http://www.edac.org/events10/ceoForecast/presentation/vod.jsp>

Download the slides: <http://www.edac.org/events10/ceoForecast/ceoSlides.jsp>



Lip-Bu Tan, CEO of Cadence, John Kibarian, CEO of PDF Solutions



Walden Rhines, CEO of Mentor Graphics, Aart de Geus, CEO of Synopsys



President's Overview



Ian Mackintosh
OCP-IP President and Chairman

Welcome to this edition of the OCP-IP newsletter. This is an extremely exciting time for OCP-IP as our Technical Vision Working Group just completed its annual strategic session to complete the overall technical planning for the organization for 2010. Our roadmap for 2010 has been set and our Working Groups will now embark on the task of setting target dates and goals for work to be completed in 2010. Several of our working groups have significant deliverables that will become available to our members shortly, and you can learn more about these items throughout this newsletter. Each of these deliverables adds additional significant value to membership in OCP-IP.

I want to take a moment to personally thank all of our working group chairs as well as our participating members for their valued contributions to make these deliverables reality. It is through their dedication, support, and effort that these deliverables are available to be shared between all our members.

As part of OCP-IP's role in industry leadership, we once again sponsored EDA Consortium's 2010 CEO Forecast and Industry Vision. In addition, we have again posted our annual online presentations from Design Automation and Test in Europe. Please take the opportunity to review presentations from members including: CAST, Cadence, HDL Dynamics, Metadata Working Group, MIPS, and Sonics. Presentations can be seen [here](#).


We'd also like to announce that Erno Saliminen of Tampere University of

Technology is the winner of the annual Outstanding Contributor of the Year Award for 2009.

The OCP-IP Governing Steering Committee presents this award each year to a member that makes key contributions to the further advancement of the OCP specification or surrounding infrastructure. The committee acknowledges Tampere University of Technology, and specifically, Erno Salminen for his leadership, commitment and contribution to OCP-IP's Network on Chip Benchmarking Working Group (NoC BWG).

Mr. Salminen played a key role by leading the NoC BWG as its Chair, helping create the 2-part NoC Benchmarking specification and is instrumental in driving the critical task of creating Network on Chip (NoC) benchmarks for the industry through OCP-IP's collaboration with the Embedded Microprocessor Benchmark Consortium (EEMBC). He is also involved in the creation of a TransactionGenerator which will be

available to OCP-IP members through the OCP-IP website later in 2010.



Creating Virtual Platform Using the OCP-IP Modeling Kit

Puneet Arora, Ruchir Bharti, Circuit Sutra, and Mark Burton, Greensocs

A Virtual Platform is created by integrating together software-models of different IPs and is used to simulate the functionality of an entire SoC. As the SoCs get more and more complex, different building blocks (IP models) are quite often sourced from different vendors. When designers will use different interfaces to meet specific requirements and optimizations, it will lead to a major rework activity during integration. The task of connecting software models together can become time consuming and error prone, especially when the individual models use non-standard interfaces.

Thus, providing standard and generic external interfaces to the individual IP models becomes crucial to their reusability and interoperability. The OCP-IP modeling kit provides a comprehensive resource that standardizes the way all IP that uses the OCP-IP bus interconnect is modeled. This greatly simplifies the difficulty in building virtual platforms.

This paper presents our work that has used the OCP-IP modeling kit[1], and shows the benefits that the OCP kit brings. The OCP-IP System Level Design Working Group is currently assessing the feasibility and opportunity to release a version of this work which developed a Virtual Platform based on the OCP modeling kit.

Overview

A Virtual Platform, can be run on a host machine(PC) to emulate the behavior of a different(guest) architecture. It can be used to boot the complete software stack as would the actual physical platform and thus is used for early software development, architecture exploration, hardware-software co-development and verification. Many methodologies and techniques are being

developed in the ESL domain, with the objective of enabling re-use of models which thereby speed up construction of Virtual Platforms.

The essential requirements on the individual building blocks(processor cores, IP peripherals, memories etc) is to provide fast simulation, and to make use of standard interfaces to ease integration; in both cases, overall turn around time for verification is greatly reduced. Both model speed and, importantly, integration times are becoming important and key factors in the design of large systems.

