

**DISTRIBUTED SIMULATION IN INDUSTRY – A SURVEY
PART 3 – THE HLA STANDARD IN INDUSTRY**

Csaba A. Boer

TBA BV
Karrepad 2,
2623 AP, Delft,
THE NETHERLANDS

Arie de Bruin

Delft University of Technology
Faculty of Electrical Engineering,
Mathematics and Computer Science
Mekelweg 4, 2628 CD Delft
THE NETHERLANDS

Alexander Verbraeck

Delft University of Technology
Faculty of Technology, Policy and
Management
Jaffalaan 5, 2628 BX Delft
THE NETHERLANDS

ABSTRACT

Distributed simulation, more specifically the HLA standard, is hardly applied in industry. We have conducted an extensive survey with COTS (commercial off-the-shelf) simulation package vendors and simulation experts, both from defence and industry, that focuses, amongst others, on the question what the reasons are behind this phenomenon. In this paper we analyze the reactions that we obtained, categorizing them into arguments related to distributed simulation in general, arguments related to HLA and arguments pertaining to the embedding of HLA concepts in COTS packages. These answers will lead us, we believe, to insights that can serve as guidelines to make distributed simulation more attractive for the industrial simulation community.

1 INTRODUCTION

In our papers (Boer, De Bruin, and Verbraeck 2006a; Boer, De Bruin, and Verbraeck 2006b), the first two papers in this series, as well as in (Boer, De Bruin, and Verbraeck 2008) we reported on a survey intending to shed light on the observation that distributed simulation, and the HLA standard, is hardly applied in industry, whereas in defence it is widely accepted and applied. The survey consisted of a questionnaire, answered by 18 COTS (commercial-of-the-shelf) simulation package vendors, and of a series of open ended interviews by phone with 16 simulation experts both from industry and defence (COTS vendors, simulation practitioners from research, industry and defence, HLA developers and vendors of HLA products). The survey focused on three main questions:

- Is it true that the HLA standard is hardly applied in industry?
- If so, why is this the case?
- What is needed to remedy this situation?

We obtained an affirmative answer to the first question. This paper deals with the answers that we obtained on the second question. We categorize these answers, providing a structured and comprehensive overview of the reactions from the experts. The third question is dealt with in (Boer 2005) and will be the subject of a forthcoming paper.

A recent survey (Straßburger, Schulze and Fujimoto 2008) partially overlaps and extends our results. The focus there is on distributed simulation in general as well as on distributed virtual environments. Most of the respondents are from research. The survey sheds light on the relevance, the latest trends and the research challenges in these fields.

The starting point in our survey is that simulation practitioners in industry mainly apply COTS packages. This statement has been validated and supported in the interview survey: all experts both from defence and industry agree. According to these experts, the main reason is that these packages are tailored for industry, they allow fast and easy design and development by hiding low level programming details from the simulation practitioners and providing a visual interface.

COTS simulation packages offer predefined building blocks or modules which help simulation practitioners to obtain high level simulation models in a short time period. This is relevant because industry, as one of the experts states, is “less interested in technology and a lot more interested in getting a solution out as fast as possible”.

It was felt that in many cases the best option to build an industrial simulation model is to create a monolithic one, designed and developed in a single appropriate COTS simulation package. On the other hand, experts mentioned characteristics of simulation projects that suggest a distributed approach. The thesis (Boer 2005) elaborates on four of these characteristics, viz. reusability, heterogeneity, collaboration and information hiding.

The industrial community has recognized the advantages of using distributed simulation in such cases. As the questionnaire survey indicates, successful projects have

been carried out in industry in which distributed simulation is applied. However, in these projects mainly homespun distributed solutions for integrating simulation models are applied and rarely use is made of the HLA standard.

The few industrial simulation projects in which the HLA standard has been applied are big and complex (one of the experts mentioned the Japanese electronic manufacturing industry and ship manufacturing for the oil industry) or they are close to defence or government type of applications like the ones from the NASA space industry. In these industrial simulation projects the simulation models required technical depth and mostly general purpose programming languages have been applied for designing and developing them.

But again, these are exceptions. As we indicated before, in this paper we aim to analyze why HLA did not gather ground in industry. We believe that this will lead us to insights that can serve as guidelines to make distributed simulation more attractive for the industrial simulation community. Analyzing the data we collected we observe that there are three main groups of arguments. In Section 2 we present the arguments related to distributed simulation in general. Next in Section 3 we present the arguments related to HLA itself, and finally in Section 4 the arguments related to the relation between HLA and COTS packages. In Section 5 we draw conclusions.

