

Adding QoS Dynamic Adaptation Services to TAO for Timeliness, High Availability, and Data Management

Russ Johnston & Trudy Morgan
U.S. Navy SPAWAR Systems Center SD
{russ,trudy}@spawar.navy.mil

Victor Fay Wolfe & Lisa DiPippo
University of Rhode Island
{wolfe,dipippo}@cs.uri.edu

Lou DiPalma, Bob Kelly & John Bagley
Raytheon Electronic Systems
{Louis_P_DiPalma, Robert_E_Kelly--Jr,
John_B_Bagley}@raytheon.com

Peter Kortmann & Ben Watson
Tri-Pacific Software
{peter,watson}@tripac.com

1 Introduction

This paper describes a recently-initiated project to extend TAO to make it better capable of flexibly and effectively adapting to provide a sufficient level of timeliness and other quality of service (QoS) to dynamic environments that have varying workloads and dynamic real-time requirements. The project is a joint effort among The U.S. Navy SPAWAR Systems Center San Diego (SPAWAR), Tri-Pacific Software, Raytheon Electronic Systems (Raytheon), and the University of Rhode Island (URI).

2 Background

This work draws on four primary areas of on-going research at the four institutions/companies: (1) dynamic real-time scheduling technologies; (2) real-time and QoS analysis tools; (3) high availability technologies; and (4) QoS data management technologies.

2.1 *Dynamic Real-Time Scheduling Technology*

For the past seven years SPAWAR, URI, and Tri-Pacific have been working together to develop real-time middleware technology. Our static scheduling service [1] incorporates a uniform static assignment of global CORBA priorities according to the well-known deadline monotonic scheduling technique. It also provides novel optimal algorithms for statically mapping CORBA priorities to the local priorities of the underlying system, as well as an adaptation of the real-time priority ceiling protocols for static locking and resource management [2] to object-based middleware systems [3]. This work was the basis for the static Scheduling Service that was standardized by the Object Management Group (OMG) in Real-Time CORBA 1.0 [4]. The scheduling service was commercialized by Tri-Pacific Software in a product called *RapidSched*, a Real-Time CORBA 1.0 compliant static scheduling service [5]. Prior to this effort, URI and SPAWAR did exploratory research and development on dynamic real-time scheduling in middleware that included dynamic global CORBA priority assignment according to a weighted (by importance) earliest-deadline-first (EDF) priority assignment, and on-going adjustments of CORBA priority due to time passage and environmental changes [6].

2.2 Real-time and QoS Analysis Tools

For over 15 years, Tri-Pacific Software has been developing real-time analysis tools that allow system developers to take an engineering approach to meeting real-time requirements in complex systems. Its original product, *PERTS*, was based on rate monotonic analysis tools originally developed at the University of Illinois and Carnegie-Mellon University [7]. Its most recent version, *RapidRMA* [5], extends the core capabilities of *PERTS* with better end-to-end capabilities, a real-time UML interface to model the system, and the equations for analyzing the real-time performance of middleware that were developed in the work described in Section 2.1 above. The *RapidRMA* tool uses real-time scheduling theory to make quantitative predictions of the predictability, performance, and timeliness of systems. It has been successfully employed in many complex military applications.

2.3 High-availability Technology

Raytheon Electronic Systems has developed software for numerous programs in support of high availability using rapid failover implemented through use of redundant nodes. An example implementation is the Configuration Management component of the Real-Time Distributed Environment for POSIX (RADEX) middleware used on the CCS MK 2 and the new Virginia Class submarine Combat Control Systems. This component manages error processing, monitoring of system node equipment status and reconfiguration due to faults or failures. It determines available nodes and then allocates both RADEX components and application programs among the available resources. Raytheon Electronic Systems has expanded its high-availability R&D effort by combining RADEX, CORBA and other emerging middleware technologies in a combat system implementation to evaluate multiple schemes of dynamic object (re)configurations.

