Real-time and Embedded Systems Workshop

Arlington, VA USA - July 11-14, 2005

Workshop Program

MONDAY – July 11, 2005

TUTORIAL TRACKS

0900 - 1230 ***Real-time CORBA***
Track 1 Kevin Buesing, Mentoring Engineer, Objective Interface Systems

The OMG Real-time CORBA specification extends CORBA for use in real-time systems. Real-time CORBA provides a clean infrastructure for building distributed applications with time constraints. The addition of Real-time CORBA 2.0, Dynamic Scheduling, provides significant extensions that allow application engineers to "plug-in" alternative workload management and scheduling capabilities. In addition, the Minimum CORBA specification offers a feature-optimized version of the CORBA specification that allows application designers to depend on the reduced feature sets of lightweight ORB implementations. This tutorial will highlight the concepts and features of each of the specifications, along with examples of their use.

1030 - 1230 ***Building and Implementing Concurrent Specifications***
Track 2 Cortland Starrett, Engineering Manager & Stephen Mellor, Chief Scientist, Mentor Graphics

Many requirements documents are described in a linear manner, which makes for “interesting times” when the system is highly concurrent and multi-threaded. Any attempt to take concurrency and distribution into account when expressing requirements can easily lead to premature design and difficulties in retargeting the application to a different environment. This tutorial describes an approach to capturing concurrency in requirements as models and providing rules to sequentialize and sequence them in an implementation on an RTOS.

1015 - 1045 Morning Refreshments

1230 - 1315 Lunch

1330 – 1730 ***Using the Lightweight CORBA Component Model to Develop Distributed Real-time and Embedded Applications (Part 1)***
Track 1
Douglas C. Schmidt, Professor of Computer Science, Vanderbilt University
Frank Pilhofer, Member of Technical Staff, Mercury Computer Systems

This tutorial will explain the key features and mechanisms in the Lightweight CORBA Component Model (CCM) specification, including in-depth examples of applying the OMG Deployment and Configuration Specification for distributed real-time and embedded (DRE) applications. Several examples will be used to demonstrate how these features and mechanisms can simplify DRE application development and integration. These examples will show how to develop CORBA components, how to assemble these components into applications, and how to deploy these applications in the Lightweight CCM run-time environment. Other examples will show how real-time extensions to Lightweight CCM can enable the development of robust, adaptive, and complex DRE applications.
By attending this tutorial, DRE system developers will get a comprehensive introduction to the Lightweight CORBA Component Model. Attendees will also learn how to use CCM features and extensions to develop CORBA applications that can meet stringent real-time QoS requirements. This tutorial will also refer to ongoing OMG standards activities to add QoS support to CCM and UML, as well as enhance CCM by integrating it with Model Driven Architecture (MDA) and Model Integration Computing (MIC) tools and techniques.

1330 – 1730 ***MDA in Eclipse for Embedded and High Performance Systems***
Track 2 Peter J. Fontana, President, Pathfinder Solutions

This tutorial outlines a process and underlying technology for developing and deploying high performance embedded systems using an Eclipse-based environment. The session has a very practical emphasis with live modeling, code generation, system construction and execution. The class will participate in completing a specific PIM component within a moderately complex system, identify and deploy to an implementation topology. A model-level execution environment will use execution-time tracing and will accumulate interface activity. Then, in an effort to improve performance, an alternative architecture will be specified, and the system redeployed to this new target set.

1500 - 1515 Afternoon Refreshments

TUESDAY - July 12, 2005

TUTORIAL TRACKS

0900 - 1200 ***Using the Lightweight CORBA Component Model to Develop***
Track 1 ***Distributed Real-time and Embedded Applications (Part 2)***
Douglas C. Schmidt, Professor of Computer Science, Vanderbilt University
Frank Pilhofer, Member of Technical Staff, Mercury Computer Systems

(Continuation of Monday afternoon session)

0900 - 1200 ***Data Distribution Service***
Track 2 Victor Giddings, Senior Scientist, Objective Interface Systems, Inc.