2 ARGUMENTS RELATED TO DISTRIBUTED SIMULATION IN INDUSTRY

According to some of the experts the observation that the HLA standard is rarely applied in industry is already a derivation, and follows from the fact that distributed simulation in general is rarely applied in industry. In the questionnaire the COTS simulation package vendors explicitly mention only a few successful distributed simulation projects. We should thus first analyze why distributed simulation in general is rarely considered by the industrial simulation practitioners.

As we perceive and interpret from the results of the survey, the reasons behind the unpopularity of distributed simulation in industry are primarily market related. The market reason has to do with the cost benefit ratio (The cost benefit ratio is defined as “the net present value of an investment divided by the investment’s initial cost”, from <http://www.trading-glossary.com>). Experts argue that currently, solving a problem using distributed simulation leads to a cost which is too high compared to the benefits that they can gain from it. According to one of the experts “the amount of asset that we would have to go to, to map our objects to their objects would probably outweigh any financial benefit that we would get”. Although one of the experts recognizes some benefits of distributed simulation, he claims that as it stands now “we see only more work”.

Distributed simulation requires a distributed design and development of the models and a tool for interoperation and this entails additional costs. We analyze these costs along three dimensions:

- Monetary cost – Direct monetary costs involved are for instance the purchase price of an existing interoperation tool or the monetary costs of developing such a solution;
- Time – Additional time has to be spent on designing and developing models in a distributed way, and also on applying (or even building) interoperability solutions;
- Quality – It might be the case that the added value of distributed simulation is countered by shortcomings on the quality of the resulting model.

2.1 Monetary Cost

When applying distributed simulation two types of direct monetary costs are involved. The practitioner has to acquire a tool using which simulation models can interoperate, and the COTS vendors have to spend effort, and thus money, to interface their packages to such a tool. For instance, if one applies HLA, the interoperation tool would be the Run Time Infrastructure (RTI).

Two experts, both from the research area, argue that the scarce use of HLA in industry might be caused by the purchase price of the architecture. One of them states that “I do not believe that HLA has a place, purely on the cost basis. The RTI costs too much and the customers are willing to pay at most 10% for an additional distributed simulation feature in a package - the vendors do not want to take steps in this direction”. Indeed, the purchase price of the architecture might be too high for academic researchers and it certainly negatively influences the cost benefit ratio.

There are counterarguments however. In spite of the fact that the DMSO RTI is a free solution that exists since 1996 almost none of the COTS simulation package vendors applies it as an additional feature for their customers. This suggests that the purchase price is not the main reason for not applying distributed simulation. Furthermore, two experts stated that even for a commercially available RTI the cost of these architectures currently is not an issue, because a 3000 euro RTI license is not so much in comparison with some of the simulation packages, which cost 30.000 euro.

If the price for an interoperability tool is not the issue, the cost may be hidden in the effort needed to integrate an existing distributed architecture in the COTS packages. Effort is needed here, even more so if one tries to hide the concepts of distributed simulation from the simulation practitioners. Currently, there is no compact tool that can be easily used to integrate simulation models from arbitrary packages, without any adjustment to the packages or the models.

2.2 Time

Distributed modelling in general is more complex than building a monolithic model, it involves additional concepts, such as operating the model at run time, synchronizing time, representing the data to be exchanged, exchanging the data, and managing ownership of an object, concepts with which the regular modeller is not familiar. In the monolithic style the simulation practitioner either does not need to take care of these issues or it is much easier. Due to this additional complexity more time is needed for design and development. Moreover, when applying a tool, like HLA, the user also has to get acquainted with the way the tool implements these concepts.

In particular, much time and effort has to be spent on the problem of the semantic inconsistencies that might occur during interoperation. As many experts pointed out, the unsolved problem of *semantic interoperability* when aligning the data models is the main obstacle that hinders developers in applying distributed simulation. According to one of the experts "Semantic interoperability is a big issue. This is a hard problem, and probably much harder than the problem for which HLA was designed and developed". Semantic interoperability is an issue for the whole system engineering community as well (Boer and Verbraeck 2003).