TLM2 technology from OSCI serves as an ideal starting point for the modeling purpose. It defines rules of communication and provides infrastructure to simulate memory-mapped buses.

Since on-chip hardware communication behavior encompasses a great variety, TLM2 is generally extended to support various requirements(for instance various types of transfers, etc). The OCP-IP modeling kit[1] provides a set of rich and configurable features for the interfaces of SystemC modules. This modeling kit is built on top of OSCI's TLM2 technology, adding support for the OCP protocol features. The kit facilitates the creation of models at various abstraction levels (TL1, TL2, TL3, TL4) and supports all the use-cases, including verification, architecture exploration and software development. The TL1 Abstraction corresponds to fully cycle-accurate modeling and TL2 provides for intra-burst timing. The TL3-TL4 are equivalent to the AT-LT methodologies of OSCI base protocol and the

Continued next page

Creating Virtual Platform Using the OCP-IP Modeling Kit (Continued)

APIs are designed to have maximum interoperability with OSCI BP. In this project we created a virtual platform at the TL4 abstraction level, using the TL3/TL4 APIs provided in the OCP Kit.

In our work, we have taken the processor model from QEMU, which is a 'C' based processor emulator. Then SystemC models of IPs like DMA, UART and InterruptController, (which form the guts of both i386 and ARM based platforms) are integrated into the Qemu emulation environment. The interface of these SystemC modules is modeled using the OCP-IP kit.

The models can be accessed by applications running on the virtual platform. Using similar approach, SystemC models of other IPs can also be plugged into the Virtual Platform and easily accessed by the application/device driver running on the system. This demonstrates the use-case of early software development and analysis.

We have also used packages from GreenSocs which provide open source infrastructure for developing standards-based tool-independent SystemC models.

Use of standard modeling kits (OCPIP TL Kit / OSCI TLM2) and opensource infrastructure (GreenSocs, QEMU, Openmoko etc..) to create the platform helps demonstrate the ease in adopting such a platform by both system designers and embedded software developers.

This platform can be used by a SoC company as a basis for creating a virtual platform for its own SoC. It is becoming a standard practice for IP vendors to provide a TLM model of their IP blocks. Our platform can be used by IP vendors to explore the models of their IP blocks and run the software stack / device driver. They can simply create the TLM model of their IP

using the OCP-IP Kit, plug their model into this virtual platform, and access it through the software running on top of the virtual platform.

By now the benefits of creating a virtual platform using SystemC/TLM2.0 have become well known in the industry. However, many companies still seem to struggle to benefit from this new technology. There is often confusion and mistrust: will it really work for their SoC, how much investment will be required etc? This project is an attempt to increase industry awareness and to accelerate the adoption of SystemC, TLM and ESL in general. With this, at least allows for creating the initial proof of concept, no heavy investment is required and no commercial ESL tool is required to be purchased.

Using Qemu and Openmoko

Qemu is a generic and open-source machine emulator and virtualizer. It allows you to run OSes and programs made for one machine(e.g. an openmoko device) on another machine(e.g. your own PC). It is 'C'-based and supports several hardware platforms, including X86, ARM, SPARC and MIPS. To use Qemu, its source code must be configured for a specific target architecture and compiled on a host machine.

For our setup, we used a particular flavor of Qemu called qemu-neo1973[3] which provides support for ARM-based hardware platform for mobile devices.

Openmoko is a hardware and software project, for creating open source mobile phones. Openmoko software is linux, kernel-based opensource OS which can be run on Openmoko hardware platforms.

(To read the rest of the article, go to <http://www.circuitsutra.com/downloads/IPESC09.pdf>)

An Analysis of Blocking versus Non-Blocking Flow Control in On-Chip Networks

By Krishnan Srinivasan, and Drew Wingard Sonics Inc

High end System-on-Chip (SoC) architectures consist of tens of processing engines. These processing engines have varied traffic profiles consisting of priority traffic that require that the latency of the traffic is minimized, controlled bandwidth traffic that require low service jitter on the throughput, and best effort traffic that can tolerate highly variable service. In this paper, we investigate the trade-off between multi-threaded non-blocking (MTNB) flow-control and single threaded tag (STT) based flow-control in the realm of Open Core Protocol (OCP) [1] specifications. Specifically, we argue that the non-blocking multi-threaded flow-control protocol is more suitable for latency minimization of the priority traffic and jitter minimization of controlled bandwidth traffic, when compared with a single threaded tag (STT) based protocol. We present experimental results comparing MTNB against STT based protocols on representative DTV data flows. On average, in the STT based system, the latency of priority traffic is increased by 2.73 times and the latency of controlled bandwidth traffic is increased by 1.14 times when compared to the MTNB system, under identical configurations: CPU Video Back End Video Decoder Blitter/DMA Access Interconnect Memory Subsystem (Memory Scheduler + Memory Controller)