Raytheon Electronic Systems is using TAO in development of OO software for next generation radar programs and is evaluating migration of combat control systems to TAO in an internal study. These programs require high availability, real-time performance and scalable computing resources. Both radar and combat control systems can expect to benefit from applying QoS, dynamic scheduling and dynamic reconfiguration services in upcoming planned releases.

2.4 Middleware QoS Data Management Techniques

Data management services have been prominent in the middleware research and development at SPAWAR, URI and Raytheon. At SPAWAR, the Common Object Framework (COF) project has used CORBA to specify interfaces for common data types to facilitate integrating distributed applications made from legacy subsystems and data types [8]. It has been successfully deployed to integrate legacy data types and services in major military applications such as the Special Operations Forces Command Intelligence and Situational Awareness Tool, in the Global Hawk UAV, and in the U.S. Navy Real-Time Retargeting (REDS) project. At Raytheon the CORBA RADEX middleware Configuration Management component employs a contact and sensor data server [9] that manages data for the entire distributed system as a centralized data service. At URI techniques have been developed for real-time object-oriented

database systems, including real-time object-oriented data models, real-time object-oriented concurrency control [10], and just-in-time data replication techniques for real-time object-oriented database systems in middleware environments [11].

3 Current R&D on Dynamic Reconfiguration Services for TAO

This project will integrate the base technology that is being developed on the R&D efforts described in Section 2 to extend TAO with four Object Services and associated analysis tools. This current R&D will provide four important extensions to TAO: (1) a dynamic scheduling service that can perform dynamic priority assignment and dynamic load adaptation; (2) a dynamic binding service that can choose to which object to bind a client when several candidate objects are possible; (3) a QoS meta-control service that performs monitoring and adaptation by setting policies and parameters of other services; and (4) data management services.

3.1 TAO Dynamic Scheduling Service

There are many different scheduling approaches that provide dynamic priority assignment. For example, both EDF and Least-slack-time-first could assign CORBA priorities to executable entities in TAO. The dynamic scheduling service that we are developing will include algorithms for dynamic priority assignment, dynamically mapping the CORBA priorities to the underlying system's real-time priorities, as well as techniques for real-time resource management (such as priority inheritance-based locking).

When a priority is assigned using a simple technique like EDF, TAO may need to adapt system load to meet dynamic QoS requirements. This adaptation can take several forms. The project will investigate various forms of load adaptation. For example, *admission control* involves keeping tasks out of the system based on dynamic QoS selection criteria. Other techniques take into account the QoS of all scheduled tasks. *Load shedding* involves aborting tasks based on dynamic QoS selection criteria such as importance and remaining execution time. While *load reduction* reduces the level of quality provided by certain tasks in order to meet overall system QoS requirements [12]. This project will measure the effectiveness of criteria used to select tasks for shedding or reduction.

The implementation of this dynamic Scheduling Service for TAO will be based on the Scheduling Service specification for Real-Time CORBA 1.0 [4] and will be consistent with the dynamic scheduling capabilities specified in the Real-Time CORBA 2.0 draft [13]. The technology developed will also facilitate associated QoS analysis to be integrated into Tri-Pacific's *RapidRMA* tool set to provide modeling and analysis tools that accompany the adaptive middleware service. Preliminary algorithms, mechanisms, and techniques for this TAO Scheduling Service will be described at the workshop.

3.2 TAO Dynamic Binding Service

The dynamic binding service will determine which objects can best provide a service required by a client based on real-time and QoS parameters and criteria, such as importance, urgency, and accuracy. The TAO implementation of the Dynamic Binding Service will be based on the standard CORBA Trader Service. The binding decisions will include a determination, based on probabilistic algorithms, of which binding yields the highest probability that the next likely dynamically generated task will meet its real-time requirements. Raytheon's existing node failover middleware technology will be extended to accommodate dynamic binding of resources for not only equipment and software faults and failures, but for conditions where dynamic load varies as well.