Furthermore, good programming knowledge alone is not sufficient to solve the problems related to distributed simulation. The interview with experts clarified that more is needed, because, as one of the experts argues, even if people "know the programming level good enough, they need to understand system architectures and how to create models of the world around them". One of the experts stated that applying the HLA standard requires high level thinking people because they need a detailed understanding of its concepts. This observation applies for distributed simulation as well. Most of the experts, both from defence and industry, claim that currently the simulation practitioners who daily design and develop simulation models using high level COTS simulation packages do not possess the knowledge that is necessary for creating distributed simulation. It requires a lot of effort to learn how to deal with it. The time spent on this entails hidden costs which might be considered too high in relation to the expected benefits.

Development time is also consumed because COTS packages do not provide high level building blocks implementing distributed system concepts. In general, using COTS packages the simulation practitioner does not need to go into low level technical details, they have been made invisible. As soon as distributed simulation comes into play, this is not the case any more, because for this issue COTS packages do not offer a similar high level transparent solution. As it stands now, in order to integrate distributed simulation models the simulation practitioner is

forced to consider technical details residing at a lower level than the standard predefined building blocks.

2.3 Quality

Next to the monetary costs and the costs related to longer development and learning times, the quality of the solution might hold people back to apply distributed simulation. Although distributed simulation can provide added value, according to some experts a distributed solution often turns out to be qualitatively deficient.

For instance, the end product may suffer from a bad performance. A monolithic simulation model implemented within one package performs faster than a collection of coupled models, because there is no need for explicit time synchronization, data exchange, etc., between the submodels (We explicitly exclude the case where a distributed model is set up to use parallel simulation techniques for speeding up the simulation run). One of the experts thinks that in this respect distributed simulation in general should be improved. On the other hand we observe that distributed simulation theory has already brought about efficient algorithms, some of them being discussed in (Fujimoto 2000). However they need to be implemented.

Apart from performance issues, also verification and validation of distributed systems is more challenging than in the monolithic case (Page, Canova, and Tufarolo 1997), and sometimes it is not even solved. So, choosing a solution based on distributed simulation could lead to a loss of quality with respect to these issues.

3 ARGUMENTS RELATED TO THE HLA STANDARD IN INDUSTRY

In order to hide the complexity of distributed simulation, the simulation community in defence designed and developed the HLA standard. One of the most important benefits of HLA, as some of the experts claimed in the interview, is that it solves the low level technical distributiveness details for the simulation practitioner. Still HLA does not catch on in industry. Analyzing the questionnaire and interview results we identify three possible reasons for this phenomenon:

- The HLA standard is too complex for industry
- The HLA standard is not transparent enough for industry
- The current implementations of the HLA standard, especially the DMSO RTI, are too inefficient for industry.

3.1 The HLA Standard is too Complex for Industry

Although the HLA standard was intended to hide the inherent complexity of distributed simulation, it seems that it did not succeed to hide it completely because, as most of

the experts both from defence and industry argue, the learning curve of applying this approach is very steep. Some of the experts from industry state that it is quite difficult to use HLA even when just trying to build some very simple examples. One of the experts argues, that “too much effort is required on behalf of a simulation practitioner, who may not be a technical person”.

One aspect of this complexity was described as follows: “the most difficult was getting used to the HLA style of thinking and programming. The complexity of the HLA interface has been an issue and developing low level task like building your attribute updates, and collecting the bytes and bits”. Apparently the HLA interface is too technical and resides on too low a level.

Another problem is the size of HLA. The HLA standard intends to solve problems for “all kinds of distributed simulations that might come into people’s mind. In order to cover so many issues, it has become quite complex”. This feature is also mentioned by one of the experts involved in the design and development of an HLA RTI who argues that “HLA provides all functionality that one can imagine for distributed simulation and maybe the weakness of HLA is that it does so much”. A COTS vendor claims that “it seems like an interesting idea but probably more attractive if you are in a large collaborative environment, possibly such as supply chain, possibly such as military”.

A common opinion is that industry does not need all functionality that HLA provides. One expert argues that “the complete HLA structure is a large structure, so normally simulation tools only need small part of it.” Another opinion: “HLA incorporates a lot of functionality that may not be useful to everybody”. An expert involved in HLA design and development argues that “something that you can do is changing the HLA specification, I believe. You could say, here is a much more limited definition of interoperability and that might make it easier for the simulation developers”. Another expert has the same opinion, he thinks that “for particular branches of industry the requirements are much narrower so probably a subset of HLA suits better, optimized for their particular purpose”. We have “to find out what kinds of functionality they need and use only those”. If the functionality which is currently not useful could be hidden or eliminated in some way the architecture would become more perspicuous, it would certainly flatten the learning curve and stimulate the applicability of HLA in industry.