The Data Distribution Service (DDS) is a recently adopted OMG specification that adds typed, multi-point data distribution between publishers and subscribers. This data-centric communications decouples publishers of information from subscribers in several dimensions. First, the need for identity of publishers and subscribers is moot; information to be shared is described by topic and type only. Second, the transfer of information from publishers to subscribers is asynchronous; subscribers decide whether to listen for changes or to use the currently available value. This decoupling of system components offers advantages for system design and evolution. System design is very flexible. For example, system components can decide to contribute or digest data based on a static system design or any number of dynamic attributes including role, capability, or physical location. Systems can also more easily evolve; data can be put to originally unanticipated uses without affecting the publishers of the data. The specification allows high-performance implementations that will scale well to very large numbers of publishers and subscribers. DDS also offers an easy to understand and use interface. This tutorial will cover the contents of the DDS specification, contrast it with other OMG specifications that offer data distribution capabilities, and discuss its possible application in different system architectures.

1015 - 1045 Morning Refreshments

1200 - 1245 Lunch

1255 – 1300 ***Opening Remarks* – Program Committee Chair**
Andrew Watson, Vice President and Technical Director, Object Management Group

1300 – 1530 **Session 1: *Design Techniques for Real-time and Embedded Applications***
Chair: Andrew Watson, Vice President and Technical Director, Object Management Group

This session covers the application of modeling and agile design techniques to help build correct, reliable and predictable real-time and embedded systems.

Model-Driven QoS Provisioning Techniques for CCM DRE Systems
Stoyan Paunov, Gan Deng, Douglas C. Schmidt, and Anirudha Gokhale, Vanderbilt University

This presentation will illustrate how Model Integrated Computing (MIC) tools can be used to configure and validate quality of service (QoS)-enabled component middleware by (1) demonstrating how the MIC paradigm can be used to model Real-time CORBA QoS policies for Lightweight CCM applications and showing how semantically invalid configurations can be validated using constraints defined in the higher level models, (2) synthesizing XML based real-time configuration information and minimizing the work required by end-users, and (3) providing opportunities for assembly-level analysis and optimization techniques that can detect errors and sources of overhead and non-determinism automatically.

Applying MDD, Generative Programming and Agile Software Techniques to the SDR Domain
Bruce Trask, Architect, PrismTech

This presentation will describe the application of Model-Driven-Development (MDD), Generative Programming (GP) and Agile Software Development (ASD) solutions, tools and techniques to the Software Defined Radio domain. It will include detailed descriptions of not only how MDD techniques and tools can be used but also descriptions of GP techniques that can be successful leveraged including domain-specific languages as applied to the SCA and the Software Radio Domain. Further, this presentation will describe how ASD techniques can not only enable fast reliable application of these technologies that result in working products but in fact are in integral part of the software technology itself.

A Performance Modeling and Simulation Approach to Software Defined Radio
Shawkang Wu, Ph.D., Scientist & Long Ho, Senior Software Engineer, The Boeing Company

An important aspect of the development for software-intensive systems, such as software defined radios, is a tradeoff design between flexibility and performance. System flexibility is mostly driven by software capabilities while system performance relies on hardware capabilities. To minimize risks and maximize chance of successful completion of the development, we apply Performance Modeling and Simulation to analyze and evaluate the co-design approach. PM&S enables to verify hardware capability, analyze system performance, and validate system portability. The radio software model is compliant to the Software Communications Architecture specification, and the hardware model is based on the Joint Tactical Radio architecture.

Leveraging Agile Software Design Techniques in the Implementation of Software Defined Radios
Vincent Rivas, MTS & Joseph N Frisina, Manager-Advanced Software Development, BAE Systems

This presentation is an experience report outlining the development activities associated with the application of Agile software design methods to a Software Defined Radio project. It will be shown how key Agile software design activities such as test driven design, refactoring and automated regression testing, used in conjunction with sound object oriented design principles and patterns, enabled the creation of a set of superior software products. The integrated development environment and key set of customized software processes used to support these Agile activities will also be presented.

