14th ICCRTS: C2 and Agility

"Utilizing Resource Brokering Within Virtual Environments to Support Distributed Collaboration and Rapid Team Configuration"

Paper 086

Topic 10: Collaborative Technologies for Network-Centric Operations

authors:
Robert Holder and R. Scott Cost
Johns Hopkins University Applied Physics Laboratory

poc: Robert Holder
Johns Hopkins University Applied Physics Laboratory
11100 Johns Hopkins Road, #2-236
Laurel, MD 20723

240.228.3263
Robert.Holder@jhuapl.edu

Utilizing Resource Brokering Within Virtual Environments to Support Distributed Collaboration and Rapid Team Configuration

Robert Holder and R. Scott Cost
Johns Hopkins University Applied Physics Laboratory
{robert.holder, scott.cost}@jhuapl.edu

Abstract

The emergence of virtual environments, in the early form of social networks, and recent, more elaborate incarnations such as Second Life and the Global Information Grid (GIG), create opportunities for new means of collaborating that negate the typical constraints of physical environments. Combined with previous work demonstrating the semi-automated formation of collaborative teams, this represents a powerful advance in information exchange and collaboration capability. However, the inherent differences between traditional physical environments and these new virtual environments require that old methods of evaluation, acquisition, and integration of resources into collaborations be updated. We introduce a scenario and a prototype implementation to illustrate the different opportunities and challenges of finding resources for use by virtual collaborations, and suggest directions for modifying the strategies used in the physical environments to be better suited for use in virtual environments.

1 Introduction

Traditional real-world collaborations, such as setting up an emergency operations center, typically locate potential resources through database accesses, registries, word of mouth, web searches, and other means. These resources include assets such as human personnel, equipment, and software. Potential resources are evaluated using metadata such the availability of a human expert, the platform requirement of a software application, or equipment specifications. After evaluation, the best set of resources are requested for acquisition.[1] Assuming a resource is actually available, it is integrated into the collaboration. At this point, one hopes that the resource is able to supply the necessary functionality. However, unexpressed constraints may make the resource unsuited to the

[1] This assumes a scenario in which all resources are acquired at one time. Other scenarios may seek to acquire one resource, and then reevaluate before seeking others.

collaboration environment, or the metadata describing the resource may be inaccurate or outdated. In either case, the collaborators must perform the resource evaluation, acquisition, and integration again. Repeated iterations can easily become expensive and delay progress.

Even if these steps are only performed once, a complex set of resource requirements or need for a large number of resources can cause the process to be slow or tedious. To ease this burden, a collaboration may employ an automated tool to facilitate the required steps. A *resource broker* (RB) retrieves and evaluates resource metadata and determines a suggested set of resources for use within a collaboration. This process contains two steps. First, an RB retrieves a set of metadata by executing a set of queries. This set of queries considers trade-offs between response time and accuracy of information. Metadata about resources, such as availability, may be accessed through various means such as a database or Outlook calendar, a web service, or directly from a human, and will have varying degrees of accuracy and query response time.

For example, consider an RB that may query either a human scheduling assistant or an Outlook calendar for an analyst's availability to provide expertise to a collaboration. An RB may expect a response from the assistant to be slower but more current than the Outlook calendar. If the primary consideration is speed, then the RB will query the Outlook calendar and accept the potential uncertainty. If is it later discovered that the resource is not available as indicated, the RB must make additional queries and requests to replace resources. On the other hand, if the analyst's presence is particularly critical to the collaboration, then it may be best for the RB to accept the delay and wait for the assistant in order to receive the most accurate information.

Once the RB has acquired the relevant metadata, it determines a set of resources that satisfy the requirements, again taking into account the potential uncertainty. In this step, the RB considers trade-offs between resources that may imperfect matches for the collaboration requirements: a resource may suit the collaboration in terms of one attribute, but not align well with the collaboration in terms of another attribute. Continuing with our previous example, let us assume that the RB has metadata describing the human analyst's availability, as well as a comparable automated tool available as a web service. The automated resource would be an ideal fit for the collaboration in terms of availability, and would be preferred over the human analyst. This is due to the human's less certain availability stemming from potential inaccuracy in the metadata, and also unforeseen elements such as illness. On the other hand, we can assume that the human analyst will be more likely to adapt to unforeseen or ambiguous collaboration needs. Thus, our RB must determine how confident it is in the precision of the match between the collaboration's needs and the web service's functionality, and also weigh the trade-off between the web service's ideal availability and the possibility that it will not meet the collaboration's functionality requirements.

The Johns Hopkins University Applied Physics Laboratory (JHU/APL) has developed a prototype RB, the *Dynamic Collaboration Action Team Resource Broker* (DCAT-RB), which is designed to determine a resource set for use in various situations such as flood disaster recovery or a chemical biological attack [Cost *et al.*, 2006; Salamacha *et al.*, 2006]. The set of resources that the DCAT-RB considers are hazardous material experts, law enforcement officials, and analysts. The DCAT-RB evaluates metadata that describes availability, years of experience, type of hazardous material expertise, and language capabilities.

Recent trends in technology have led to increasing use of virtual environments, such as Second Life, in which users can collaborate virtually. These environments offer many advantages such as convenience, flexibility, and cost. However, these virtual collaborations are often not constrained by the same basic premises governing traditional collaborations, such as the constraint that a resource is limited to one location at any given time. As the use of virtual environments becomes more pervasive, we must examine these differences between physical and virtual environments, and then determine how the differences impact the type of metadata we require for resources, and how they impact the use of a traditional RB built for use in physical environments. Below, we discuss differences between the physical and virtual environments that impact an RB, and offer an illustrative scenario.