3.2 The HLA Standard is not Transparent enough for Industry

Some of the experts complain that HLA is not a transparent solution. Before starting the discussion around transparency we like to mention another expert’s comment on this point, namely “I think what industry is saying is that it is

too complex”. So basically this reason relates to the previous point.

Regarding the transparency of the architecture an interesting difference in viewpoint between experts from defence and from industry surfaced. Simulation practitioners from the defence community complain that the HLA standard is not transparent in the sense that they cannot perceive what is happening in the RTI implementation. In order to solve this problem MÄK designed a tool called RTI Spy which sits on the top of the RTI functionality. This allows the modeller to analyze and debug, to help federation developers to find out what is going on, where the connectivity problems are, who is not making a right RTI call, etc. In contrast to modellers from defence, simulation practitioners from industry complain that HLA is not transparent in the sense that they are confronted with too many functions that are not relevant for them. It seems that, while for defence transparency means *visibility* of the underlying layers, for the industry it means *invisibility*.

In information technology the concept of transparency has the latter meaning, viz. *invisibility* or *undetectability*. An implementation layer is said to be transparent when the user is, or could be, unaware of it. For instance, in (Tanenbaum 1995) transparency refers to the extent distributiveness is hidden from the user. Transparency is an especially desirable feature in situations where users who are not particularly technically inclined would tend to be confused by seeing or having to interact directly with low level programming components. This is the case for most simulation practitioners, who are experts in COTS simulation packages but not necessarily familiar with low level technical details. So basically what industry needs is a transparent or invisible architecture that helps simulation practitioners to integrate different simulation models without requiring distributed system knowledge.

Transparency might be achieved in several ways. Some experts suggest to include distributiveness (or HLA-) concepts directly into the modelling paradigm of a COTS simulation tool. Other experts argue that implementing the concept of HLA in a commercial tool is not beneficial for the end user, because then the end user in fact again has to deal with HLA and the HLA mechanisms, thus defying transparency.

Other experts propose to achieve transparency by using adaptors or middleware, high level building blocks on the same level as the traditional COTS building blocks to handle distributiveness. We have been made acquainted with advanced HLA adaptors or middleware for COTS simulation packages, designed and developed for Arena, Simple++ and Taylor ED in the IMS MISSION project (McLean and Riddick 2000; Rabe, Jaekel, and De Gurtubai 2001), for SLX and Quest at Magdenburg University (Straßburger 2001), and for Arena at Genoa University (Revetria, Blomjous, and Van Houten 2003). One of the experts gave another argument that this is a better solution

because with complete integration “you save the adaptor, but you are even less open than when having the adaptor as a separate tool. You could even run the same commercial system with 2 or 3 types of adaptors if it is necessary and if the adaptor is inside the tool you are limited to one adaptor type”. He underpins his point of view by arguing that “the RTI’s which are now on the market are not really compatible. That means if you exchange the RTI you have to change the adaptor a little bit. So if the adaptor is integrated in the commercial tool you are limited to one specific commercial RTP”. Using intermediate flexible adaptors we have more chance for compatibility.

Although these are improvements, many people express the opinion that the transparency problem is not yet satisfactorily solved by adaptors or middleware as available today. One of the experts who developed adaptors himself stated “the systems available today for distributed simulation are all in a prototypic status. I think our adaptor, even if I would say that is near a product, I would still call it somehow prototypic”. Another expert claims that the current adaptors or middleware are “workable solutions”, which prove the concept but they are not end results. One of the COTS simulation package vendors similarly believes that even when having adaptors distributiveness is not completely invisible as “you have to write your own connection using HLA between two different simulation programs and this seems to be really hard coded. So the soft coding, the easy connection is not what I currently see”. According to him currently the only people who can easily deal with integrating COTS simulation models through adaptors to HLA are those people who designed and developed the adaptors. Many experts agree that we need clearer standards for interfacing COTS simulation packages.

3.3 The Current Implementations of the HLA Standard, especially the DMSO RTI, are too Inefficient for Industry

Another issue that hampers the acceptance of the HLA standard in industry is the perceived performance problems of the currently available HLA RTI implementations. Almost all experts who applied the HLA standard made use of the DMSO RTI, so first of all we focus on the performance of this version.

The opinions of the experts in defence are mixed, there is a group who is in general satisfied, while there is another group that thinks that performance should certainly be improved. For the experts from industry the opinion is more univocal, all experts from this area complain about the performance of the DMSO RTI.