1500 – 1900 ***Demonstration Area Open***

1530 – 1600 Afternoon Refreshments

1600 – 1800 **Session 2: *High Assurance***

Chair: Dock Allen, Lead Software Systems Engineer, MITRE

High-assurance DRE systems are designed to be robust against various types of failure, and to degrade gracefully under adverse conditions. This session describes research and experience reports from differing techniques for engineering high-assurance systems to resist different types of failure, including failures of the security perimeter around mission-critical DRE systems, and failures of parts of the hardware platforms on which they run.

High Assurance CORBA and DDS Profiles

Victor Giddings, Senior Scientist, Objective Interface Systems, Inc.

Creating high-assurance safety-critical applications is a non-trivial effort. The criteria such systems must adhere to results in a system designer having a limited choice of operating system, programming language, development tools, third party libraries, etc. Additionally, there are certification organizations that have rigorous requirements for documentation and artifact production that must be met. These certification requirements are the result of extensive research, both past and present, into high-assurance systems. This presentation will describe the efforts to develop a high assurance profile for CORBA and DDS. The constraints on the languages, tools and operating systems used in creating high-assurance distributed systems will be described. In particular, the language profiles, the consequent IDL mapping profiles, and the CORBA profiles will be discussed.

Multiple Independent Levels of Safety and Security: High Assurance Architecture

Gordon Uchenick, Mentor/Principal Engineer, Objective Interface Systems, Inc.

The Multiple Independent Levels of Security/Safety (MILS) architecture greatly reduces the amount of security enforcing code while simultaneously making that code more effective. MILS enables security functions to be layered among a kernel, middleware, and applications. MILS also enables application developers to implement their own security policy enforcement and be guaranteed their protection functions are Non-Bypassable, Evaluatable, Always Invoked, and Tamper-Proof. Application level security policy enforcement can be layered, preserving prior certifications for the kernel and middleware.

Living Realistically with Nondeterminism in Fault-Tolerant Middleware

Joseph G. Slember and Priya Narasimhan, Carnegie Mellon University

Through MEAD's program-analysis framework, we aim to exploit, and derive insights from, the application sourcecode, in order to provide for more effective handling of nondeterminism. This is accomplished through a synergistic combination of compile-time program analysis tools with a run-time distributed fault-tolerant infrastructure in order to gain application-level insights that can lead to more targeted fault-tolerance without compromising consistency. We have currently implemented MEAD's program-analysis-based approach to nondeterminism for both active and warm passive replication styles in a distributed CORBA system. This presentation focuses on the key features of this approach, along with our promising preliminary results that indicate advances in our ability to support the development of real-world fault-tolerant middleware systems.

1800 – 1900 ***Demonstration Area Reception***

hosted by PRISMTECH ***and*** BAE SYSTEMS

WEDNESDAY, July 13, 2005

0900 - 1200 Session 3: *Experiences with COTS Implementations of the Data Distribution Service (DDS)*

Chair: Steve Jennis, SVP, Corporate Development, PrismTech

To quote the OMG's web site description of the DDS, "Predictable distribution of data with minimal overhead is of primary concern to real-time applications. Since it is not feasible to infinitely extend the needed resources, it is important to be able to specify the available resources and provide policies that allow the middleware to align the resources to the most critical requirements." This session will start with an overview comparison of CORBA and DDS highlighting: selection criteria, synergies and coexistance. It will go on to explore the potential of DDS and experiences with benchmarking, data consistency and configuration of DDS implementations with respect to predictability, fault tolerance and resource optimization.

CORBA / DDS - Competing or Complementary Technologies?
Ramzi Karoui, EMEA Technical Supervisor, PrismTech

An in depth analysis of distributed applications reveals that canonical interactions between them are either Service-Centric, where applications make queries and expect results, or Information-Centric, where applications share information by publishing or reading it. Both CORBA and DDS middlewares can be selected to implement one of those canonical interactions, however communication efficiency, performance, scalability and application maintainability may suffer when the middleware technology is used beyond it's initial application scope. This presentation will compare and define the application scope of each technology.