2 What is a virtual world?

"Virtual" refers to functionality that exists outside the confines of a physical embodiment [Igbaria, 1999]. Wikipedia defines a virtual world as "a computer-based simulated environment intended for its users to inhabit and interact via avatars."[2] Virtual Worlds Review (VWR) defines a virtual world as "an interactive simulated environment accessed by multiple users through an online interface."[3] VWR also defines the common characteristics of virtual worlds as having a *shared space* for multiple-user participation, being *GUI-based*, supporting *immediacy* by having real-time interaction, *interactivity* between users, *persistence* of the world independent of individual users being logged on, and encouragement of *socialization and community.*

The most prevalent use of virtual worlds is for interactive games. Some have specific themes, such as World of Warcraft. Other virtual world games, such as Second Life or The Sims Online, are more general and are meant for more creative, open-ended interaction.[4] Although these specific types of utilizations are new, the use of virtual environment technology dates back to 1985 in the form of text-based chatrooms such as The Well [Babcock, 2007]. Current trends point to the expansion of this technology to applications in education, military training, and simulation.[5]

For our purposes, we will adopt the definition used by VWR, but relax the requirement that a virtual world includes a GUI interface. We are particularly interested in the identification and use of resources within collaborations in a virtual world. One could define a virtual environment as the more general construct, of which the virtual world is an implementation. However, it the context of this writing, the distinction is usually not relevant, and we use the terms interchangeably.

Virtual worlds are an increasing means by which the DoD may facilitate effective collaborations for

[2] http://en.wikipedia.org/wiki/Virtual_world

[3] http://www.virtualworldsreview.com/info/whatis.shtml

[4] Camp gives an in-depth discussion on "structured" versus "unstructured" games [Camp, 2007].

[5] One interesting simulation application includes experiments which explore the effects of different social norms or political structures on a community.

C2 operations. Virtual worlds offer advantages over the traditional communication mediums such as telephone, text, and videoconferencing. While valuable tools, research shows that teams that rely on these forms of communication miss important communication cues such as body language and facial expressions. Also, in collaborative teams that include both remote and local members, the remote team members many times experience a sense of being excluded. As recent research shows, there is a positive correlation between social cohesion and small team performance [Huber *et al.*, 2007], and thus these effects on communication and feelings of inclusiveness can have a direct impact on the productivity of a collaboration. Other tools currently used for distributed collaborations, such as file sharing and whiteboards, have a limited effectiveness in their ability to convey complex ideas [Kuper and Giurelli, 2007]. Virtual worlds offer a solution to these shortcomings by providing a three-dimensional collaborative environment in which all team members participate. A room in this virtual world can include operational models, whiteboards, powerpoint presentations, documents, and most other collaborative needs. As an avatar in a virtual world, collaborators are in an environment with other collaborators' avatar representations, and have the ability to move around, look, gesture, look at documents and presentations, use analytical tools, and simulate many aspects of the richness one expects from real-world interactions.

In addition to providing virtual versions of typical collaborative facilities, virtual worlds offer features typically not available in real-world collaborations. For example, research shows that parallel text-based communication is particularly beneficial for complex missions [Alberts, 2007]. In real-world collaborations, engaging in simultaneous conversations is difficult, and monitoring and contributing to relevant side conversations is even more difficult. This is particularly important given the gradual shift from classic C2 to "edge organization" [Alberts, 2007], in which information analysis and decision making functions are distributed throughout an organization instead of residing at a central point. The ability for the most appropriate team members to be involved in relevant discussions will be critical to a team's success. Virtual worlds typically offer the ability to communicate both in voice and text medium. Also, because communications are mediated through the virtual world, it is not difficult for an avatar to be automatically notified about a conversation for which his or her expertise could be beneficial.

3 Resource allocation in virtual environments

When virtual collaborations require resources, the differences between virtual and physical environments may cause the traditional methods of resource evaluation, acquisition, and integration to be less effective. These differences can affect the ability to evaluate the resources for use in a collaboration, or they may affect the effectiveness of resources within the collaboration itself. Some variants are differences in connectivity, lack of stability of collaborations, increased diversity of environmental culture, external constraints to consider, boundaries between environments, and differences in security considerations.

3.1 More diverse environmental cultures

More so than collaborations based in the real world, collaborations in virtual environments tend to span economic, political, and policy boundaries [Igbaria *et al.*, 1999]. Given that virtual collaborations transcend typical physical constraints such a location, this is reasonably intuitive. Given the advantages of virtual worlds, it is not unreasonable that the people and resources involved in potential collaborations will increase in diversity. As a consequence, the potential scope of the required metadata about resources is larger, both in terms of the type and value range of the metadata. Examples can include differences in languages, customs, culture, formatting, units, vocabulary, or laws. Normally many of these items would be standardized as part of physical locale. For example, it would be less critical for two collaborators in the United States using a US-based tool to specify whether they expect data to be in US customary, metric, or Imperial units than a collaboration spanning the US and Europe. For an RB gathering resources for an international collaboration, it must have the ability to consider or even reason over additional types of metadata.

Resources may also find themselves being utilized in unexpected domains. The potential effect of this increases as resources become more elaborate.[6] Whether or not an RB can determine if a resource is suitable for an unexpected domain depends on the type of metadata available, and what the RB is able to infer. One might consider a case in which a resource that depends on access to a web service may find that a virtual collaboration has a firewall preventing the server from being accessed.