Off Chip DRAM Memory (DDR-1/DDR-2/DDR-3 Single or Multi-threaded OCP)

A System-on-Chip (SoC) consists of several processing engines integrated onto the same chip. The traffic flow from each of the processing engines can have different performance requirements. For example, a CPU whose traffic to the external DRAM is dominated by cacheline-fills would require that the latency of the traffic is minimized. On the other hand, traffic flow from a video decoder engine requires that the traffic is serviced with low jitter, such that the amount of buffering inside the engine can be minimized, while ensuring that the engine never starves for data.

Figure 1: depicts a typical SoC architecture for HDTV systems [2]. For clarity, only the major blocks are included. The traffic flows from the different processing engines (henceforth called Initiators) converge into the DRAM memory. The memory subsystem consists of a memory scheduler, a memory controller, and the off-chip DRAM. The

Continued next page

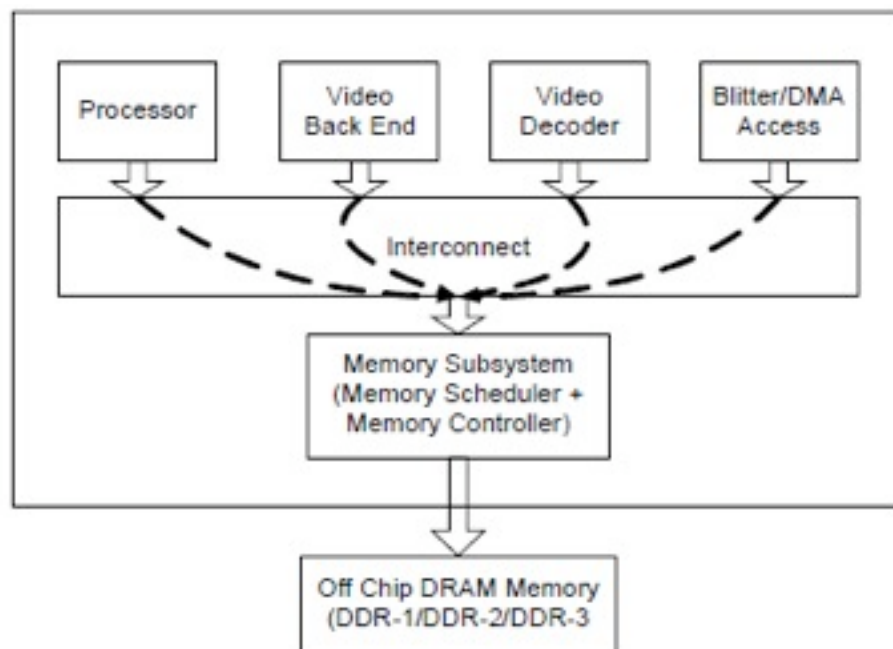


Figure 1: A typical video SoC

An Analysis of Blocking versus Non-Blocking Flow Control (Continued)

memory scheduler arbitrates among multiple traffic flows to generate requests to the controller, which are then issued to the DRAM. Cost and performance requirements of SoCs dictate that the DRAM system operate at high transfer efficiency while minimizing latency for high priority flows and guaranteeing throughput for real-time flow. SoC memory schedulers rely upon transaction re-ordering to optimize these tradeoffs. OCP provides two different out-of-order transaction models: a single threaded tag (STT) based system, or a multi-threaded non-blocking flow-control (MTNB) based system. In the STT system, a single queue is used, in which the insertion is in-order, and the de-queuing operation can be out of order. The memory scheduler looks into the queue and re-orders the requests such that the performance goals are met.

The MTNB system uses threads, which are similar to virtual channels [3]. A MTNB memory subsystem manages per-thread queues. The memory scheduler looks at the head of the queue of each thread, and schedules a request to the memory controller from one of the threads. Therefore, it maintains the order within each thread, while interleaving the requests from different threads.