A part of this effort will be to extend the Tri-Pacific *RapidRMA* tool set with load balancing capabilities that incorporate the algorithms of this project as well as other effective load balancing mechanisms so that the dynamic binding done by the system can be analyzed for QoS criteria such as timeliness and availability.

3.3 TAO QoS Meta-Control Service

The QoS-Meta-Control Service will use application-level requirements to reflectively "watch" TAO and adapt the Dynamic Scheduling Service and its parameters to best suit the application and/or dynamic environmental conditions. For instance, in a military command and control application the QoS-Meta-Control Service may choose a different Scheduling Service configuration (e.g. EDF vs. other policies, "aging" priorities by increasing them as time passes, changing weights on criteria used in load shedding decisions, etc) for surveillance applications than it would for combat applications. The QoS-Meta-Control Service would similarly configure and adapt the policies and parameters of the TAO Dynamic Binding Service described in Section 3 to allow *rebinding* clients to CORBA objects, based on application requirements. For instance, in a dynamic overload situation the TAO QoS-Meta-Control Service may use the TAO Dynamic Binding Service to rebind a client to a CORBA object or object service with a lower accuracy QoS value in order to meet timeliness QoS requirements. Or, in a failure scenario, the QoS-Meta-Control Service might use the TAO Dynamic Binding Service to reallocate resources using rapid failover techniques to meet availability QoS requirements. The TAO QoS Meta-Control Service will leverage preliminary work at URI/SPAWAR to employ *real-time agent technology* [14] to implement this reflection and adaptation in TAO.

3.4 TAO Data Management Services

In addition to the core middleware services described in 3.1, 3.2, and 3.3, the project will focus on data management services for this dynamic QoS enabled version of TAO. We will leverage the data server technology that Raytheon has effectively deployed in submarine combat control system middleware for the past 10 years. Included will be R&D of novel checkpointing solutions in TAO. SPAWAR has developed on its Common Object Framework projects to integrate legacy data types and services. With this technology, we will integrate the underlying real-time

support in services from the SPAWAR COF project, the ability to move legacy data sharing formerly managed in RADEX middleware to being managed in TAO, and with the ability to support real-time object-based data using QoS/real-time concurrency control and data replication techniques developed at URI. The result will be a TAO Data Management Service that is general enough to adapt to a wide range of applications that have real-time QoS requirements. Novel techniques that will be investigated include *semantic real-time object-based concurrency control* [10], *frontier-based adaptive scheduling* [15] that trades off real-time and accuracy optimality criteria, and *just-in-time data replication* [11]. In addition, this work will involve integrating real-time UML object modeling work done at URI [16] with the UML profile for Schedulability, Performance and Time work developed at Tri-Pacific Software and being standardized by the OMG.

4 Prototyping and Application

The technology resulting from the above R&D will be prototyped as CORBA Services and implemented as services on TAO. The services will be designed to be implemented independently, and in combination with each other, and in combination with existing TAO Object Services.

In addition, a suite of associated tools will be developed to allow modeling and analysis of dynamic QoS performance of TAO-based applications. These tools will continue Tri-Pacific Software's current effort to refine UML as a QoS middleware modeling language. In particular, this project will develop an "engineering layer" for real-time UML. The engineering layer bridges the gap between the expression of abstract models and the implementation of those models. It gives a designer/architect access to the entire UML to model the system and then provides for restricting the implementation or realization of the model by specifying QoS constraints. Furthermore, this project will expand the modeling and analysis tools to add the ability to do efficient automatic TAO-based code generation that comprehends dynamic QoS requirements.

We expect that Raytheon will provide initial testbeds for the dynamically adaptable TAO services and for the associated tools developed on this project by applying them to their existing Virginia Class submarine combat system, International Combat System, and radar programs. All of these are complex applications that have real-time, high-availability, and data management requirements. Also, the real-time COF-based TAO Data Service will be used with the SPAWAR Real-Time Re-Targeting (REDS) project and with the SPAWAR Special Operation Forces Command (SOCOM) projects. Ultimately, we hope to transition the prototyped dynamically adaptable TAO middleware to the Boeing Open Experimental Platform For The Model-Based Integration of Embedded Software (MoBIES), and to the Naval Surface Warfare Center on their HiperD platform, as well to commercial products, and to groups such as the Object Management Group for possible standardization.

