Session on Real-Time, Fault Tolerance, and Distribution

Chair: Michael González Harbour Rapporteur: Juan Antonio de la Puente

1 Scope and aims

Three position papers were included in this session [1, 2, 3]. The aims of the session were stated by the chair as follows:

1. To discuss support for fault tolerance in the Ravenscar profile.
2. To discuss the implications of response time analysis in the Ravenscar profile.
3. To examine the ways that the Real-Time and Distributed System annexes can be integrated.
4. To explore the relationships of the Ravenscar profile and the Distributed Systems Annex.

The discussion was centered on the above topics.

2 Support for fault tolerance in the Ravenscar profile

Miguel Pinho introduced the topic by showing how a middleware for fault tolerant replication in distributed systems can be written, in such a way that replication is transparent to the application code. This is a classical approach to fault tolerance, but the challenge now is how to implement it with the Ravenscar profile, as the support for some required constructs may be lacking. Preliminary work on a prototype middleware has shown that implementing such a middleware on the Ravenscar profile is possible, although the advantages of having a simpler run-time systems may be overcome by the increased complexity of the upper layers. Actually, replication-based fault tolerance with a Ravenscar run-time may be useful in applications that require an efficient platform, but are not necessarily critical.

The issue that a superset of the Ravenscar profile might be better for non-critical systems, in order to reduce the complexity of the applications, was raised by some participants. The following possible extensions of the profile were discussed:

- Timed entry calls. This mechanism has also a strong demand from some evaluators of Ravenscar products in the European aerospace sector, as reported by Tullio Vardanega. After some discussion, it was concluded that absolute delay timed entry calls would be enough for most applications.
- Dynamic priorities. This is a straightforward extension of the profile, which would add little complexity to the runtime.
- Dynamic creation of tasks and protected objects. This is useful for dynamic distributed systems, with variable configurations. A possible workaround is to use a static array of tasks, which are dynamically allocated to application code.

The group discussed how these extensions could be related to the Ravenscar profile, and the conclusion was that a separate profile, which could be specified with the proposed pragma Profile, was better.

The group conclusion on this issue is this kind of extension is interesting, but need not be proposed to be included in the standard yet. The primary interest of the group is to have the Ravenscar profile integrated into the Ada standard first.

3 Response time analysis and the Ravenscar profile

Juan de la Puente presented the work which is being carried out at the Technical University of Madrid on how to adapt response time analysis (RTA) techniques to Ravenscar programs. For a precise modeling, some metrics of the underlying real-time kernel and runtime system are required. Some of the metrics are already required by Annex D, and some other ones can be measured form benchmarks at the application level, but there are still some metrics which can only be measured by having access to the kernel source. Consequently, he proposed that Annex D should include some additional

metrics for Ravenscar kernels, so that a precise timing analysis can be performed.

While there was some support for the proposal, some of the attendants argued that requiring vendors to provide the metrics for validation could make it difficult to develop commercial implementations of the Ravenscar profile. George Romanski and Joyce Tokar made the point that some language constructs may be translated differently to executable code depending on the context, which would make it difficult or impractical to provide metrics for them. On the other hand, Andy Wellings and Neil Audsley stressed the fact that the Ravenscar profile was developed with static timing analysis in mind. However, the precise nature of the metrics that are required depends on the granularity if the analysis. Juan Zamorano said that the clock model was most relevant, together with interrupt latency and context switch times.

The conclusion was that a recommendation to the vendors on the set of metrics to be provided for timing analysis with Ravenscar kernels would be more useful than a strict requirement in the standard itself. The recommendation could be included in the Ravenscar users' guide which is being written by the HRG. Juan de la Puente and Tullio Vardanega will work on writing a document for this purpose.

4 Integration of the Real-Time and Distributed Systems annexes

Javier Gutiérrez reported that the team of the University of Cantabria has been working on this topic, following the recommendations of the previous workshop (IRTAW'10). A document has been written with a proposal for a new DSA, and a prototype implementation has been built, tested, and analyzed for temporal behavior. The analysis model has already been included in the UML-MAST toolset. However, there is still some work to be done before the new proposal can be put in AI form.

The audience expressed some doubts about the interest of the approach. On one side, there is the possibility of using RT-Corba. However, Michael González pointed out that the requirements for RT-Corba are quite different, and make it unsuitable for hard real-time systems or systems built with the Ravenscar profile due to its complexity.

Other objections to the proposal came from the observation that the DSA is only used by a small number of users. However, the reason for this may be that most Ada users are in the critical systems domain, where the current DSA is not of very great value. Tullio Vardanega supported this view and said that the ESA is moving towards distributed systems with an approach similar to the DSA.

As a conclusion, the group decided that the work is interesting. The Cantabria group should continue working on this and report back to the next IRTAW when the work is more mature.

5 Relationship of the Ravenscar with the Distributed Systems Annex

There has been no activity on this issue, even though there was a recommendation from the last workshop.

References

[1] Luis Miguel Pinho and Francisco Vasques. Using Ravenscar to support fault tolerant real-time applications. *Ada Letters*, XXII(this issue), 2002.

[2] Juan Zamorano and Juan A. de la Puente. Precise response time analysis for Ravenscar kernels. *Ada Letters*, XXII(this issue), 2002.

[3] J. Javier Gutiérrez, José M. Drake, Michael González Harbour, and Julio L. Medina. Modeling and schedulability analysis in the development of real-time distributed ada systems. *Ada Letters*, XXII(this issue), 2002.