Real-time and Embedded Systems Workshop

Arlington, VA USA - July 9-12, 2007

Program

MONDAY - July 9, 2007 - *Tutorials*

0900 - 1230 Track 1

Model-Driven Engineering for Distributed Real-time and Embedded Systems

Aniruddha Gokhale, Assistant Professor, Vanderbilt University

Reusable software components and standards-based component models are increasingly being used to develop large-scale distributed real-time and embedded (DRE) systems. This trend, however, also introduces new complexities, including the need to design consistent component interface definitions, validate interactions between components, generate component deployment descriptors and verify that the selected configuration and deployment satisfies end-to-end QoS requirements. Model-Driven Engineering (MDE) has emerged as a promising means to address these issues by combining domain-specific modeling languages (DSMLs) with generators that analyze certain aspects of the models and then synthesize some of the required code, deployment descriptors, and so on. This tutorial provides an overview of MDE for DRE systems, focusing on:

- Fundamental concepts of MDE including DSMLs, system execution modeling (SEM) and generative programming.
- How MDE tools and metamodeling are applicable to DRE systems.
- Role of code generation and model-to-model transformation in meeting QoS requirements of DRE systems.
- Role of MDE in design-time analysis of DRE system properties.
- Deploying and configuring middleware and DRE applications using MDE tools.

Many of the topics mentioned above will be introduced using examples and case studies from production DRE systems. Wherever possible, we'll show live demos of using MDE tools in the tutorial.

0900 - 1230 Track 2

Real-time & Embedded CORBA Tutorial

Victor Giddings, Senior Scientist, Objective Interface Systems
Kevin Buesing, Mentoring Engineer, Objective Interface Systems

This tutorial covers two CORBA variants adapted for RTE systems:

- Real-time CORBA extends CORBA for use in real-time systems, providing a clean infrastructure for building distributed applications with time constraints. Revision 1.2 of the specification adds Dynamic Scheduling, which provides significant extensions that allow application engineers to "plug-in" alternative workload management and scheduling capabilities.
- CORBA/e offers a feature-optimized version of CORBA with a reduced feature set, permitting lightweight ORB implementations with a reduced footprint.

This tutorial will highlight the concepts and features of each of the specifications, along with examples of their use.

1030 - 1045 Morning Refreshments

1230 - 1315 Lunch

1330 – 1700 ***Real-time Data Distribution Service Tutorial***
Track 1 Hans van't Hag, Product Manager, PrismTech

In the large class of network-centric systems, real-time availability of information is of utmost importance. Information generated from multiple sources must be distributed and made available to 'interested parties' taking into account Quality of Service (QoS) offerings by information-producers and requests by information-consumers. Especially in real-time and mission-critical systems, getting 'the right data at the right time at the right place' is not a trivial task yet targeted by the OMG's DDS specification. In this Tutorial we will provide an overview of the complete DDS specification covering Pub/Sub concepts, background of the specification, introductions to DCPS, DLRL and related data-modeling and QoS policies. All this will be completed with some real-world use-cases for the many facets and features of the total DDS specification.

1330 – 1700 ***Using the Lightweight CORBA Component Model to Develop***
Track 2 ***Distributed Real-time and Embedded Applications***
Douglas C. Schmidt, Professor of Computer Science & Will Otte, Vanderbilt University

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, showing 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. Further 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.

1500 - 1515 Afternoon Refreshments

1700 – 1800 **BOF: *TAO Users Session***
Moderator: Douglas C. Schmidt, Professor of Computer Science & Aniruddha Gokhale, Assistant Professor, Vanderbilt University

---

TUESDAY – July 10, 2007 - *Tutorials & Presentation Sessions*

0900 - 1200 ***DDS Advanced Tutorial: Using DDS QoS to Solve Real-world***
Track 1 ***Problems***
Gordon Hunt, Chief Applications Engineer, Real-Time Innovations, Inc.