OMG-DDS "Exploiting The Potential"
Hans van 't Hag, Product Manager Middleware, Thales Naval Netherlands

With the finalization of the OMG-DDS specification, an important milestone has been reached in providing a complementary (to CORBA) set of services to aid in the design and implementation of mission-critical and distributed real-time systems. The potential of an 'information-centric approach' in the area of dynamic and spontaneous systems that require loosely coupled infrastructures is already being recognized and the challenge now lies into 'how to exploit' the OMG-DDS specification and its many services to transform this 'potential' into real-life systems. This presentation starts with the requirements that are imposed on the class of mission-critical real-time distributed systems (such as Naval Combat Systems) and then explains how an information-centric approach, as is supported by the DDS, addresses requirements for these systems as they manifest themselves throughout their lifecycle-stages.

Open Architecture Computing Environment: DDS Middleware Benchmarking
Bruce McCormick, Engineer, EG&G Inc. &
Leslie Madden, Technical Lead Scientist, Naval Surface Warfare Center Dahlgren Division

This presentation will discuss the results of a standalone benchmarking effort involving two precursors to DDS products, Splice2 and NDDS 3.1. The purpose of the effort was to characterize the performance of different products, for multiple configurations. This project provided insight and data on key implementation issues, and performance. Additionally, we will discuss a related effort to incorporate early release features of DDS into the benchmark tool. The products used included the DDS APIs in SPLICE2 and the NDDS Limited Access Release 2 (LAR2). This effort provided insight into some of the issues of porting from existing publish/subscribe middleware to DDS-compliant middleware.

Data Consistency with the SPLICE Middleware
Leslie Madden Technical Lead Scientist & Chad Offenbacker, Scientist,
Naval Surface Warfare Center Dahlgren Division

SPLICE is an emerging DDS-compliant, real-time, publish/subscribe middleware. It has a distributed computing infrastructure which offers distributed data store functionality for storing and retrieving subscribed data. To examine how this capability can be leveraged to develop fault-tolerant combat system applications, an experiment was constructed to ascertain the extent to which transient state data is maintained across distributed data consumers under a variety of fault conditions. This presentation will discuss the design of the experiment and present the results obtained through the comparison of state value over time across multiple hosts under different fault conditions.

1000 - 1600 ***Demonstration Area Open***

1030 - 1100 Morning Refreshments

1200 – 1230 **Sponsor Presentation – PrismTech**
RTE Middleware versus Homegrown Solutions. Are We Winning?
Steve Jennis, SVP Corporate Development, PrismTech Corporation

Unlike for enterprise systems integration, where CORBA and then J2EE became the dominant non-proprietary integration standards in the 1990s, embedded systems integration is still based overwhelmingly on proprietary and often homegrown architectures. In this presentation, we will explore how standards evolution (e.g. CORBA/e, DDS, RTSJ, SCA) and new COTS implementations (e.g. OpenFusion e*ORB SDR over switch fabrics, Splice DDS, OpenFusion RT Java ORB over Sun's Mackinac, GIOP on FPGA, Spectra advanced modeling and code generation tools) that directly address past barriers to adoption (e.g. overhead and performance issues versus proprietary solutions) are slowly but surely fostering a revolution in RTE systems development that could parallel the success of standardization in enterprise systems.

1230 - 1315 Lunch

1330 – 1400 **Sponsor Presentation – BAE Systems**
Case Study Comparing Traditional VHDL Design to Rapid Development (Model Based) Design Using the Xilinx System Generator Tool
Dr. David Haessig, Sr. MTS, BAE Systems CNIR

1400 – 1600 **Session 4: *Real-time Java: The New World of RTSJ***
Chair: Joseph M. Jacob, Senior Vice President, Objective Interface Systems, Inc.

With the advent of multiple different RTSJ (Real-Time Specification for Java) implementations finally hitting the market, real-time and embedded developers now have the opportunity to build hard real-time systems with Java. However, RTSJ is not classic Java in the traditional sense, and creates potentially different paradigms for garbage collection, memory management, etc. This session will explore the benefits and pitfalls of using RTSJ, as well as opportunities for improving the experience going forward.

Experience Migrating a Representative Naval Combat Systems Application to Real-Time Java
Tim Childress, Scientist & Barbara Doyal, Lead Scientist, Naval Surface Warfare Center Dahlgren Division

While standard Java is a powerful and convenient programming language, it is not suitable for creating time deterministic applications. Vendors of real-time Java have originated two technical approaches, RTSJ and real-time garbage collectors, to address this deficiency. We will present the design and results of an experiment where a real-time weapons system application was ported to real-time Java under each of the two paradigms. In conclusion, we will discuss the possible implications for use of Java in U.S. Navy Open Architecture applications.