Experts from defence came up with applications where, although early versions of the DMSO RTI indeed showed performance problems, the performance of the newer DMSO versions was satisfactory. One expert claims

that “using the DMSO RTI we have supported simulations with tens of thousands of entities”. We also obtained positive feedback from the HLA vendors: “the DMSO RTI now is suitable for large federations”. Another defence simulation practitioner states that he had experience with training related distributed simulation where the RTI performed satisfactory.

Some defence experts criticized the negative opinions on the DMSO RTI. According to one of the HLA designers and developers “people who complain make a lot more noise than people who do not complain”. Both HLA RTI designers and developers argue that sometimes people blame the performance of the RTI when this bad performance might arise from inefficient design and development of the simulation models or inefficiencies of the network. According to one of these experts these complaints stem from the fact that “RTI is hiding the network from you, so if you do not have enough network capability, then it appears that the RTI is slow, and you say that HLA is slow”.

But negative judgments could be heard from practitioners in defence as well. One defence expert claims “the DMSO RTI had a pretty poor performance and very little control over the performance”. An HLA designer pointed to cases where he can imagine that the DMSO RTI imposes latencies which slow down the whole simulation, for example if somebody is doing visualization and frequent refresh of the display is required.

From the industrial side the reports are negative without exception. One expert complains that it is very difficult to get an acceptable performance when coupling COTS simulation models. Another expert was confronted with an unsatisfactory performance when he coupled two Arena simulation models. He ascribes the problem partially to the performance of the DMSO RTI, and partially to the huge overhead due to time synchronization. He mentioned that changing to another Arena version helped. A COTS simulation package vendor reports that they have done distributed simulation using Extend for aerospace companies simulating large scale communication system, and as he states “HLA was rejected for these projects because of performance issues”. An expert, whose activity at SimTech was to research on the performance of distributed simulation in general, has explicitly analyzed the performance of the DMSO RTI. His conclusion was that “there are bottlenecks in terms of long delay when a federate intends to join a federation and there is also a lot of overhead in terms of broadcasting messages or point to point message sending. So you get the feeling that the implementation is inefficient”.

In view of the diverging opinions it is difficult to draw a clear conclusion about performance. However for industry it seems that there is a need for an approach which performs better than the DMSO RTI. That is where the HLA vendors stepped in. The RTI that DMSO produced, as one expert argues, was designed and developed to be a general

purpose architecture, it was supposed to implement all services possible. DMSO's vision on how the market would evolve was that at some point the full RTI will be taken over by the commercial market and the commercial market would produce RTI's for specific user domains with a performance satisfactory for that community. Actually this is the idea behind the RTI that was produced by MÄK Technologies. Its primary aim was to serve the real-time defence community. Both HLA vendors that we interviewed claim that their HLA implementation, Pitch RTI and MÄK RTI respectively, performs better than the DMSO RTI. The differences between the performance of the DMSO RTI and the MÄK RTI is described in (Burks et al. 2002). Not only the vendors but other experts as well declare that the commercial RTI's perform better than the one from DMSO. However, as a defence expert argues "it still should be improved".

Another idea is subsetting the HLA standard. One of the experts claims that this is beneficial "you can gain efficiency by trimming the specification". At GeorgiaTech experiments have been conducted with a research version of RTI, called RTI-Kit, which had less functionality than the DMSO RTI and performed better. As the expert involved in this research states "DMSO argues that it does not matter how much functionality one is using, the performance will remain the same. My opinion diverges, if you use less functionality you will get a better performance". More experts voice this opinion, they think that there is functionality that not everybody uses and needs which costs a lot of performance. Earlier we found that there is need for a simpler solution than the HLA standard providing only the functionality that is currently needed for industry. The above remarks lead us to expect that the implementation of such a less complex solution will perform better than the RTI's that are currently available.

4 ARGUMENTS RELATED TO THE RELATION BETWEEN HLA AND COTS

The arguments discussed above were related to features of HLA. In this section we focus on how COTS simulation packages support HLA. Our starting point is the observation, validated in the survey, that hardly any COTS simulation package allows appropriate creation of HLA compliant simulation models.

A confusing outcome of the questionnaire was that 11 out of the 18 COTS vendors stated that models built in their package had been combined through HLA, and 11 vendors that either their package supports HLA as a feature or will support HLA in the future. This generated the need for a definition of the notion HLA compliance of a COTS package. In the interview survey we interrogated the experts on this issue.