This tutorial complements Monday's DDS tutorial, and assumes a basic understanding of DDS's information model and quality-of-service parameters.
From fail-over and redundancy, system startup and discovery, reliability and synchronization, to data and topic design, QoS settings affect all aspects of an application. With over 30 different quality-of-service parameters, many of them with coupled functionality, properly setting and leveraging all of the capabilities in DDS can be challenging. Understanding these QoS parameters is the first step to utilizing the full capabilities of DDS. Applying and combining them together with appropriate settings is the second. This tutorial will cover many common application design patterns and corresponding DDS Qos configurations and will show how common distributed-system requirements can be mapped to DDS QoS settings.

0900 - 1200 Track 2 ***An Introduction to the OMG Systems Modeling Language (OMG SysML™)***

Sanford Friedenthal, Principal Systems Engineer, Lockheed Martin

The OMG Systems Modeling Language (SysML) is a general-purpose graphical modeling language for specifying, analyzing, designing, and verifying complex systems that may include hardware, software, information, personnel, procedures, and facilities. In particular, the language provides graphical representations with a semantic foundation for modeling system requirements, behavior, structure, and integration with a broad range of engineering analysis. SysML represents a subset of UML 2.0 with extensions needed to satisfy the requirements of the UML(tm) for Systems Engineering (SE) RFP.

This tutorial provides an introduction to how the SysML extension to UML can address the needs of the systems engineer. It includes background and motivation, and provides a summary of the different SysML diagram types and extensions to UML.

1030 - 1045 Morning Refreshments

1200 - 1245 Lunch

1255 - 1300 ***Welcome & Opening Remarks*** – **Program Committee Chair**

Andrew Watson, Vice President & Technical Director, Object Management Group

1300 – 1330 Sponsor Presentation – ***Vanderbilt University***

1330 – 1530 **Session 1:** ***Distributed Real-Time Java Middleware***

Chair: Angelo Corsaro, Scientist, SELEX-SI

Industrial strength implementations of Real-Time Java have been available on the market for a few years. However, the same cannot be argued for Distributed Real-Time Java Middleware. The intent of this session is to provide an update on the state of the art in industrial and academic R&D on Distributed Real-Time Java Middleware.

Building Successful Real-time Distributed Systems in Java

Andrew Foster, Product Manager, PrismTech

This presentation will discuss the on-going development of a Real-time CORBA ORB for the Java platform. The talk will highlight the key concerns that must be addressed for users to successfully develop and deploy Real-time CORBA systems in a Java. Issues discussed will include: improving the IDL to Java language mapping in order to reduce garbage collection, supporting multiple Real-time modes of ORB operations (non Real-time, soft Real-time and hard Real-time) and threads of execution from the same application process, techniques for maximizing re-use of legacy code in a Real-time Java application, leveraging automatic garbage collection to improve system predictability and improve ease-of-use, what are the tradeoffs that can be made between RTSJ features, real-time predictability and ease-of-use, and finally a comparison between the Real-time performance that can be achieved for different combinations of ORB modes of operation, RTSJ features, and garbage collection strategies.

Design and Evaluation of a High-Integrity DDS Profile

Vincenzo Caruso, University of Bologna & Angelo Corsaro, Scientist, SELEX-SI

Can challenges introduced by High Integrity Systems be addressed by means of a DDS profile? In this presentation we will show the results of our work in defining and evaluating a DDS profile for real-time, embedded and high-integrity systems. It is important to have extremely predictable, uniform and low-jittered data dissemination while keeping a low middleware overhead and parsimonious use of memory and computational resources. With these objectives in mind, we have elaborated a high-integrity profile for DDS by: making assumptions about the operational context; imposing constraints on network protocols and topology; defining a reference architecture; defining a method to distribute data in an efficient and predictable way through protocol-based multicast; extending DDS APIs in order to allow static participant discovery; and excluding the most complex and time-consuming QoS policies.

