

Time Management Interoperability and Re-use

David Bruce

Defence Evaluation and Research Agency

St. Andrews Road, MALVERN WR14 3PS, U.K.

dib@dera.gov.uk

Abstract

Interoperability is a deceptive beast, far subtler than it might at first glance appear. After the relatively easy win of “technical interoperability” afforded by the likes of DIS and the HLA, the distributed simulation community is now coming to recognize what will be needed in order to realize the wider notion of “substantive interoperability”. At the Fall 1999 Simulation Interoperability Workshop, for example, Dahmann et al.’s forward-looking paper [99F-SIW-073] drew attention to several of these issues.

In the realm of time management, however, distributed simulation practice continues to distinguish sharply between (real-time, interactive) training simulations and (logical-time, repeatable) analytic simulations. Each is relatively well understood but, not least due to the formidable performance engineering challenges both present, they have evolved quite different programming styles and time management techniques. One might therefore anticipate that the two approaches are fundamentally incompatible — and indeed, extreme forms of time management interoperability (in particular, simulations that must exhibit rigorous analytic properties while simultaneously operating strictly in real time) remain beyond the state of the art, at best a formidable research problem.

Rather than giving up on this aspect of the interoperability challenge, however, I identify room for a more modest advance — namely for federates that can be used either in an analytic or in a real time fashion, according to the needs of the federation as a whole — and illustrate how federates might use the HLA’s time management services to effect such versatility.

By enabling simulations to operate under multiple time management regimes, one gains the ability to interoperate federates that might otherwise be locked into disparate ‘native’ time management schemes, and more generally facilitates their re-use in widely different contexts. Though all but neglected until now, it seems likely that such flexibility will prove a core capability of genuinely re-usable simulations.

About the speaker

David Bruce has been at DERA since graduating with a B.Sc. in Mathematics from the University of Kent at Canterbury in 1988. He is currently investigating emerging software construction techniques and their application to simulation interoperability and re-use; other research interests include parallel and distributed simulation, and programming language design and implementation. He has served on the international Parallel and Distributed Simulation (PADS) workshop steering and programme committees since 1997, and was the programme chair for PADS 2000.

Acknowledgements

This work has been supported by the UK Ministry of Defence under Corporate Research TG5 projects 6.2.3 and 6.2.701. © British Crown Copyright 2001 / DERA.