The choice of implementing an STT-based flow-control protocol or an MTNB-based flow-control protocol is driven by performance requirements and area minimization goals within the SoC. Proponents of STT-based flow-control protocol argue that it is efficient because transaction storage is shared across all initiators and all pending requests are candidates for re-ordering to optimize performance.

An MTNB-based protocol allocates separate queues for each traffic flow. Therefore, the designer can optimize the queue depths based on the distribution of traffic among the threads. Moreover, the MTNB system isolates the performance characteristics of threads from each other. An MTNB scheduler can back-pressure the traffic on over-burdened threads to ensure that well-behaved threads receive their allocated service.

In contrast to the MTNB-based protocol, the STT protocol is highly sensitive to the traffic shaping characteristics at the initiators and the interconnection network. The memory scheduler has only one queue, and has to backpressure the interconnection network whenever this queue is filled. Thus, the scheduler has no control over the type of traffic arriving into its queues. Hence, all the entries in the queue may be occupied by a collection of bursty best-effort traffic, causing the scheduler to back-pressure the network, thus preventing higher priority traffic from reaching the scheduler. Therefore, the STT-based scheduler can

only reliably achieve the required system performance when the traffic from the different initiators is carefully shaped – both within an initiator and across entire service classes. Conversely, the achievable performance for the latency and bandwidth sensitive data flows can be severely impacted when the traffic is composed of heterogeneous flows with diverse performance requirements.

DRAM chips used in high-end SoCs operate at 533 MHz (with data pin rates twice that) or higher. The memory scheduler should normally operate at similar frequencies to deliver acceptable throughput. In a MTNB system, the number of threads defines the fan-in to the scheduler, and thus helps determine the achievable frequency. The system-level SoC design challenge includes optimal mapping of the traffic flows to a smaller number of scheduler threads to obtain the desired performance-area-timing trade-off. Mapping algorithms normally allocate flows with similar characteristics and constraints to the same threads, and allocate relatively fewer flows per low latency or controlled bandwidth thread than per best-effort thread. Such a mapping will minimize the number of threads and thus help in achieving higher operating frequency in the scheduler and minimize area overhead.

In the STT system, the depth of the single queue defines the fan-in to the scheduler. This depth is typically much higher than the number of threads in the MTNB system, so the STT scheduler can minimize the frequency of full queues back-pressuring into the network. Since the SoC designer does not have the flexibility to optimize the size of the queue per initiator, the single queue size must satisfy all operating scenarios. This can force the designer to assume worst-case conditions and over-estimate the size of the queue. Since the STT scheduler examines all the entries in the queue to make its scheduling decision, the increased queue depth increases the complexity of the scheduler and results in lower operating frequencies or deeper, less efficient scheduling pipelines than the MTNB system.

The remaining part of the paper is organized as follows: In Section 2, we describe related work. In Section 3, we provide an overview of thread versus tag, vis-à-vis the OCP protocol. In Section 4, we provide details of the memory subsystem. In Section 5, we describe the STT protocol in detail. In Section 6, we describe the MTNB protocol in detail. In Section 7, we present experimental results, and finally, in Section 8, we conclude the paper.

Continued next page

Previous Work

In the past, researchers have spent considerable effort attempting to achieve performance guarantees in a NoC system. Goossens et al. [4] presented the Aethereal network-on-chip that provides guaranteed throughput and bounded latency. Vellanki et al. [5] presented a NoC architecture where the traffic profile is divided into two categories: best effort, and guaranteed throughput, and the NoC assigns higher priority to the guaranteed throughput traffic over the best effort traffic. Weber et al. [6] presented an analytical analysis to achieve bounded latency and service jitter of a heterogeneous traffic system. Marescaux et al. [7] presented an architecture called SuperGT, where time slots are reserved for guaranteed throughput and best effort traffic respectively, to achieve the required performance guarantees. Bolotin et al. [8] and Srinivasan et al. [9] address the latency minimization problem at network design time by routing latency-sensitive traffic through a minimum number of router hops. All these papers apply QoS at the interconnection network and do not take its effect on the memory subsystem into account. For example, in order to service latency-sensitive traffic, the interconnection may inject traffic into the memory subsystem that is detrimental to DRAM efficiency (due to high page miss rate and read-to-write direction turnaround). In contrast, our work in this paper uses a scheme that closely follows the technique presented in [6], and applies it at the memory subsystem. Thus, it is able to gracefully trade-off the QoS requirements of the latency sensitive traffic, while maintaining high memory utilization.