The Impact of a Real-Time JVM on Middleware Performance: Lessons from Implementing DDS on IBM's J9
Ken Brophy, Senior Applications Engineer & Rick Warren, Lead Software Engineer, Real-Time Innovations

We describe the lessons learned from implementing a real-time middleware (OMG DDS) on top of a RTSJ Java-Virtual Machine (IBM's J9). The conclusions and development patterns presented affect both middleware developers and users. These include lessons on the RTSJ APIs themselves and their semantics, but also more subtle pitfalls, such as (a) real-time threading and the Java Native Interface (JNI); (b) class loading, static data, and memory areas; and (c) RTSJ compatibility with the Java standard libraries and Java 5 language features, including differences among JVM implementations. A critical component of this presentation is a quantitative analysis of the impact of the RTSJ JVM on the performance and real-time predictability of the resulting middleware and application.

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

1530 - 1600 Afternoon Refreshments in Demonstration Area

1600 – 1800 **Session 2: *Tools for Real-time Software***
Chair: Aniruddha Gokhale, Assistant Professor, Vanderbilt University

Advanced tools play an increasingly important role in RTE software development. This session includes three experience reports on the benefits of tool assistance in different parts of the software development process.

Using the Eclipse Framework to Build an MDD Tool Chain for the DDS Application Domain
Hans van't Hag, Product Manager, PrismTech

Model Driven Development (MDD) is increasingly gaining the attention of both industry and research communities. MDD stresses the use of models in the software development life-cycle and argues automation via model execution, model transformation and code generation techniques. This presentation will describe how the Eclipse frameworks can be applied to the Data Distribution Service (DDS) problem space to provide a MDD tool chain that can be used to address the inherent complexities of the domain by providing an intuitive visual environment including both UML and domain-specific views to support DLRL object-modeling as well as DCPS information/topic-modeling, application code-generation and (future) round-trip engineering supported by dynamic XML/Soap based interfaces with the (remote) DDS target-system.

Architectural and Optimization Techniques for Scalable, Real-time and Robust Deployment and Configuration in Large Systems
Gan Deng, Douglas C. Schmidt, and Aniruddha Gokhale, Department of EECS, Vanderbilt University

This presentation will describe the architectural and optimization patterns we used for scalable, real-time and robust deployment and configuration (D&C) of large-scale distributed real-time and embedded (DRE) systems. To meet the these requirements, we enhanced the OMG standards-based D&C model to exploit the parallel processing and the concurrent processing techniques simultaneously and leveraged Real-time CORBA features to avoid priority inversion. In addition, the design also ensures the robustness of the D&C even in the presence of partial deployment failure. Our results show these enhancements enable open DRE systems to meet their D&C QoS requirements.

US Navy Automated Software Testing for Real-Time Mission-Critical Systems
Elfriede Dustin, Senior Engineer, Innovative Defense Technologies
Charles Fudge, Lead Scientist & Leslie Madden, Lead Scientist, Naval Surface Warfare Center Dahlgren Division

Automation technologies for software testing are increasingly being adopted by industry to reduce testing costs and improve test quality, however, challenges exist in the introduction and usage of these technologies for mission-critical systems. Automated Test and Re-Test (ATRT) is a Navy initiative to examine the return on investment (ROI) that can be realized through the use of automation technologies for testing naval applications. This presentation will review ATRT efforts to date and examine the processes used to introduce automated testing technologies as well as the testing architecture. The role of standards in automated software testing will also be discussed.

1800 – 1900 ***Demonstration Area Reception*** ***hosted by*** PRISMTECH

1900 – 2000 **BOF: *SOA in Real-Time***
Moderator: Gerardo Pardo-Castellote, Chief Technology Officer, Real-Time Innovations, Inc.