Experiences in Developing a Real-Time Java Object Request Broker
John Russell, Product Manager, PrismTech

PrismTech has recently undertaken the development of a Java ORB designed to conform with the Real-Time CORBA specification and to meet stringent performance requirements in the defense sector. A detailed investigation of the performance of distributed real-time Java systems has been carried out and in the presentation, PrismTech will describe the design goals and the investigation methodology adopted. The test scenarios will be described and the results obtained will be presented together with their analysis. Aspects of the design of the Real-Time ORB will be described.

Recommendations for a CORBA Language Mapping for RTSJ
Victor Giddings, Senior Scientist, Objective Interface Systems

With the emergence of RTSJ implementations, it has become apparent that full utilization of this technology in Real-Time applications would be enhanced by certain changes to the IDL to Java language mapping. The possible mapping changes occur because of these RTSJ considerations: memory management, thread types and thread priorities. These three features combine to cause problems in the Java language mapping in three different areas: stubs and skeletons, thread type matching, and priority mapping: This presentation will explain these problems and discuss potential solutions.

1600 – 1630 Afternoon Refreshments

1630 – 1800 **Panel: *Real-time Java***
Moderator: Joseph M. Jacob, Senior Vice President, Objective Interface Systems, Inc.

The buzz surrounding Real-time Java was big a few years ago, then died down as it took some time for RTSJ implementations to come to market. Now, with several robust RTSJ products becoming available this year, Real-time Java is back. Is it a viable technology to build robust, high-performance embedded systems? What do developers need to be thinking about as they decide whether to adopt the technology? What benefits and pitfalls await Real-time Java programmers? This panel will discuss these issues interactively with members of the audience.

Panelists: Dr. David F. Bacon, Research Staff Member, IBM's T.J. Watson Research Center
Dennis Govoni, Technical Director, Sun Microsystems Inc.
Steve Kautz, Ph.D., Senior Software Engineer, Aonix
Additional Panelists TBA

1800 - 2000 ***Workshop Reception***

hosted by PRISMTECH ***and*** BAE SYSTEMS

THURSDAY, July 14, 2005

0900 - 1215 **Session 5:** ***Real-time Middleware***

Chair: Doug Jensen, Consulting Scientist, MITRE Corporation

Real-time middleware continues to be a fruitful area for research and development, especially in the area of component-based approaches to configuring DRE systems, and papers in this session reflect this.

The COMPARE Project: Component-based Approach for Real-time & Embedded Systems
The COMPARE Consortium - Presenter: Vincent Seignole, Engineer, THALES

The COMPARE project will adopt the OMG CORBA Component Model (CCM)/Lightweight CCM standard, and adapt it to the Real-time and Embedded application domain. It is expected that this approach will provide a well-defined component model architecture that will enforce a distinct separation of concerns. This presentation will show the advantages of the COMPARE model by highlighting its features. A description of the proposed technology demonstrators for the COMPARE project will also be given; these include proof-of-concept test-beds for an SDR Waveform (whose specification is available in the public domain) deployed onto SDR hardware (including DSP and GPP devices), and a software defined electrical-circuit breaker also requiring a real-time capability.

A QoS-aware Deployment & Configuration Engine for Component-Based DRE Systems
Gan Deng, Jaiganesh Balasubramanian, Arvind Krishna, William R. Otte, Krishnakumar Balasubramanian, Anirudha Gokhale, and Douglas Schmidt, ISIS, Vanderbilt University & Nanbor Wang, Tech-X Corp

This presentation describes techniques for enhancing the OMG's D&C specification to compose real-time Lightweight CCM-based applications and show how they these techniques make it easier to develop, validate, and evolve such applications. In particular, we describe (1) how we enhanced the D&C specification to provide a flexible configuration and deployment mechanism of components, (2) how application developers can specify real-time policies to provision the QoS CCM components via an XML-based metadata format, and (3) how these added capabilities are being applied to develop applications in shipboard computing, avionics mission computing, process control, software defined radio, and inventory tracking systems.