The attempt to standardize on-chip communication interfaces has led to the advent of protocols such as OCP and AMBA AXI [10]. While OCP supports both MTNB and STT-based protocols, AXI supports only STT-based protocol. To the best of our knowledge, there is a dearth of qualitative and quantitative analysis to determine the relative advantages and disadvantages of the two protocols. In fact, going by prior published work, there seems to be a general lack of clarity on MTNB and STT-based protocols and their implications on the performance of the SoC. For example, adopting an AMBA protocol automatically forces the communication interfaces to be STT-based, and the initiators threads must be collapsed into a single target thread. However, adopting an OCP protocol may allow the initiators to be routed along individual threads. These system-level decisions play a critical role in the optimization of the NoC for the desired power, area and performance. The primary focus of our work is to describe the two protocols and

then analyze them qualitatively and quantitatively in a real video SoC environment.

OCP Tags versus OCP Threads

The Open Core Protocol (OCP) [1] is a point-to-point communication standard developed for interfacing between on-chip components such as initiators, communication networks, and memory schedulers. The protocol provides a layered set of communication semantics that define threads, tags, transactions, transfer phases and flow control. It also provides a high degree of configurability to choose a subset of the full semantics that best meets the needs of a specific interface. The highest layer in OCP is the thread, which defines sequences of transactions that are unordered with respect to transactions on other threads. OCP supports optional per-thread flow control, enabling the delivery of fully non-blocking interfaces in a manner analogous to virtual channels with virtual flit flow control [3]. In a MTNB interface, the receiver of each transfer phase implements a FIFO per thread, guaranteeing that the interface cannot block by asserting flow control on those threads with a full FIFO. Within a thread, transactions may have independent ordering tags; transactions with the same tag (including the common degenerate case of interfaces supporting the null tag) are ordered. The ordering restrictions apply to both transaction completion at the target (ensuring memory consistency) and response ordering at the initiator.

The OCP tag mechanism is similar in principle to the ID mechanism proposed in the AMBA AXI system architecture [10]. Since tagged transactions have shared flow control, a shared queue model is implied where re-ordering is typically only practiced at the target, since deadlock-free re-ordering in the network would require storage for complete transactions. Tag-based systems should thus minimize transaction lengths to prevent a lower priority transaction from blocking a later, higher priority request from reaching the target. While both threads and tags can coexist on the same OCP interface, existing schedulers implement either the MTNB (multiple threads, but only a null tag per thread) or the STT (a single thread with multiple tags) approaches to minimize ...

To read the rest of the article, here. http://www.ocpip.org/uploads/documents/Non_Blocking_Flow_Control_Long_Version_April_2010.pdf

Working Group Updates

Technical Vision Working Group

The Technical Vision Working Group held its annual strategic session to complete the overall technical planning for the organization in March 2010. The next step is for all OCP-IP working groups to confirm their specific roadmaps and goals for work to be completed in 2010. The OCP-IP Governing Steering Committee will continue to monitor and support these targets throughout the year.

Debug Working Group:

The Debug Working Group continues to identify required instrumentation signaling extensions and example cases for debug of multi-core systems that support OCP 3.0 cache and power management features. These may include SMP, AMP, multi-threaded, and other system architectures. Preliminary goals are for OCP 3.0 compatible debug systems white papers to be published in 2QCY10 with updated specifications being developed later in the year.

Meta Data Working Group:

The Metadata Working Group (MDWG), has finalized the enhancements to fully capture OCP interfaces using the IP-XACT format defined by SPIRIT Consortium. The Group now has an updated version of the package which is both IP-XACT 1.4 and IEEE standard (IP-XACT 1.5) compatible. The package will be available to OCP-IP members in 2Q 2010. The MDWG is also created OCP configuration and interface compatibility checkers.