WEDNESDAY, July 11, 2007

0900 – 1230 Session 3: *Software Performance & Scalability*

Chair: Gerardo Pardo-Castellote, Chief Technology Officer, Real-Time Innovations, Inc.

Papers in this session address different aspects of the problem of developing and testing complex, large-scale distributed systems designed to meet stringent QoS & performance requirements.

DDS Scalability: "One size fits all"?
Hans van't Hag, Product Manager, PrismTech

Now that the OMG Data Distribution Service (DDS) specification has matured a few years, it is time to look back and reflect on its applicability in the real-time domain that ranges from small-scale embedded to large-scale mission-critical systems. This presentation will present an overview of requirements from different application domains ranging from platform-centric naval combat systems to net-centric metropolitan transportation systems and will provide an assessment of relevant QoS policies in the DDS specification that drive its applicability in environments where both high-volume as well as low latency must be provided within a single system.

Clock Synchronization in Large Scale Dynamic Distributed Systems
Angelo Corsaro, Scientist, SELEX-SI & Roberto Baldoni, Leonardo Querzoni, Sirio Scipioni, Sara Tucci-Piergiovanni, Department of Computer Engineering, University of Rome

Clock synchronization is a fundamental building block for many distributed applications. The work presented here is motivated by an emergent class of applications and services, operating in very challenging settings, for which the problem of clock synchronization is far from being solved. A promising approach is to embrace a fully decentralized paradigm in which peers implement all the required functionalities, by running so called gossip-based algorithms. We consider a time-invariant coupling factor identical for all oscillators. Different constant coupling factors are then evaluated to investigate their effect on system perturbations specific to target settings. We improve the system behavior by using an adaptive coupling factor, which reduces the impact of system perturbations. Finally, we show how this kind of solution can be applied to the OMG DDS, and specifically how peers can internally synchronize by taking advantage of the DDS heart-beating mechanism.

Management of Applications in Large and Heterogeneous Systems
Fabrizio Morciano, Software Designer & Dario Di Crescenzo, Software Designer, SELEX Sistemi Integrati S.p.A.

Naval CMS is a specific domain inside the C4I domain (Command, Control, Communications, Computers, Intelligence) that covers systems located on a naval ship that assist the command team in its responsibility for mission execution. In this presentation we discuss the forthcoming Application Management and System Monitoring (AMSM) specification and how the lesson learned during the development of Cardamom's System Management by Selex S.p.A. and Thales, have been employed in AMSM.

1000 – 1600 *Demonstration Area Open*

1030 - 1100 Morning Refreshments in Demonstration Area

The Middleware NetMARK-RT: A Quantitative Metric for Middleware Performance with Emphasis on RT-Relevant Characteristics
Edwin De Jong, Principal Engineer, Real-Time Innovations

Processor performance is measured using standardized benchmarks such as SPEC CPU2006. Similar comprehensive benchmarks do not exist for real-time middleware performance. Typical performance comparisons focus on simple throughput and latency measurements which are hard to compare as these depend on message size, size of the network, QoS, etc. In addition for RT systems, average throughput and latency have limited value if the jitter is large. This presentation contributes towards defining a set of metrics, use-cases, and benchmarks for assessing real-time middleware and proposes three benchmarks: NetMARK-raw, NetMARK-RT & NetMARK-scale. These measure raw performance, performance within real-time considerations, and scalability. Experimental results are presented for several middleware products.

The Use of Web Services Technologies in the U.S. Surface Navy Domain
Fred Weindelmayer, Naval Surface Warfare Center & Paul Haynes, Ministry of Defence

The implementation of basic and extended Web Services (WS) specifications, including WS-Notification (WS-N), WS Reliable Messaging (WSRM), WS Security (WSS) and others, has led to a base of technologies providing a powerful framework for distributed computing among disparate heterogeneous systems. Though there are many benefits to such technologies, questions remain with respect to their applicability in the real-time domain. In this presentation we document our experience using Web Services technologies in U.S. Surface Navy domain. Their benefits and limitations and the context in which they can be utilized will be discussed. We conclude with a discussion of possible implications on the use of Web Services in future Open Architecture weapons systems.