Virtual worlds offer the ability for an individual to effectively contribute to a collaboration without physically being present. This presents the opportunity for an expert to simultaneously lend guidance to many disparate collaborations. However, this means the collaborator must have the ability to quickly shift context to properly engage these different cultures.

3.2 Physical constraints on virtual resources

Ultimately, a virtual resource will be subject to physical constraints due to its servers being based in the physical world. Physical phenomena such as power outages, natural disasters, and similar issues present obvious physical constraints that impact the availability of any virtual resource. More subtle effects will impact a resource's speed or effectiveness. Constraints that originate from outside a resource's environment may be difficult for an RB to consider, or even to acquire.

Virtual resources, not being constrained by time and space, may need to account for the short- and long-term effectiveness of a resource that is utilized outside of "normal working hours" [Igbaria *et al.*, 1999]. For example, if a resource is a virtual interface for a human or human-maintained physical entity, then that resource's ability to effectively support the collaboration will be reduced outside of normal working hours.

[6] As a simple example, consider an ASCII-based email versus an HTML-based email that has an embedded picture. The text email will display on any email client. However, the more elaborate HTML email will become mangled if it is unexpectedly rendered by text-based email client.

Bandwidth is another physical constraint on a virtual resource. Two otherwise unrelated virtual resources may be supported by the same physical server. In general, any large load on a physical server will impact the performance of the virtual resource. Therefore, the manner in which a one resource is used may impact the effectiveness of another resource. In this case, there would appear to be a correlation between the use of two resources that would not be explainable from the perspective within the virtual environments.

These examples illustrate the challenge of accounting for uncertainty in resource capability when the cause of the uncertainty is not available through the resource's environment. We must be able to identify the relevant factors stemming from the physical world and be able to model their impact on resource capabilities. This may be as simple as decreasing a resource's capability by some percentage when some physical effect is active, or we may have to consider a more elaborate evaluation of both the current physical and virtual situation. In either case, we require an RB to consider a model that encompasses more than just its immediate environment.

Physical constraints impact the human participants as well. Collaborations should take advantage of the increased availability of a participant due to virtual technology, but must also still consider a human's physical and technological constraints. With the proper equipment and the necessary mental dexterity, one individual can contribute to several virtual collaborations at one time. However, the needs of a particular individual may affect the resources that should be acquired by the RB. For example, if a collaboration acquires a virtual visualization resource and several monitors, then a participant may be able to support more virtual collaborations than a participant that has to manage a local tool and rely on a single monitor.

3.3 Security and privacy

Addressing the issue of security within virtual environments requires us to reconsider the traditional paradigms. Traditionally, security can rely upon physical properties of nature. For example, barriers such as walls and locks are commonly used to control access, and we can soundproof rooms to foil potential eavesdroppers. In virtual worlds, however, physical properties are not enforced by natural law, but rather by the instructions within programs running on servers. In one sense, it may be easier to foil eavesdroppers in a world like Second Life by performing a programmatic search for all avatars within proximity of a location of interest. On the other hand, an adversary who compromises the server may have the ability to cause more disruption than an adversary who has defeated a single lock.

Given that a virtual world resides on one or more servers, everything that happens in a virtual environment, as well as anything that has *ever* happened is subject to inspection by anyone to whom the server's logs are available, or to anyone who can sniff the server's network connections. In addition to security, another immediate consequence relates to privacy. The ability to sift through server logs to discover what an avatar is doing, or what it has *ever* done, will be a benefit to some collaborations as an archive of its activities. However, sensitive activities, such as classified discussions or simulations of experimental technology, should not be readily available to all users.

Technologies such as encryption and private information retrieval (PIR) [Chor *et al.*, 1995] may offer solutions to these challenges.

In addition to securing the infrastructure with secure hardware and software, we must secure the functionality that runs on top of that infrastructure. Voting, financial transactions, secure communications, and other sensitive uses of virtual environments will require secure network protocols, avatar verification, validated script functionality, and other means to prove that a virtual entity represents itself honestly. In addition to avatar verification, virtual worlds must also consider how to enforce access protocols. For example, if a verified avatar is making a financial transaction with a bank based in Second Life, how is an arbitrary avatar prevented from teleporting or flying nearby and observing the transaction? Typically software is used to enforce this, but we cannot completely rely on software solutions. A recent reminder of this occurred when security researchers demonstrated a hole in the QuickTime software utilized by Second Life allows in-world objects to steal resources from and take control of other avatars. [7]

The sharing of information within virtual partnerships between business or other individuals must be regulated to prevent trade secrets from being distributed outside the company. This is particularly important because virtual partnerships may involve only portions of a company; as Igbaria [1999] notes, "two companies in partnership on one project may be bitter rivals in another." Thus, technology is required that allows the use of sensitive information in one virtual collaboration, but prevents that information from being exploited in another virtual collaboration.

As always, the human aspects of security represent the greatest risk. In particular, human participants are present in both the physical and corresponding virtual environment. Thus, there is always a bridge between worlds that must be managed.

3.4 Collaborations spanning multiple virtual environments

Until a standard for communication between virtual environments exists, collaborations spanning multiple virtual environments will be difficult. This poses a challenge both for user interactions and resource functionality. Although simple data such as a string can be transmitted via HTTP or other protocols, the bigger challenge is that a resource rendered in one virtual environment may not be accessible in another virtual environment. A resource built specifically for one virtual environment may find that its properties are not useful, applicable, or even available in another virtual environment. In general, one must consider that a user may need to use a resource in more than one virtual environment, in which case the difficulty is two-fold: how does one develop a resource that is deployable to arbitrary virtual environments, and how may a user transfer the state of a resource across virtual environments?