Functional Verification Working Group:

The Functional Verification WG is continuing the activity started at the end of 2009 targeting the definition of the functional verification checks and coverage to support OCP 3.0. Current work is focused on the Cache Coherence Extensions, specifically targeting configuration checks, signal checks and transaction-level checks.

NoC Benchmarking Working Group:

The NoC BWG is currently updating the Transaction Generator (TG) tool. The main focus of the work is incorporating the OCP-IP TLM Modeling kit to simplify connecting the TG to various network-on-chips that are to be evaluated. Additional work is focused on ease-of-use; hence the installation is being simplified and a tutorial is being prepared. Finally, co-simulation of SystemC (TG) and VHDL (NoC) is being studied.

Specification Working Group:

Following the release of OCP 3.0 the Specification Working Group is now focused on minor Errata items and prioritizing goals for 2010 as provided by OCP-IP's Technical Vision Working Group. Be sure to watch future editions of the OCP-IP newsletter for announcements about OCP 3.1 and future plans of the Spec Working Group.

Marketing Working Group:

The Marketing Working Group has recently completed presentations at: NASCUG, International Symposium on FPGAs, EDAC's CEO Forecast, DATE 2010 presentations and IP-SoC Days. In addition, we remain very active helping member companies compose and place their OCP-related articles and conference papers, while publishing the OCP-IP newsletter and regular press releases. If your company would like assistance placing an article, prominently targeted to the industry or directly to our OCP-IP focused community, please contact admin@ocpip.org.

Spotlight on ENSTA

The École Nationale Supérieure de Techniques Avancées enjoys a special position in the French educational system as one of the foremost schools of engineering in the country.

ENSTA is currently engaged with OCP-IP's Network on Chip Benchmarking Working Group and is currently using EEMBC's MultiBench and related materials to create a set of traffic profiles that incorporate the data movements occurring in a multiprocessor system. In the future, these Traffic Profiles will be available on OCP-IP's website and used for benchmarking Network-on-chip designs. This is an important step in standardization of NoC benchmarks as previously, models of parallel application have not been available in the industry.

ENSTA offers its students a broad education in engineering with the aim of enabling them to design, carry out and manage complex technical projects, while meeting economic constraints in an international environment. In this view, the institute provides a high-level scientific and technological education, which is constantly updated to keep pace with changes in the leading edge technologies, and supplemented by languages, humanities, and the skills needed in business life such as law, communication, economics, accounting and management.

Educational practice at ENSTA combines lectures with group courses. Lectures are given by the institute's professors and researchers, with the participation of numerous auxiliary teachers from the economic and industrial world, familiar with the latest technical developments in their field.

Research is ENSTA's other primary mission. Its five departments carry out research in many areas in partnership with French, European and international universities and public research bodies. A large number of research staff from the CNRS, INSERM and École Polytechnique also work at ENSTA alongside the institute's own professors. The departments mostly perform applied research addressing industrial problems, but are also involved in fundamental developments in the interest of scientific knowledge and aimed at technological breakthroughs. Research makes a dynamic contribution to the School's educational work, underpinning its connection with scientific state-of-the-art.

ENSTA delivers about 170 engineering degrees every year. The broad nature of their education enables graduates to find a career in a large number of sectors: mechanical industries such as the automotive or naval industry, the electronic and IT sector, the energy sector, process and environment industries, but also consulting, finance and defence. ENSTA's graduate engineers are much sought-after by companies. A majority of them find their first job in R&D departments and design offices, rapidly moving towards supervisory and project management positions.

ENSTA is a public educational and research establishment, self-governed under the supervision of the Ministry of Defence.

ENSTA is a founding member of ParisTech, the federation of high-level engineering and business schools located in and around Paris. Covering almost the complete range of science and technology, ParisTech is a leading entity in French higher education and a partner to some of the most famous universities of science and technology worldwide.

D&R Hosts First US Event

D&R IP-SOC Days

After 10 successful years the IP-SoC conference and exhibition in Grenoble, France, D&R pursued industry suggestions to hold a similar event in one of the hottest IP-SoC locations in the world – Silicon Valley. This first event outside of Grenoble took place in Santa Clara, March 23 and 24 with 40 IP consumer companies and 35 provider companies.