A preliminary remark was made by some experts who stated that, strictly speaking, "because HLA compliance is

always bound to a certain federate and to a certain simulation object model, you cannot say that a tool is HLA compliant, but only a federate is HLA compliant". Based on this observation, we might define a package to be HLA compliant, if it provides a way to create HLA compliant models.

This definition is somewhat vacuous though. Interpreted this way, the COTS vendors have a point in claiming that their package is HLA compliant, because HLA compliant models can be built by using some low level communication protocol, like WinSock, provided in the package. In a follow up of the questionnaire we obtained evidence that although there seems to be active interest in HLA by about half of the COTS vendors, this interest is rather tentative and aiming at low level solutions. In the interview the experts agreed that we can only speak of "HLA compliance" if the package enables the model builder to construct an HLA compliant model in a relatively easy way. Opinions differed between the experts on what "relatively easy" means.

One group of experts interprets this on an intermediate level. According to one of these experts an HLA compliant COTS simulation package should "actually have an HLA interface and have a well-defined FOM or SOM that you support". Another expert has more or less the same opinion, he states that "the simulation package should have built in HLA interfaces". An HLA designer and developer states that "in order to be HLA compliant you would need to supply to the COTS simulation package a mapping between the things that are in your simulation, objects within the simulation, you need to be able to tell the package which ones to expose to the HLA federation, and how should they be exposed, how do your things inside the simulation map to the FOM". The same expert then proceeds with stating that "unfortunately it may not be that easy, it may be a long way towards hiding the details of HLA".

Some of the experts suggested a definition of HLA compliance on a higher level "it is important that we have a specific system which allows the automatic adaptation of the model into a specific federation". Common opinion, validated by all experts, is that hardly any COTS simulation package provides interfaces that enable a user to create HLA compliant simulation models with small effort.

A complicating factor is that there are concepts in the HLA standard which currently are not implemented in COTS simulation packages, for instance ownership. One expert who has experience on this issue claims that ownership management is "a concept that is not really common to these tools. So you have to find a way of mapping the logic to the mechanism of the package". Another expert agrees "you cannot implement it within the basic structure of the COTS simulator without going into the internals of the simulator and then it is no longer a COTS simulation package".

HLA compliance is severely hampered by the closedness of the COTS packages. Even assuming that a simulation practitioner exists with high level programming knowledge and deep knowledge about HLA, who wants to access HLA from a COTS package, his task is still not simple due to the fact that COTS simulation packages do not allow access to internal variables which are essential for distributed simulation. Such a variable is, for example, the event calendar which is needed in time scheduling. According to one expert “you need to understand the internal scheduling system of the simulation system, because this interferes heavily with the HLA RTI”.

From the interview results we conclude that the COTS vendors are not willing to open up their packages. Most of the vendors want to bind their customers by providing more functionality in their own package, instead of integrating it with other packages. For example, an HLA RTI vendor argues that “when COTS simulation tool vendors develop these products and add modules for more and more types of simulations there maybe an interest in keeping the users inside their own architecture and not interoperate with other types of architectures. They do not want you to connect to competitors, so they want you to stay with the same simulation tool and the simulation tool should cover more and more functionality, which leads to a customer lock situation”. An HLA RTI designer and developer has the same opinion, he notices that “if you are the developer of a COTS simulation package or general framework, you are trying to argue that your package is extremely general purpose, does everything under the sun, it has all the capability, so this makes it hard for you to envision how simulations built in your package are going to deal with other simulations across the HLA”. An expert from defence claims that “there is no incentive for different vendors of different simulation packages to agree on a common interoperability standard”. An expert from industry observes that COTS simulation package vendors do not want to take too much initiative because “simulation vendors are just trying to avoid that somebody uses other systems than theirs”. A COTS vendor states that “Why would we encourage our users to integrate our products with our competitor’s products? That is the basic question you have to ask. What is the economic reason for working with a competitor? So simple as that”.

Another reason why COTS vendors are careful in investing in HLA is the cost benefit ratio. One COTS vendor did a market research regarding HLA support and he realized that “there is no money in it. There is no organization who wants to pay for it. We do not see a market where today the small or middle size companies would pay for it”. Another representative of a COTS vendor argues “I can only think that there has not yet been a convincing financial argument for why any particular simulation vendor should modify their product or their documentation to explain to their users why is in their benefit to talk to HLA”.