The challenge of virtual environment integration may be analogous to website integration in that individuals can build their own virtual worlds to support their unique needs, but this leads to

[7] http://www.mercextra.com/blogs/takahashi/2007/11/30/

difficulty when disparate virtual worlds need to interact. Multiverse Networks[8] starts to address this challenge by offering a free platform from which virtual worlds may be developed. Virtual worlds built upon this platform share a de facto standardization that will allow the exchange of information. The OpenSimulator Project[9] is working on an open source solution. Other organizations working toward virtual environment interoperability include Media Grid[10], Virtual World Interoperability[11], and Architecture Working Group[12]. Until these or other integration efforts reach maturity, the virtual worlds that do not share a common platform will have to find a way to exchange information, perhaps using middleware such as XML or CORBA, or web interoperability technologies such as OWL and web services.

3.5 Resources spanning both virtual and physical environments

During early explorations into the use of virtual environments, resources used within collaborations are likely to have both real-world and virtual world analogues. An RB that is considering both aspects for inclusion into a collaboration must compare the merits of two types of metadata. For example, a real-world resource is likely to have an attribute that describes its location. A virtual resource may have a similar attribute. However, the real-world location is likely to be a geospacial location, whereas the virtual world location is likely to be a URI or some other Internet-based location. For the RB to determine which resource is more suited to the needs of the collaboration, it must weigh the benefit of having a resource locally located, and perhaps more reliably available; or remotely located and widely available, but perhaps with higher possibility of losing access if the network connection suffers.

Another open question related to resources with dual representations is how to evaluate the cost versus benefit of moving those resources across boundaries between virtual and physical environments. For example, two physical entities, each of which has control over its own virtual avatar, may need to exchange a document. If a traditional text or binary document, then the entities may exchange the document either virtually or physically. The optimal means of exchange then depends on the current state of the document, the cost of transitioning the document to a new environment, and the recipient's preferred means of acquiring the document. This process may become non-trivial when the document is non-traditional. For example, a natively virtual document may contain functionality, such a audio or a three-dimensional component, that is not readily replicated in a physical environment. In other cases, a single very large virtual document may be supported by many servers residing in many different physical locations, making it more expensive to coordinate all server activity to produce a consolidated physical document. In the virtual environment, this cost is mitigated by querying the appropriate server only when a specific portion of the document is being used. There is no physical equivalent to this on-demand type of retrieval; either one

[8]http://www.multiverse.net/
[9]http://opensimulator.org/
[10]http://www.mediagrid.org/
[11]http://vwinterop.wikidot.com/
[12]http://wiki.secondlife.com/wiki/Architecture_Working_Group/

continuously has the complete document, or a less useful incomplete document.

3.6 Connectivity

Within virtual environments we rely on the assumption that information we need is linked. That is, we have the ability to easily access information from a related information source. Thus, it becomes easy to find information regarding a particular resource because there are many paths that reach it. The Internet is a prime example of this in terms of general information, and social networking is an example regarding information specific to people. However, this assumption of connectivity is not always valid. When modeling a social network, is it often more accurate to create a disjointed network than a fully connected one [de Weerdt *et al.*, 2007]. In many applications, matching resources to a collaboration tends to be a centralized process, with all necessary information available from a central source. Because of this centralization, the links to potential resources from peers may not be as necessary, creating a network that is not fully connected.

This lack of connectivity has an analogue with the physical world in that businesses or individuals tend to rely on a few trusted associates for collaboration [Gulati, 1995]. The advertised strength of social networks is the ability to reach a wider range of resources. However, we cannot rely on all those resources being connected. Because we rely so heavily on the assumption of connectivity to capture the advantage of automated resource discovery, this lack of connectivity is one challenge that must be addressed.

Ironically, another challenge of resource evaluation, acquisition, and integration in virtual environments is too much connectivity. In this case, a collaboration may be overwhelmed by the sheer amount of potential resources that it can consider. In the physical world, the number of potential resources to consider can also be overwhelming, but typically these resources are filtered by inherent constraints such as proximity. In virtual environments, physical restrictions such as geography are not barriers; the limits of potential resources to consider are as vast as the reach of digital communications.

3.7 Less stable collaborations

In physical environments, businesses tend to repeat business with the same partners partially because the cost of identifying and trying different partners is high [Mowshowitz, 1997]. Mowshowitz points out that services in virtual environments tend to be more suited to automated consideration of various options, and switching from one option to another is low. Additionally, new resources are more quickly available for consideration. Thus, collaborations tend to be shorter-lived unless there is evidence that keeping the partnership has some tangible benefit; secondary motivations like inertia or loyalty tend to have less effect.