5 Conclusion

This paper has presented a summary of the newly-initiated joint effort to extend TAO with QoS dynamic adaptation services that support real-time, QoS, and data services. We believe that this project will provide an important technology enhancement to TAO that will allow it to respond flexibly and effectively in dynamic environments that have varying workloads, dynamic real-time requirements, and data management requirements.

References

-
- [1] L. DiPippo, V. F. Wolfe, L. Esibov, G. Cooper, R. Johnston, B. Thuraisingham, J. Mauer. Scheduling and Priority Mapping for Static Real-Time Middleware, *Real-Time Systems Journal*, 20, 155-182, 2001.
 - [2] R. Rajkumar, Synchronization in Real-Time Systems: A Priority Inheritance Approach. Kluwer Academic Publishers, Boston, MA, 1991.
 - [3] M. Squadrito, L. Esibov, L. DiPippo, V. F. Wolfe, G. Cooper, B. Thuraisingham, P. Krupp, M. Milligan, R. Johnston, R. Bethmangalkar. The Affected Set Priority Ceiling Protocols for Real-Time Object-Oriented Systems. *International Journal of Computer Systems Science and Engineering - Special Issue on Object-Oriented Real-Time Distributed Systems*. Mar. 1999.
 - [4] OMG. *Realtime CORBA*. Electronic document at <http://www.omg.org/docs/orbos/98-10-05.pdf>.
 - [5] TriPacific Software, Inc. www.tripac.com.
 - [6] V. Fay-Wolfe, L. DiPippo, R. Ginis, M. Squadrito, S. Wohlever, I. Zyk. Expressing and Enforcing Timing Constraints in a Dynamic Real-Time CORBA System. *The International Journal of Time-Critical Computing Systems*, 16, 253-280 (1999).
 - [7] Liu, Jane W.S. et. al. PERTS : A Prototyping Environment for Real-Time Systems. Technical Report UIUCDCS-R-93-1082, The University of Illinois, Urbana, May 1993.
 - [8] Trudy Morgan and George Green, *COF Architecture Description* Technical Report, U.S. Navy SPAWAR Systems Center San Diego, 2001.
 - [9] T. Jones, NPES Working Group: Contact Data Management Technical Concept of Operations. NPES working Group, Virginia Class. CONOPS-CONTACTS-07.
 - [10] L. Cingiser DiPippo, V. Fay Wolfe. Object-Based Semantic Real-Time Concurrency Control with Bounded Imprecision. *IEEE Transactions on Knowledge and Data Engineering*, vol. 9 no. 1. Feb. 1997.
 - [11] P. Peddi. *A Replication Strategy for Distributed Real-Time Object-Oriented Databases*, University of Rhode Island Technical Report TR01-282, May 2001.
 - [12] Ethan Hodys. *A Scheduling Algorithm for a Real-Time Multi-Agent System*, University of Rhode Island Technical Report – TR00-275.
 - [13] OMG *rtcorba 2.0??*
 - [14] L.C. DiPippo, V. F.-Wolfe, L. Nair, E. Hodys and O. Uvarov. A Real-Time Multi-Agent System Architecture for E-Commerce Applications, *Proceedings of the The Fifth International Symposium on Autonomous Decentralized Systems*, March 2001.
 - [15] G. Jones, L. C. DiPippo, V. F. Wolfe. Towards Unifying Database Scheduling and Concurrency Control: A Frontier Approach. *Real-Time Databases and Information Systems: Research Advances*, Kluwer Academic Press. Sept. 1997.
 - [16] L. C. DiPippo and L. Ma. A UML Package for Specifying Real-Time Objects, *The Journal of Computer Standards and Interfaces*; 22 (2000), 307-321..