1230 - 1315 Lunch

1330 – 1400 Sponsor Presentation - ***PrismTech***

1400 – 1730 Session 4: *Data Distribution Service: Directions, Extensions and Experiences*

Chair: Victor Giddings, Senior Scientist, Objective Interface Systems

The OMG Data-Distribution Service for Real-Time Systems specification has enjoyed rapid acceptance within the real-time community as a basis for building information-centric systems. DDS is characterized by its flexibility and features: (1) the richness of quality of service, (2) the high-level data-centric distributed-programming abstraction, (3) the performance and modularity afforded by its publish-subscribe model, (4) the built-in fault-tolerance, filtering, etc. The presentations in this session will explore present and future aspects of DDS: first with a practical experience report, then on to possible future directions, including protocol and integration with other technologies.

Developing MilSOFT DDS Middleware
Huseyin Kutluca-Lead Software Engineer, Izzet Emre Cetin-Lead Software Engineer, Ertan Deniz-Software Engineer, Baris Bal-Software Engineer, Murat Kilic-Project Manager, Ugur Cakir-Project Manager, MilSOFT

The MilSOFT DDS implementation contains a full implementation of the Minimum Profile and some of the optional profiles features. In addition, DDS Automatic Code Generator and DDS Spy software tools have been developed. MilSOFT DDS middleware has found applications in multiple projects including combat management systems, multilink and UAV image exploitation systems. This presentation will be focusing on MilSOFT DDS architecture, DDS tools, application areas and performance results.

Scalable DDS Discovery Protocols Based on Bloom Filters
Javier Sanchez-Monedero, Javier Povedano & Juan M. Lopez-Soler, Associate Professor
Signal Theory, Telematics and Communications Department, University of Granada

The OMG Data Distribution Service (DDS) standard has been recently adopted for real-time data communications. In any DDS platform, to deal with heterogeneous environments and to provide transparent communications, an entity discovery protocol is required. Broadly speaking, an exhaustive discovery approach could result in quadratic complexity. Obviously, the scalability of DDS infrastructure is somehow limited by the adopted discovery protocol. In this work, inspired by peer-to-peer and agent based networking services, we propose a preliminary study on scalable discovery protocols for DDS. In particular, we propose and evaluate solutions to the discovery DDS scalability problem by using Bloom Filters (BF).

Analysis of the Advanced Message Queuing Protocol (AMQP) and Comparison with the Real-Time Publish Subscribe Protocol (DDS-RTPS Interoperability Protocol)
Gerardo Pardo-Castellote, Chief Technology Officer, Real-Time Innovations & Angelo Corsaro, Scientist, SELEX-SI

Both DDS-RTPS and AMQP claim to support advanced features such as zero-copy, message batching, streaming, arbitrary message sizes, one-to-many information flows, scalability to large systems and high-performance. They were, however, designed with different application domains in mind. DDS-RTPS targets real-time and high performance mission-critical systems, which can be highly dynamic in their applications and interconnections (e.g. applications may join and leave, there may be a mix of wired and wireless transports and a mix of platforms from servers to PDAs). On the other hand, AMQP targets financial systems where the topology and message routing is fairly static and there are two classes of machines: clients and servers. This presentation analyzes the two approaches from the structural, functional, and operational point of views.

1530 – 1600 Afternoon Refreshments

Applications of XML in Conjunction with DDS for the Development of Distributed Real-Time Software
Daniel D. Schneider, Software Engineer & John J. Rusnak, Software Lead, Lockheed Martin Space Systems

We explore the use of XML as a software specification language for real-time distributed simulation software, and in particular with the real-time publish-subscribe interface provided as part of the DDS standard. We use XML to formally specify both message definitions (types and topics) and public object behavior (operations) in a coherent fashion. Our XML specification allows developers to express component behavioral interfaces independent of any underlying system API. The mapping of this specification to the underlying simulation API is then performed automatically through use of XSLT.