As the availability of virtual environments and the general use of digital technology increases, this contributes to the generally increasing pace of business [Gould, 1997]. This is another factor

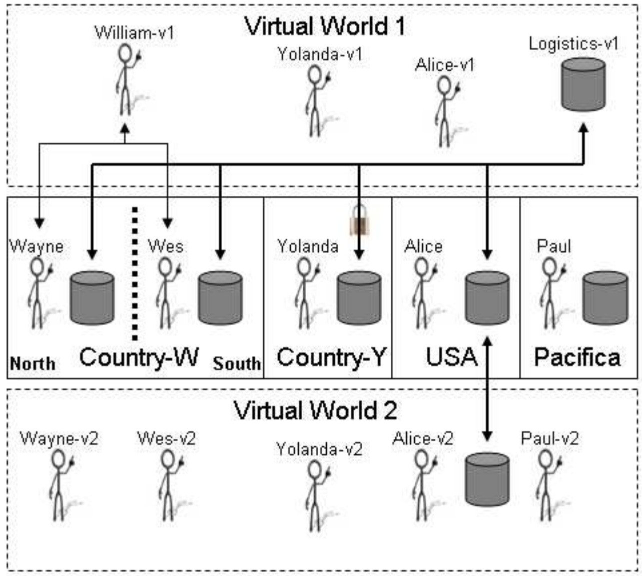

Figure 1: Multiple resources for potential use in collaborations spanning five countries and two virtual environments

contributing to the reduced stability of such collaborations. Instead of projects lasting years, they may only require a limited periodic activity. This also can imply that, as collaborators rapidly achieve goals and subsequently move to new goals, their requirements for resources change, perhaps even while still engaged in other collaborations.

As a consequence of both of these factors, there is a higher probability that an RB will be consulted repeatedly in order to find resources or new collaborators. In addition to using the RB to create new collaborations, it becomes more likely that an RB will be queried to find resources for partially formed collaborations as well. Given that repeated evaluation of resources may be expensive, it may be advantageous for an RB to return more than one solution in response to a single evaluation request. Multiple solutions could be appropriate if the suggested set of resources cannot be acquired, if a subsequent request is submitted for a similar collaboration, or if an existing collaboration requires a replacement for a resource that becomes unavailable.

4 Scenario

To illustrate some of the challenges described above, we present an illustrative example, depicted in figure 1. This C2-like scenario is derived from the PRECiS (Planning, Reactive Execution, and Constraint Satisfaction) Environment [Reece *et al.*, 1993], and includes the availability of two virtual environments in which representatives of various countries may collaborate. The scenario contains four countries: Country-W, which is divided politically into Northern and Southern regions; Country-Y, which is not an ally of the other countries due to ongoing territorial disputes; the United

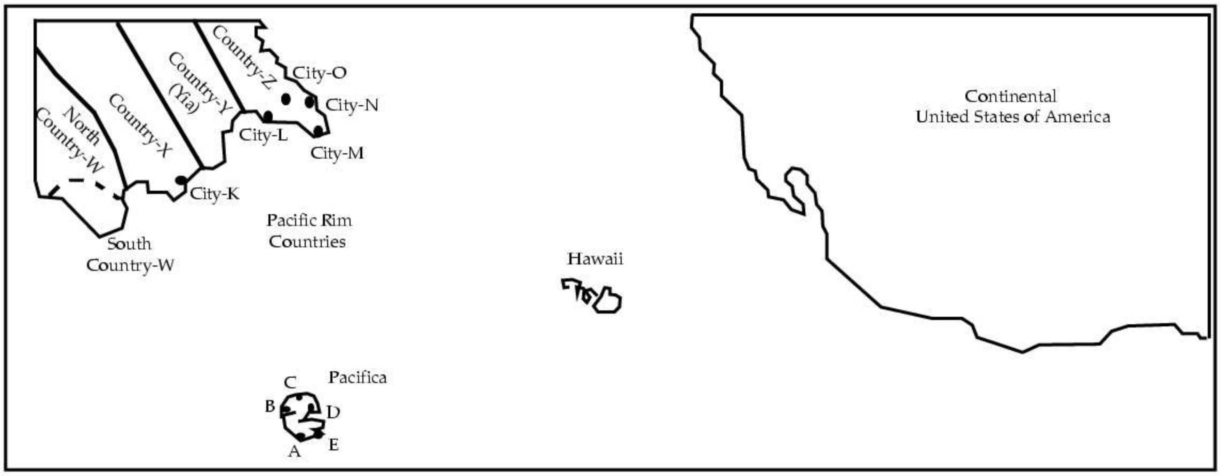

Figure 2: Pacific Theater, as defined in the PRECiS environment [Reece *et al.*, 1993]

States, and Pacifica. As is shown in figure 2, Country-W and Country-Y make up part of the Pacific Rim, and Pacifica is an island between the Pacific Rim and the United States.

Each country maintains one server that hosts a logistics computation service, except for Country-W which maintains one server in each political region. Every analyst can locally access his or her logistics tool, but no analyst has direct remote access to another analyst's tool. Physical exchange of data is not feasible between any two analysts except for those in Country-W. Although the analysts in Country-W are geographically close enough to physically exchange data, a firewall exists between their servers which prevents direct remote access.

In our scenario, the countries of the Pacific Rim (Country-W and Country-Y) are being assisted by the USA to recover from severe hurricane damage. To help with logistics planning, the countries have established a virtual collaborative environment, "Virtual World 1." The virtual worlds contain avatars supported by representatives from the collaborating countries, and virtual resources supported by servers in those countries. "Yolanda-v1" and "Alice-v1" are virtual representatives of, respectively, Yolanda of Country-Y and Alice of the USA. Wayne of North Country-W and Wes of South Country-W have similar enough geographical interests such that they aggregate their contributions into the actions of the "William-v1" avatar. The "Logistic-v1" virtual resource is supported by the combined efforts of the total four servers in Country-W, Country-Y, and the USA, and is able aggregate information to give more accurate results than any individual server. However, depending on the speed of moving data from the physical world to the virtual world, particularly over Country-Y's secure connection, the results can be delayed or even inaccessible at times.