The 15 sponsor talks (including one by OCP-IP president and chairman Ian Mackintosh) offered a good balance ranging from large players such as Synopsys and TSMC, to medium sized companies including Arasan, GDA Technologies, and Sonics, to small IP houses such as Eureka. The sponsor talks also showed a good balance between technical and business topics. For more information visit <http://www.design-reuse.com/ipsocdays/santaclara/>

At the conclusion of the Santa Clara show a meeting was held to discuss the future of the Tour. A significant need arose among large companies interested in acquiring IPs. There are so many issues in that activity that it is impossible to get all necessary answers without face-to-face meetings, but it is difficult, if not impossible to get face-to-face meetings with providers in the short window before a product must get to market. What they need is an efficient and effective face-to-face meeting forum.

The level of satisfaction with the Santa Clara event has led to a D&R Tour, after the Tel Aviv event in April, to India, Taiwan and China at the end of September.

D&R offers its experience in organizing IP-SoC days around the world with the same level of quality of content as IP-SoC in Grenoble and supported by an intensive B2B dating process on the D&R website (www.design-reuse.com).

If you are interested in becoming a sponsor participant for any of these events please contact Design and Reuse at: <http://www.design-reuse.com/contact/> or gabriele.saucier@design-reuse.com

Recent Publications

Available at <http://www.ocpip.org/articles.php>

Press Releases

March 11, 2010 [OCP-IP Announces Online DATE Virtual-Conference Presentations](#)

January 27, 2010 [OCP-IP Delivers Even More OSCI TLM 2.0 Compatibility in Advanced SystemC TLM Kit](#)

January 19, 2010 [Tampere University of Technology Wins OCP-IP Contributor of the Year Award](#)

Announcements: Now Available

NEW! OCP 3.0 Specification

The Specification Working Group released the OCP 3.0 Specification in November. This latest version contains extensions to support cache coherence and more aggressive power management, as well as an additional high-speed consensus profile and other new elements. For a copy complete our [Research License Agreement](#).

OCP Checker Now Part of CoreCreator II

The OCP checker is a fourth-generation solution for validating protocol compliance of master and slave devices using OCP. It is based on SystemVerilog assertions (SVA) and can be used with all major logic simulators. It supports the complete set of protocol compliance checks defined in the OCP specification and spans the full range of OCP socket configuration options. The OCP checker can also generate trace files in the standard “.ocp” format for post-processing. It can be obtained, as part of CoreCreator II [here](#). For a free copy contact admin@ocpip.org

Debug Specification Version 1.0

The Debug Specification provides guidelines and recommended signal interfaces for on-chip debug of OCP-based systems and related multicore architectures. It describes a debug socket as a framework for IP and tools providers to develop comprehensive and re-usable debug and instrumentation environments that provide on-chip analysis and control features. These include trace, triggering, multicore synchronization, etc., along with recommendations for integration within ESL environments. For a copy of the spec click [here](#).

Articles

[Thumbs up for Open Core Protocol \(OCP\) version 3.0 - Techbites](#)

[Development OCP wrapper for non-OCP reuse IP cores](#)

[An Analysis of Blocking versus Non-Blocking Flow Control in On-Chip Networks - D&R](#)

NoC Benchmarking Specification, Part 2

Part 2 of the two-part NoC Benchmarking Specification presents a generic NoC architecture, a comprehensive set of synthetic workloads as micro-benchmarks, workload scenarios and evaluation criteria. These micro-benchmarks enable you to measure and pinpoint particular properties of NoC architectures, complementing application benchmarks. For more information, [click here](#).

NoC Benchmarking Specification, Part 1

The specification presents a modeling methodology for applications running on multicore systems and it defines an XML format for documenting and distributing NoC benchmarks. It defines a black-box view of the processing elements that discloses only the relevant computational aspects for interacting with the on-chip data transport mechanism. Download our [NoC white paper](#) for more information

Transaction Analysis Tool

The Tool is an innovative, detailed OCP transaction viewer that enables fine-grained analysis of bus transactions. A complete transaction sequence can be traced from request to response along with a host of related information about the transaction. For a free copy contact admin@ocpip.org

OCP SystemC TLM Kits

The new kit is the first, and most advanced TLM-2.0 based, industry-ready kit in existence today. The kits significantly increase performance, ease of use and ensures alignment with the OSCI 2.0 standard. The kits are free as part of OCP-IP membership. For more information contact admin@ocpip.org