COTS vendors do not necessarily exclude themselves from supporting distributed simulation, however they do not want to invest too much in it, because the benefit is not clear. “If we could find a way that it would help increase our sales as well as theirs, perhaps, then you might have a business case for it, but otherwise there is no financial motivation to do that”. Both COTS vendors from the interview explicitly stated that they are open and interested to support any distributed simulation standard, and if it is the case and they see benefit they will build alliances with other vendors, like Brooks Automation did with Simul8.

Finally, we should not just blame the COTS vendors, we should also take into account the complexity of the HLA standard and the domain for which it was designed and developed. An expert from an organization that is quite willing to open up its simulation package for distributed simulation purposes, claims that “HLA is very military specific and so is weighted down by support for features not required in many cases”. Another expert from industry thinks that “when defence designed HLA they did not have in mind any COTS simulation package”. His observation is supported by an expert from defence who states that the defence simulation community does not focus on COTS simulation packages because “a lot of these big military simulations are really running on the edge of what can be done from a performance standpoint, and you cannot tune your system unless you own it all”. Regarding these two groups, COTS and HLA, we are led to agree with one of the expert’s observation, who states that “the developers of these big simulation systems are just suspicious of the COTS simulation environment, and the COTS simulation environment people are suspicious of the HLA.”

5 CONCLUSION

A recurring theme in this paper is that the main reason why industrial practitioners do not apply distributed simulation is the perceived cost benefit ratio. Although one of the aims of HLA was to reduce complexity and thus cost, we found that industry still perceives high costs. The problem is not so much HLA specific, but more general and relates in fact to the reluctance in industry to use distributed simulation.

The cost benefit comparison however, would not be fair if we would focus only on the cost side and not on the benefit side as well. For this reason, we should examine whether simulation practitioners are fully aware of the benefits. In (Boer 2005) we pointed out in that this is not the case.

The invisibility of the benefits, according to a COTS vendor, relates to the business model currently applied, namely designing and developing “throw away” simulation models. The primary aim of industry as most experts state is to get a solution as fast and as cheap as possible and they have a tool, COTS simulation packages, to achieve this

aim. However, these tools are designed and developed with the above business model in mind, they are not supporting design and development of distributed simulation in a proper way.

Very often added value can only be discovered through trial or experimenting with a feasible prototype. Simulation practitioners who are dependent on COTS packages do not have the possibility to carry out distributed simulation, they do not have a chance to do so, and benefits of distributed simulation that they might discover remains invisible. On the other hand, for the few simulation practitioners who do perceive benefits of distributed simulation the inherent complexity of distributed design and development and the absence of tools which might handle this complexity lead to the fact that they scarcely apply it.

Neither practitioners nor vendors are willing to take a risk to invest in invisible benefits. This is reflected by the quote in Section 2.1 that COTS customers are not willing to pay more than an additional 10% for distributiveness to be included. This reflects a minor interest. The question here is: who is going to take the first step, the vendors or the practitioners?

Basically we experience a deadlock, a chicken and egg scenario. On the one hand simulation practitioners do not see the benefits of distributed simulation because the tools they are used to are not appropriate to experiment with distributed simulation and therefore they do not request from the tool vendor an upgrade that supports distributed simulation. On the other hand COTS vendors do not provide a tool because the end users do not request it.

Nevertheless, time seems to be ripe: “industry does not yet understand that this is a topic of the future”, as an expert said. One of the COTS vendors gave a similar comment, he argues that distributed simulation is a “quite large area, but remains as yet unexploited”. According to him it is a matter of marketing. There is room for distributed simulation and for protocols that implement the concept. The most important need is to create a market for it. By providing easy and appropriate technical solutions for end users who can then play with it, a market will be created. So basically a technology push will create a market pull.

In (Boer 2005) we responded to this challenge. Based on the reactions from the survey on our third question (“What is needed to overcome the deadlock situation?”) we developed a requirement list for a light-weight COTS based distributed simulation architecture on which we based our “FAMAS Simulation Backbone Architecture”. This will be the subject of a forthcoming paper.

ACKNOWLEDGMENTS

We would like to thank the Commercial-Off-The-Shelf simulation package vendors for participating in the ques-

tionnaire survey and the simulation experts for the useful discussions during the interview survey.