"Virtual World 2" is a collaboration space to support the treaty negotiations between Country-Y and Pacifica, with the USA and Country-W as moderators. Unlike in Virtual World 1, each of Country-W's regions is represented by its own avatar because of their distinct political interests. In addition, one avatar from each of the USA, Country-Y, and Pacifica are present in this virtual world. The server in the USA supports a virtual political analysis tool, "Analysis-v2," to help with the treaty negotiations.

4.1 "Simple" resource allocation

Consider the simple collaboration in Virtual World 1 where "William-v1," "Yolanda-v1," and "Alice-v1" work together to coordinate logistics for the disaster relief effort in the Pacific Rim, and assume the collaboration includes a logistics planning tool. Even in this simple scenario, there are *thirty-six* different options to consider when determining which logistics tool or tools to use!

- For Walt's contribution to "William-v1", he has the option of consulting either of the logistics tools in Virtual World 1, North Country-W, or South Country-W (if he chooses to physically travel).
- For Wes's contribution to "William-v1", he has the option of consulting either of the logistics tools in Virtual World 1, South Country-W, or North Country-W (if he chooses to physically travel).
- Yolanda may consult the logistics tool in either Virtual World 1 or Country-Y
- Alice may consult the logistics tool in either Virtual World 1 or the USA

Walt and Wes each have three options and Yolanda and Alice each have two options, resulting in thirty-six combinations of resource configurations.

The appropriate resource or resources to use depends on the cost of moving information between the physical and virtual environments, the cost of various translations, and the required accuracy and speed of those translation. Additionally, some allocations will allow an entity to utilize either a local, physical resource or a remote, virtual one. In this case, the scheme employed also affects which allocation is optimal.

4.2 External constraints on resources

Typically, an RB will apply constraints determined by the collaboration's needs to the selection of appropriate resources. Uncertainty within constraints and resource capability is not uncommon. However, when dealing with multiple environments, these constraints may be manifestations of influences from another, apparently unrelated, environment. Again considering the scenario from figure 1, note that the server in the USA supports both the Logistics-v1 and Analysis-v2 resources. If an analytic resource (either physical or virtual) is allocated to another user or collaboration, then the additional load on the server may impact the effectiveness of the Logistics-v1 resource in Virtual World 1. From the perspective of resource allocation in the virtual world, this type of relationship is impossible to consider because, from the virtual point of view, there is no relationship between the Logistics-v1 and Analysis-v2 resources.

This type of phenomenon can be modeled as uncertainty with respect to resource capability, however, it would be better if we were able to consider the actual source of the uncertainty in determining resource allocation. The means of including constraints from external environments is an open question.

Another aspect of external constraints on virtual resources is the collaborators' ability to control multiple avatars. As can be seen in the scenario in figure 1, Alice will need to control both Alice-v1 and Alice-v2 if the negotiations in the virtual worlds are conducted simultaneously. Although virtual technology allows Alice to "be" in multiple "places" at one time, care must be taken that this multitasking is not done to the point where Alice's productivity experiences diminishing or negative returns.

4.3 Security

There is a common interest in the disaster relief efforts between Country-Y and the rest of the countries, but the relationship is complicated by the adversarial political relationship. Country-Y and the other countries of the Pacific Rim are interested in aggregating logistical insight to support the disaster relief effort, but both parties are also concerned about information supplied for the relief effort in Virtual World 1 becoming an advantage to the other party in the political negotiations in Virtual World 2.

In addition to the information itself, both parties might like to conceal possession of an expertise or set of knowledge. Conversely, both parties might like to avoid advertising interest in certain types of information. For example, perhaps there is a mustard gas facility in the disaster recovery area. The country that owns the facility would like to query for information about mustard gas mitigation and cleanup without revealing its interest or need. Analogously, a country may be interested in supplying mustard gas knowledge for the good of all involved, but would not want to reveal the fact that it has collected a mass of knowledge in that area. Ideally, the parties would like to anonymously supply and query for information, and at the same time feel confident in the validity of the information.

How to apply research in information assurance and security to virtual environments is an open and important question.

5 Integrating a resource broker within a virtual environment

The Johns Hopkins University Applied Physics Laboratory (JHU/APL) has developed a prototype resource broker (RB), the *Dynamic Collaboration Action Team Resource Broker* (DCAT-RB), to locate and request resources for use in predefined situations such as flood disaster recovery or a chemical biological attack. This prototype finds resources to satisfy collaboration needs within physical environments.

Recently, we created a prototype integration of the DCAT-RB with a Croquet-based virtual world. Croquet is a software development kit for virtual world development. It is built upon the Squeak development environment, which is a Smalltalk graphical environment.

Our prototype demonstrates two capabilities. First, it shows the general ability to make a resource based in the physical world available in a virtual world. Second, it demonstrates the ability for a resource broker to programmatically access contextual information about its encompassing virtual environment and apply that information when calculating resource recommendations. This type of awareness is a difficult challenge in the real world, but is readily available in virtual environment for resources that are designed to exploit it.

5.1 Scenario

To motivate our design, we describe a simplified version of our working scenario. In this case, we have an avatar residing in a world containing an interface to the external RB service. A second avatar enters the world to request recommendations to complete the collaborative team that will meet in the virtual world. The second avatar is knowledgeable of team needs and configures the DCAT-RB to create a team that will address a remote flood disaster. When the avatar submits the request, the DCAT-RB interface surveys the local world for resources that are already present, and incorporates that information into the request. The DCAT-RB returns a suggested team consisting only of resources that satisfy the needs of the team that do not already exist in the world. By virtue of the DCAT-RB having an awareness of the current context, there is no need for either avatar to exhaustively list current resources, remove redundant resource requests, or waste resources on overprovisioning.