Simplifying the Development of QoS-aware EJB Applications via Model-Integrated Computing
Jules White, Douglas Schmidt and Aniruddha Gokhale, Department of EECS, Vanderbilt University

This presentation describes a model-driven development (MDD) approach to developing adaptive DRE computing systems using Enterprise Java Beans (EJBs). First, we will cover the structure and functionality of J2EEML, which is a domain-specific modeling language (DSML) that captures the design of EJB systems, their quality of service (QoS) requirements, and the adaptive properties that will be applied to the EJBs. Second, we cover how MDD tools can generate code to plug EJBs into JFense, which is a Java component framework that monitors, configures, and executes EJBs and their adaptation strategies at run-time.

1030 - 1100 Morning Refreshments

Techniques for Dynamic Swapping in the Lightweight CORBA Component Model
Jaiganesh Balasubramanian, Gan Deng, Balachandran Natarajan, Aniruddha Gokhale, and Douglas Schmidt, Department of EECS, Vanderbilt University

This presentation describes the structure, functionality, and performance of a dynamic component implementation swapping capability in Lightweight CORBA Component Model (CCM). Our focus is on the middleware infrastructure concerned with connection management, container policies, and IDL/CIDL enhancement mechanisms that enable components in an assembly to be replaced without significantly affecting the performance and availability of the entire component system. We show how these added capabilities has been applied to DRE applications in real-world shipboard computing, avionics mission computing and process control systems.

Evaluating Adaptive Resource Management for Distributed Real-Time Embedded Systems
Nishanth Shankaran, Xenofon Koutsoukos, Douglas C. Schmidt, & Aniruddha Gokhale, EECS Dept., Vanderbilt U.

This presentation describes the structure and functionality of the Hybrid Adaptive Resource Management Middleware (HyARM), which provides advanced distributed resource management based on hybrid control techniques. The relationship between HyARM and Real-time CORBA is presented, along with a demonstration of HyARM's adaptive behavior on a DRE multimedia system that provides real-time video distribution. This demonstration indicates that the HyARM can yield predictable, stable, and better system performance, even in the face of fluctuating resource availability.

1215 - 1300 Lunch

1315 - 1435 **Session 6:** ***Middleware for Specialized Devices***

Chair: Gordon Uchenick, Mentor/Principal Engineer, Objective Interface Systems

Open architecture frameworks built on Real-time CORBA, such as the SCA, have enabled significant portability for code running on General Purpose Processors (GPPs). However, performance requirements and SWaP (Size, Weight, and Power) constraints typically dictate that the system designer utilize specialized devices such as Digital Signal Processors (DSPs) or Field Programmable Gate Arrays (FPGAs) to meet RF signal processing requirements. Unfortunately the development tools for these devices rarely provide any standards-based vendor independence of resulting logic. The result is that a significant portion of waveforms is specific to a vendor-specific device architecture. Version 3.0 of the SCA introduced a C transport abstraction called the Hardware Abstraction Layer (HAL-C) to abstract away some communication details of a DSP or FPGA device communicating with a GPP. Version 3.1 of the SCA attempts to further refine and extend these abstractions. The speakers will present alternative approaches to providing waveform portability for the portion of the waveform logic developed on DSPs and FPGAs.

Using CORBA on DSPs, FPGAs & Microcontrollers - Synthesizing an Ubiquitous SCA Machine in Soft-Radio Devices Andrew Foster, Product Manager, PrismTech

Current debate concerning the use of hardware abstraction layers on software-defined radio platforms currently perceived to be non-CORBA stations has given rise to two schools of thought: 1) defining a HAL (Hardware Abstraction Layer) specification or 2) using an approach based on the OMG's PIM and PSM for Software Radio Components, where these components and their interfaces are defined in CORBA IDL. The SCA today is one of the more powerful models being used to define the canonical structure of the soft-radio device. However in terms of implementation, SCA makes little stipulation about the use of CORBA on the more exotic elements and platforms in the radio's signal processing chain – viz. the DSP and the FPGA. This presentation will detail the latest advances in high performance real-time ORB technology based on MIPS normalized design methodologies.