Integrating DDS and COTS Complex-Event Processing Engines: Semantic Gap and Performance Characteristics
Supreet Oberoi, V.P. of Engineering & Gerardo Pardo-Castellote, Chief Technology Officer, Real-Time Innovations

This presentation will start with a discussion of complex-event semantics and analyze them in terms of fundamental semantic data-model constructs such as generalization, aggregation, classification, and association. Then we will discuss new semantic concepts introduced by CEP: conditional generalization, classification, situations, and reactions. Finally, we will discuss where and how the complex-event model semantically maps to the DDS-data model, and how a system can be built that leverages the capabilities of DDS and CEP. To characterize the class of problems that can be solved using these commercially available pre-built components, this presentation will also analyze the performance characteristics of typical CEP engines and the performance achievable once it is integrated with the middleware.

1730 - 1930 ***Workshop Reception*** ***hosted by*** 

1930 – 2030+ **BOF: *CORBA Future, CORBA/e, C++ Mapping***
Moderator: Victor Giddings, Senior Scientist, Objective Interface Systems

---

THURSDAY, July 12, 2007

0900 – 1015 Session 5: *Security Models for Distributed Publish-Subscribe Messaging Systems*

Chair: Charles Rush, Director of Technology Services, Objective Interface Systems

Distributed systems face significant security challenges due to their requirements for information sharing between address spaces and across processing nodes. Every system must provide capabilities for creating, changing, and enforcing various security policies. Middleware technologies can support security access control policies by providing methods for authentication, authorization, and auditing throughout a distributed system. This session presents two information security models for distributed systems utilizing publish and subscribe messaging architecture.

Security Model and Enforcement for Data Centric Publish/Subscribe Applications with High Information Assurance Requirements

Sebastian Staamann, Director of Security Products, PrismTech

We introduce a information-centric security model for DDS-based distributed applications interaction. The model has three abstraction layers for the definition and enforcement of security policies according to various established security models for high information assurance. At the bottom, the implementation and deployment ensures that data are separated so that no data can infiltrate or exfiltrate from entities of one classification to another. At the middle layer, access control is expressed and enforced in terms of subscriptions and the authorization to publish. The top layer enables the expression of higher-level access control policies which are automatically mapped to middle-layer access control policies.

A Security Model Suitable for Net-Centric for Pub-Sub and Data Distribution Systems

Gerardo Pardo-Castellote, Chief Technology Officer, Real-Time Innovations

This presentation introduces recent research in developing a security model and approach that maps the commonly accepted security concepts (authentication, non-repudiation, confidentiality), security policies (role-based access control, mandatory access control) and technologies (Public and Private Key Encryption, Certificates and Certificate Authorities, TSL/SSL) into a cohesive whole applicable to data-centric publish-subscribe systems in general, and to systems designed using the OMG Data-Distribution Service in particular.

1015 - 1030 Morning Refreshments

1030 – 1145 Session 6: *Verification*

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

Papers in this session focus on techniques for verifying that RTE system performance meets design goals.

QUICKER: A Model-driven QoS Mapping Tool for QoS-enabled Component Middleware

Amogh Kavimandan, Krishnakumar Balasubramanian, Nishanth Shankaran, Aniruddha Gokhale & Douglas C. Schmidt, Department of EECS, Vanderbilt University

The development of distributed real-time and embedded (DRE) systems can be improved by using component middleware. The flexibility of component middleware, however, can also complicate DRE system development since system quality of service (QoS) now depends on how the middleware is configured. In this presentation, we outline the challenges associated with mapping the platform-independent QoS policies into platform-dependent values of QoS parameters. We describe the QUality of service pICKER (QUICKER) toolchain we have developed to address these QoS mapping challenges. We show how we applied model-checking to verify the properties of the transformation. Our experience indicates that our automated QoS mapping process provides significant benefits compared with conventional configuration techniques.