5.2 Selection of Croquet

In addition to Croquet, we considered integrating with SecondLife and Sun Wonderland. We did create a proof-of-concept implementation in SecondLife. However, we felt that a collaboration between parties involving sensitive information, such a group of company employees designing a new product, would preclude the use of SecondLife due to its use of third-party servers. As mentioned in section 3.3, servers can record and control interactions within the virtual worlds they host. Thus, those who use virtual environments for sensitive exchanges would likely want to host the virtual world on their own servers. Secondly, we found that importing arbitrary applications into SecondLife is not supported. This is reasonable as otherwise Linden Labs might end up becoming a large farm of application servers. Finally, we investigated the use of HTTP requests within SecondLife. We found that communications with external services, although invoked from the local client, originate from the Linden servers. Because our application resides behind a firewall, it could not receive HTTP requests originating from Linden, and thus this technique was not an option.

Unlike SecondLife, Croquet is a distributed virtual environment with no central server. Thus, Croquet worlds can be hosted on any machine running the Croquet software. Also, Croquet allows one to import "projects" into a world that consist of arbitrary code. We have not yet explored how to import project that are not based on Croquet's native Smalltalk, but we expect that Smalltalk has such facilities. Lastly, because Croquet runs on individual machines inside our firewall, we are able to communicate with locally hosted services.

We were optimistic about the possibility of directly integrating our Java-based application into the Java-based Sun Wonderland project. However, that functionality is currently not supported on Microsoft Windows-based systems. Given that setback as well as time constraints, we did not explore that integration possibility any further.

5.3 Implementation

We implemented two versions of our demonstration. One takes advantage of Croquet's "Squeaklet" capability, which allows small projects written in Smalltalk to be imported into a Croquet virtual environment. This technique is very user-friendly, as much of the creating of the Squeaklet may be done graphically. However, these imported projects are not able to access information about the enclosing environment other than information explicitly supplied by avatars through the Squeaklet's interface. The other version of our demonstration is coded directly into the definition of our the virtual environment. By virtue of this implementation, the resource broker interface is able to survey the local environment and respond appropriately. The disadvantage of this method is a the requirement of a stronger familiarity with Croquet code, and it is more difficult to distribute the interface to other environments. Additional details of our work follow.

Our DCAT-RB implementation is Java-based, so we have not yet attempted a full application import into the Smalltalk-based environment of Croquet. As an alternative, we implemented a version of the DCAT-RB such that it could be invoked and respond via an HTTP connection. Within Croquet, we created a Squeaklet project with text boxes for input and output, and a button to invoke the call to the DCAT-RB. For the sake of time, our initial approach was to import this project into a supplied sample Croquet world, "MPEG Demo (Master)" which has a project import menu. After we import the project into the world, we are able to create a second virtual world, open a portal to the MPEG world, and have a second avatar enter the MPEG world. That avatar is able then able to use the DCAT-RB application that has been made available by the user from the first world. Unfortunately, we could not discover a way for the project to be aware of the environment in which it was embedded. This was not a problem as long as relevant information could be supplied to the interface, but in complex cases where some resources are already available in the virtual environment, we may not want an avatar to have to comprehensively document the state of the environment.

Our second attempt directly augments the code that generates the Croquet world. In this case, we created a copy of the supplied "SimpleDemo (Master)" world and modified the initialization code to create a cube, that serves as the interface to the DCAT-RB. This class, CubicRBInterface, extends Croquet's TCube to allow it to handle calls to the DCAT-RB. To support awareness of the environment, we pass the TSpace reference representing the local Croquet environment to the CubicRBInterface. When the cube is clicked, it uses the TSpace reference to determine which avatars are in the local environment, and passes this information to the DCAT-RB. If an avatar that represents a resource that the team requires already exists in the world, then a request for that resource is not generated. For large teams, it is advantageous to automate the step of reporting previously acquired resources. The primary disadvantage of this method is that we cannot drag and drop our service interface into a world as we could with the Squeaklet. Also, new functionality must be added programmatically instead of graphically, which can make complex widget tedious to implement.

5.4 Discussion

This incorporation allows multiple avatars, potentially representing diverse remote collaborators, to work together to evaluate team needs. Additionally, inclusion of the DCAT-RB within the virtual environment allows it to be contextually aware, allowing it consider local resources when generating suggestions for additional resource needs. Our work prototypes a specific example of a generally applicable technique for establishing collaborative and contextually aware applications in virtual environment. This demonstration is implemented as a Java-based application invoked from within a Croquet multi-user virtual environment (MUVE).

5.5 Future Work

Our proof of concept implementation suggests additional work to enhance both the virtual environment and the DCAT tool. The virtual environment we chose was a free software development kit that suited our initial goals. However, there are more fully-featured and high level tools available to build virtual environments. As a next step, we would like to investigate including our DCAT product in a more sophisticated virtual environment. OpenSimulator[13] is one interesting option. OpenSimulator is a free, open source, server that supports Second Life-like virtual worlds. For many potential users, the ability to house one's own server is essential, particularly when clients may be exchanging proprietary information. As mentioned earlier, the usefullness of integrationg our DCAT tool into a Second Life environment was limited by the fact that all communications ultimately had to pass through the third-party servers that support Second Life. With OpenSimulator, this constraint would be removed. Unfortunately, OpenSimulator is alpha software with limited functionality, which may limit the success of our integration efforts.