Middleware for DSPs and FPGAs
Bill Beckwith, CEO/CTO, Objective Interface Systems

To minimize deployment cost while achieving the required current and future functionality, system designers must balance the following factors: initial ramp up cost, period to amortize the initial research and development, cost per deployed unit, power consumption, legacy intellectual property, existing engineering skills and paradigms and performance. The appropriate communications middleware allows system designers to construct flexible, maintainable systems that can accommodate the widest possible range of waveform computing loads while maintaining economic goals for a target platform cost profile. This presentation will compare and contrast multiple approaches to using communications middleware on specialized devices such as DSPs, FPGAs and ASICs.

1435 - 1450 Afternoon Refreshments

1450 - 1620 **Session 7:** ***Protocols***

Chair: Douglas C. Schmidt, Professor of Computer Science, Vanderbilt University

For years, developers of DRE systems have built applications that ran over special-purpose networking protocols to meet their end-to-end QoS requirements. These types of systems are increasingly being built using middleware and networking standards. To support the end-to-end QoS requirements of these DRE systems, the middleware must be integrated seamlessly with the underlying protocols. This session focuses on advances in the networking and ORB layers that significantly enhance the support for stringent QoS requirements of DRE systems.

Practical Experiences Using the OMG's Extensible Transport Framework (ETF) Under a Real-Time CORBA ORB to Implement QoS Sensitive Custom Transports
Andrew Foster, Product Manager, PrismTech

The use of novel and unusual transports under CORBA middleware has been a subject of much interest for the real-time CORBA community. The use of specialized QoS-sensitive transports has been shown to provide significant benefits in terms of performance and reliability. The OMG specification for extensible transports aims to simplify and unify the thinking behind the idea of plugging in custom transports under an ORB. Whilst this specification offers major benefits, some aspects that emerge from it require further clarification. This presentation will highlight both the successes and shortcomings of the ETF specification based on practical experience from the field. The presentation will also provide real-world examples of the use of custom transports for secure and narrow-bandwidth ORB communications for classes of application requiring a high degree of QoS control.

The Time Triggered (TT) Ethernet
Astrit Ademaj, Hermann Kopetz, Petr Grillinger & Klaus Stainhammer, Vienna University of Technology

In this presentation we introduce the design of a time-triggered (TT) Ethernet, which unifies real-time and non-real-time traffic into a single communication architecture. TT Ethernet distinguishes between two traffic categories: the standard event-triggered Ethernet traffic and the time-triggered traffic that is temporally guaranteed. The event triggered traffic in TT Ethernet is handled in conformance with the existing Ethernet standards of the IEEE. TT Ethernet is intended to support different types of distributed applications, from simple data acquisition systems, to multimedia systems up to the most demanding safety-critical real-time control systems which require a fault-tolerant communication service that must be certified.

Real-time and Embedded Systems Workshop Program Committee

Chair: Andrew Watson, *Object Management Group*

Dock Allen, *MITRE*
Shahzad Aslam-Mir, *PrismTech*
Bill Beckwith, *Objective Interface Systems*
Carol Burt, *2AB*
Ben Calloni, *Lockheed Martin Aeronautics*
Lisa DiPippo, *University of Rhode Island*
Andrew Foster, *PrismTech*
Victor Giddings, *Objective Interface Systems*
Andy Gokhale, *Vanderbilt University*
Joe Jacob, *Objective Interface Systems*
Doug Jensen, *MITRE*
Vana Kalogeraki, *University of California, Riverside*
Hermann Kopetz, *TTTech*
Kevin Loughry, *Object Management Group*
David McKinnon, *Washington State University*
Steve Mellor, *Mentor Graphics*
Jishnu Mukerji, *Hewlett Packard*
Steve Osselton, *PrismTech*
Irfan Pyarali, *OOMWorks*
Doug Schmidt, *Vanderbilt University*
Bran Selic, *IBM Rational Software*
Dave Stringer, *Borland*
Jon Siegel, *Object Management Group*
Richard Soley, *Object Management Group*