Architectural Considerations for Validation of Run-Time Application Control Capabilities for RT Systems
Paul V. Werme & Antonio L. Samuel, Naval Surface Warfare Center, Dahlgren Division

We will focus on the architectural issues, components, interfaces, data flows and design guidance necessary for validation of both the run-time application control capabilities and of the control-related requirements of the applications being deployed and managed. Issues to be discussed include: required controls, granularity of controls, control semantics, control flows and interdependencies, coordination of control information across components within the architecture, sufficiency of control mechanisms, effectiveness of control mechanisms, coordination of controls for multiple application models, and status and state information required at the application, control infrastructure, middleware, operating system, and network levels. Examples, results, and lessons-learned from both historical and recent NSWCDD Application Management efforts will be presented, as well as a discussion of existing Application Management and Deployment & Configuration standards within the context of support for validation of run-time application control capabilities.

1145 - 1230 Lunch

1245 – 1400 Session 7: *Applying New Methods to Build Better Distributed Real-time Embedded Systems*

Chair: Andrew Foster, Product Manager, PrismTech

This session will focus on applying innovative new methods to build better DRE systems. Presentations in this session will describe the latest thinking on new approaches to solving issues such as trouble-shooting distributed systems and applying techniques to improve the provisioning of network QoS by integrating modeling techniques with QoS enabled middleware.

A Systematic Process and Tool Strategy for Trouble-shooting Distributed Real-time Applications
Gordon Hunt, Chief Applications Engineer, Real-Time Innovations, Inc.

We will introduce the "Systematic Distributed-System Troubleshooting Process" (SDSTP), which has been recently developed based on field experience troubleshooting hundreds of real-time distributed systems. SDSTP identifies a number of tools (both open-source and COTS), and drives the troubleshooting process using a series of questions, each designed to narrow the search for distributed system errors, glitches, and non-deterministic behavior. Each group of questions is designed to focus on specific areas of the distributed system, and is assisted by a set of the tools. The process allows the practitioner to narrowly focus on one aspect at a time while still maintaining a 'system view'.

Middleware-based Network QoS Provisioning for Enterprise Distributed Real-time and Embedded Systems
Jaiganesh Balasubramanian, Sumant Tambe, Aniruddha Gokhale, Assistant Professor, Douglas Schmidt, Professor, Department of EECS, Vanderbilt University
Shrirang Gadgil, Frederick Porter, Balakrishnan Dasarathy, Telcordia Technologies
Nanbor Wang, Tech-X Corporation

This talk describes declarative techniques that integrate modeling and provisioning of network QoS with a QoS-enabled component middleware for enterprise DRE systems. These techniques enhance prior work that predominantly used host resources to provision end-system QoS by integrating a bandwidth broker to provide network QoS into QoS-enabled middleware. We demonstrate and evaluate the effectiveness of our techniques in the context of a representative enterprise DRE system with a single Layer-3/Layer-2 network that supports different types of traffic. Our empirical results show that the capabilities provided by integrating QoS-enabled middleware with network-level QoS mechanisms yield a predictable and efficient system for QoS sensitive applications even in the face of changing workloads and resource availability in the network.

1400 – 1415 Afternoon Refreshments

1415 – 1645 **Session 8:** ***CORBA and CCM for Embedded Systems Development***

Chair: Andrew Watson, Vice President & Technical Director, Object Management Group

This session will show how CORBA and its Component Model deal with three main challenges that new embedded systems need to address:

- Necessity for devices to be integrated in larger scale systems. The first presentation will make a comprehensive comparison between SOA and CORBA as an integration technique
- Seamless development over the whole platform: classically reconfigurable hardware is considered fully separate from the rest of the software running on GPP or DSP, which constitutes a break in the overall architecture and its development. The second and third presentations will introduce how CORBA could be implemented on FPGA, allowing thus a better integration of this part of the system in the whole architecture.
- Increase in software complexity and constraints heterogeneity. The last presentation will show on a practical use-case how CCM and its separation of concerns can help to better organize the software and thus provide actual component reuse.