The concepts exhibited by our DCAT tool are a promising benefit to collaborations that will take place within virtual worlds. However, our current implementation uses a rudimentary set of metadata to describe resources. A general solution would require the creation of a richer set of metadata.

[13]http://www.opensimulator.org

We expect that determining the information represented within the metadata, and the means to evaluate resource compatibility based on that metadata, to be an iterative process. Our next steps would include choosing a domain for which to generate appropriate metadata, followed by soliciting potential users for feedback.

6 Summary

As the use of virtual environments becomes more prevalent, collaborations will more frequently require resources that have some basis in the virtual world. The ability to take maximal advantage of such resources will be a determining factor in the effectiveness of these collaborations. We have described some relevant issues to consider when evaluating, acquiring, and integrating virtual resources for use in collaborations, and we have also presented a brief scenario to concretely illustrate these issues.

Finally, we have offered brief example of how a traditional application for resource allocation in the physical world might be adapted for use within a virtual environment. Our proof-of-concept implementation demonstrates the potential of the utility of resource brokering in virtual worlds. However, for these techniques to be generally useful, the challenges previously discussed will have to be resolved.

When collaborations and resources have the potential to span both the physical world and possibly multiple virtual environments, the impact on the use of those resources becomes non-trivial. We must determine when there is a benefit to bearing the cost associated with transitioning a resource across a boundary between environments. Additionally, the functionality or value of a resource may be impacted by the extent to which either it or its associated collaboration is not limited to just one world. Also, a resource may be affected by events in another environment, and the mechanism to recognize those events and model their impact is an open question.

Additionally, a resource may have a presence in both the physical and virtual environment. The potential benefit of a resource may vary depending upon the version (physical or virtual) that the collaboration uses. For example, a physical resource located far away from the collaboration will likely be of little benefit; however, a virtual version of the resource may be very useful. A characterization of which resource attributes are relevant in the physical world versus virtual environments would therefore be very helpful in determining how to evaluate a resource for use by a collaboration.

Because the connectivity of the network in which resources are located can either be either very disconnected with resources more difficult to find, or connected and a large number of resources to consider, an automated resource acquisition system (RB) should have the ability to adapt to at least one of these non-ideal situations. Ideally, an RB might have the ability to diagnose the nature of a particular network in which it is searching for resources (presumably, different resources could reside in distinct networks), and adapt its acquisition strategy accordingly.

The security implications of all these potentially diverse resource configurations is an issue that will have to be resolved. As resources are represented and accessed in both physical and virtual environments and move between different environments, we must be able to guarantee necessary synchronization, while allowing only appropriate accesses.

References

David S. Alberts. Agility, focus, and convergence: The future of command and control. *The International C2 Journal*, 1(1):113–144, Spring 2007.

Charles Babcock. What's the greatest web software ever written? *InformationWeek*, May 5, 2007.

Bryan T. Camp. The Play's the Thing: A Theory of Taxing Virtual Worlds. *SSRN eLibrary*, 2007.

Benny Chor, Oded Goldreich, Eyal Kushilevitz, and Madhu Sudan. Private information retrieval. In *IEEE Symposium on Foundations of Computer Science*, pages 41–50, 1995.

R. Scott Cost, Markus E. Dale, David P. Glock, James Mayfield, Christine O. Salamacha, and David P. Silberberg. Support for dynamic collaborative action teams. In *Proceedings of the Command and Control Research and Technology Symposium (CCRTS 06)*, 2006.

Mathijs M. de Weerdt, Yingqian Zhang, and Tomas Klos. Distributed task allocation in social networks. In *Proceedings of the AAMAS*. IFAAMAS, 2007.

David A. Gould. *Leadership in virtual teams*. PhD thesis, Seattle University, available from Ann Arbor, MI: UMI Dissertation Services, 1997.

Ranjay Gulati. Does familiarity breed trust? The implications of repeated ties for contractual choice in alliances. *The Academy of Management Journal*, 38(1):85–112, February 1995.

Reiner K. Huber, Petra M. Eggenhofer, Jens Römer, Sebastian Schäfer, and Klaus Titze. Effects of individual and team characteristics on the performance of small networked teams. *The International C2 Journal*, 1(1):113–144, Spring 2007.

Magid Igbaria, Conrad Shayo, and Lorne Olfman. On becoming virtual: The driving forces and arrangements. In *SIGCPR '99: Proceedings of the 1999 ACM SIGCPR Conference on Computer Personnel Research*, pages 27–41, New York, NY, USA, 1999. ACM Press.

Magid Igbaria. The driving forces in the virtual society. *Communications of the ACM*, 42(12):64–70, 1999.

Samuel R. Kuper and Beth L. Giurelli. Custom work aids for distributed command and control teams: A key to enabling highly effective teams. *The International C2 Journal*, 1(2):113–144, Fall 2007.

Abbe Mowshowitz. Virtual organization. *Communications of the ACM*, 40, Issue 9(9):30–37, September 1997.

G. Reece, A. Tate, D. Brown, and M. Hoffman. The PRECiS Environment. In *Proceedings of National Conference on Artificial Intelligence (AAAI-93) DARPA-RL Planning Initiative Workshop*, 1993.

Christine O. Salamacha, R. Scott Cost, Michael Cramer, Timothy Frey, and David P. Glock. Dynamic collaborative action teams: Implementing a transformational concept. In *Proceedings of Command and Control Research and Technology Symposium (CCRTS 06)*, 2006.