REFERENCES

- Boer, C. A. and A. Verbraeck. 2003. Distributed Simulation With COTS Simulation Packages. In *Proceedings of the 2003 Winter Simulation Conference*, ed. S. Chick, P. J. Sanchez, D. Ferrin and D. J. Morrice, 829-837. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.
- Boer, C. A. 2005. *Distributed Simulation in Industry*. ERIM Ph.D. Series Research in Management: Rotterdam, The Netherlands. Available via <http://hdl.handle.net/1765/6925> [accessed June 29, 2008].
- Boer, C. A., A. de Bruin, and A. Verbraeck. 2006a. Distributed Simulation in Industry – a Survey, Part 1 – The COTS Vendors. In *Proceedings of the 2006 Winter Simulation Conference*, ed. L. F. Perrone, F. P. Wieland, J. Liu, B. G. Lawson, D. M. Nicol, and R. M. Fujimoto, 1053-1060. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.
- Boer, C. A., A. de Bruin, and A. Verbraeck. 2006b. Distributed Simulation in Industry – a Survey, Part 2 – Experts on Distributed Simulation. In *Proceedings of the 2006 Winter Simulation Conference*, ed. L. F. Perrone, F. P. Wieland, J. Liu, B. G. Lawson, D. M. Nicol, and R. M. Fujimoto, 1061-1068. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.
- Boer, C. A., A. de Bruin, and A. Verbraeck. 2008. A Survey on Distributed Simulation in Industry. Accepted for publication in *Journal of Simulation*.
- Burks, T., T. Alexander, K. Lessmann, and K. G. LeSueur. 2002. Latency Performance of Various HLA RTI Implementations, Technical Report, White paper.
- Fujimoto, R. M. 2000. *Parallel and Distributed Simulation Systems*, New York, USA: John Wiley and Sons, Inc.
- McLean, C. and F. Riddick. 2000. The IMS MISSION Architecture for Distributed Manufacturing Simulation. In *Proceedings of the 2000 Winter Simulation Conference*, ed. J. A. Joines, R. R. Barton, K. Kang, and P. A. Fishwick, 1539 - 1548. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.
- Page, E. H., B. S. Canova and J. A. Tufarolo. 1997. A Case Study of Verification, Validation and Accreditation for Advanced Distributed Simulation, *ACM Transactions on Modeling and Computer Simulation*, 7(3), pp. 393-424.
- Rabe, M., F. W. Jaekel and G. G. De Gurtubai. 2001. Modelling and Simulation for Globally Distributed Enterprises. In *Proceedings of the 4th International EUROSIM 2001 Congress*, ed. A. Heemink, L. Dek-

- ker, H. de Swaan Arons, I. Smit and T van Stijn. Delft, The Netherlands.
- Revetria, R., P. E. J. N. Blomjous and S. P. A. van Houten. 2003. An HLA Federation for Evaluating Multi-Drop Strategies in Logistics. In *Proceedings of European Simulation Symposium and Exhibition Conference*, ed. A. Verbraeck and V. Hlupic, 450-455. Delft, The Netherlands.
- Straßburger, S. 2001. *Distributed Simulation Based on the High Level Architecture in Civilian Application Domains*, PhD thesis, University Otto-von-Guericke: Magdeburg, Germany.
- Straßburger, S., T. Schulze and R. M. Fujimoto, 2008. Future Trends in Distributed Simulation and Distributed Virtual Environments, Peer study report. Available via <http://www.sisostds.org> [accessed June 29, 2008].
- Tanenbaum, A. S. 1995. *Distributed Operating Systems*, New Jersey, USA: Prentice-Hall International, Inc.

AUTHOR BIOGRAPHIES

CSABA A. BOER is a senior product manager at TBA BV, one of the leading logistics and simulation consultancy firms in Europe. He holds a Ph.D. in Computer Science and Logistics from Erasmus University Rotterdam. His research interests include distributed simulation, distributed virtual environments, port logistics, and port simulation and emulation. His e-mail address is csaba.boer@tba.nl.

ARIE DE BRUIN is a full professor at Delft University of Technology. He got his Ph.D. on semantics of programming languages and since then he has specialized in investigating implementation issues from a theoretician's point of view. One of his current research interest is distributed simulation, most notably the essential concepts thereof, and the soundness of the implementation of these concepts. His e-mail address is a.debruin@tudelft.nl.

ALEXANDER VERBRAECK is a full professor in the Systems Engineering Group of the Faculty of Technology, Policy and Management of Delft University of Technology, and a part-time full professor in supply chain management at the R.H. Smith School of Business of the University of Maryland. He is a specialist in discrete event simulation for real-time control of complex transportation systems and for modeling business systems. His current research focus is on development of open and generic libraries of object oriented simulation building blocks in Java. His e-mail address is alexandv@tbm.tudelft.nl.