A Comparison of the Common Object Request Broker Architecture to a Service Oriented Architecture
Roy M. Bell, Principle Software Engineer, Raytheon

This presentation will discuss SOA software alternatives for embedded systems. The advantages, disadvantages, and issues associated with embedded SOA systems will be compared to CORBA. An objective comparison of the rough order of magnitude difference in size and performance will be made. Some non-objective comparisons will also be made for things such as the development process, ease of debugging, and different implementation strategies.

Using a Hardware ORB to Facilitate Seamless Communication with FPGAs in DRE Systems
Andrew Foster, Product Manager, PrismTech

This presentation will discuss the development of a hardware implementation of a CORBA ORB. Specifically targeted for use in FPGAs and ASICs, ICO - the Integrated Circuit ORB provides a subset of CORBA functions required to support the most commonly-used communication patterns. While specifically targeted for use in high performance Software Defined Radio applications and may be used to provide SCA compatibility, it is primarily a CORBA core and may also be used in pure CORBA applications with no SCA requirements.

CORBA Across Devices
Victor Giddings, Senior Scientist, Objective Interface Systems

Recent developments have extended the architectural benefits of CORBA to new classes of devices: DSPs and FPGAs. Efficient and small CORBA implementations on DSPs and FPGAs allow "software objects" on these specialized processors to interface with each other and with general purpose microprocessors using a consistent and efficient infrastructure. The benefits of location transparency are extended to technology transparency: interfaces can now be developed irrespective of implementation location or processing technology, allowing reconfiguration across different processing classes without affecting other software components. This processing mobility reduces time to market and ameliorates risk in the development of real-time and embedded devices.

Experience Report on Implementing and Applying a Standard Real-time Embedded Component Platform
Vincent Seignole, O. Hachet, J.-L. Gilbert, H. Balp, G. Haik, THALES

We will report on the adoption of the OMG open software components standard (OMG Lw-CCM) as a basis for a modular component platform targeting a broad spectrum of real-time and embedded systems requirements, and adaptable to a large range of devices. A synopsis of the technical approach that has been defined and implemented will be provided. An experience report on integration in software-defined radio will be given, and some key benefits in the software development process will be highlighted. We will also provide some insights about other types of operational systems related to this technology. As a conclusion, a European innovation outlook will be given, covering topics including: extensions to new targets (FPGAs, partitioning kernels), reconfiguration, flexible real-time scheduling, scheduling analysis.

Program Committee

Chair: Andrew Watson, Object Management Group

Dock Allen, MITRE
Bill Beckwith, Objective Interface Systems
Angelo Corsaro, SELEX SI
Andrew Foster, PrismTech
Charles Fudge, Naval Surface Warfare Center
Sébastien Gérard, CEA-LIST
Victor Giddings, Objective Interface Systems
Andy Gokhale, Vanderbilt University
Doug Jensen, MITRE
Vana Kalogeraki, University of California, Riverside
James Kulp, Mercury Computer Systems
Stephen Mellor
Jishnu Mukerji, Hewlett Packard
Steve Osselton, PrismTech
Gerardo Pardo-Castellote, Real-Time Innovations
Irfan Pyarali, OOMWorks
Doug Schmidt, Vanderbilt University
Bran Selic, IBM Rational Software
Dave Stringer, Borland
Venkita Subramonian, Washington University
Fred Waskiewicz, Object Management Group
Virginie Watine, Thales
Ben Watson, Lockheed Martin Aeronautics
Johnny Willemsen, Remedy IT