Tool Support for Security-oriented Virtual Research Collaborations

John Watt, Richard Sinnott, Jipu Jiang,
Tom Doherty
National e-Science Centre
University of Glasgow
Glasgow, G12 8QQ, UK
e-mail: j.watt@nesc.gla.ac.uk

Chris Higgins, Michael Koutroumpas
EDINA
160 Causewayhead
Edinburgh, EH9 1PR, UK

Abstract—Collaboration is at the heart of e-Science and e-Research more generally. Successful collaborations must address both the needs of the end user researchers and the providers that make resources available. Usability and security are two fundamental requirements that are demanded by many collaborations and both concerns must be considered from both the researcher and resource provider perspective. In this paper we outline tools and methods developed at the National e-Science Centre (NeSC) that provide users with seamless, secure access to distributed resources through security-oriented research environments, whilst also allowing resource providers to define and enforce their own local access and usage policies through intuitive user interfaces. We describe these tools and illustrate their application in the ESRC-funded Data Management through e-Social Science (DAMES) and the JISC-funded SeeGEO projects.

Keywords—*shibboleth, permis, nesc, shintau, seegeo*

I. INTRODUCTION

It is fairly common to stumble across stories of data compromise and outright carelessness with sensitive information in the national newspapers of most countries these days. More likely than not, the story will detail how a CD or laptop with the personal details (and in some cases, bank account numbers) of thousands of employees or citizens was either abandoned unprotected on public transport, or lost in transit by a courier. While the free movement of data is critical to nearly all business and research collaborations, controlling who gets this data, and what they can do with it once acquired, represents a major research challenge.

There is a constant trade-off between security and usability in developing and supporting e-Infrastructures, especially those that provide access to federated data. On the one hand, the need for state-of-the-art middleware that gives collaborators the level of security they need is an unavoidable requirement. Yet this technology can be baffling for non-IT savvy end users who may not know what an X.509 digital certificate[1] is, let alone possess the critical know how to use it properly to access distributed resources. Many projects have adopted approaches based upon virtual research environments where portals are supported that provide a one-stop shop offering access to distributed information and services tailored to the individual research domain. Whilst overcoming aspects of usability this model

introduces other concerns including access control. Thus a remote provider, especially in more security-oriented domains not simply delegate access control completely to a remote portal. Rather they will still wish to remain autonomous and define their own access and usage policy when requests are received from the portal. This paper describes a security model and supporting set of tools being adopted across a wide range of UK e-Science projects at NeSC in the biomedical, geospatial and clinical domains. The focus of the model is the integration of complex middleware like Shibboleth[2], PERMIS[3], portals frameworks and the Globus Toolkit (globus.org), to support local site administrators, virtual organisation administrators and essentially, making the interaction of these technologies as transparent to the end user researchers as possible.

The rest of the paper is structured as follows. Section II provides an outline of the underlying technologies upon which the security-oriented tools have been developed. Section III describes the tools themselves and their functionality. Section IV describes the application of these tools to support different security-oriented application domains in the social/clinical sciences and in the social/geospatial sciences.

II. UNDERLYING TECHNOLOGIES

Overcoming user issues in access to distributed resources is often achieved through development and support of portals. Whilst addressing user interface challenges, portals do not in themselves address all usability and security concerns. The notion of autonomy (where potentially remote sites can make their own authorisation decisions) and single sign-on (where users are not required to authenticate at each distributed resource) are two areas in particular that must be supported to support *distributed* virtual research collaborations.

A. Portals

First, confirm Portals enable a service (or collection of services) to be accessed through a regular web browser. This is a major gain in usability as internet browsers are ubiquitous and thus familiar to all researchers. Instead of complicated command line switches and environments, users are presented with a project-tailored user interface that allows them to submit jobs or run queries by filling in forms or clicking buttons. The portal can also manage and

aggregate the results, and display them in a graph or output file.

A portal is a container for hosting *portlets*, which are fragments of markup generated to allow access to an application across the internet. An application or service will typically be exposed by a single portlet, whereas if a project depends on several backend services, a portal can aggregate these services into a single page for the user.

The JSR-168 specification[4] describes a framework for writing portlets which should be independent of the chosen portal framework, of which there are many to choose from. Apache Pluto runs a reference implementation of JSR-168, whereas other portal containers like GridSphere (gridsphere.org) contain custom tags intended to make programming the user interface easier. By sticking as close to JSR-168 as possible, portlets may be authored which can be deployed in any portal container.

A recent extension to the specification, JSR-286, adds functionality for inter-portlet communication, plus defines some portlet security features which were once handled purely by the portal container. These new features are of particular interest to the projects described in this paper, and early adoption now may make the work we are doing fully transferrable to future portal frameworks with minimal re-writes.

Considering the portal model more generally, the migration of portal functionality into the JSR-286 portlet specification hints at a wider change towards portal technology being a client-side process which consumes remote portlets, rather than a service-side gateway into local portlets. For this reason, our portal tools have attempted to stick as close to the standards-based approach as much as possible, since portal framework technologies themselves tend to be fluid and subject to change/evolution.

Another important point to be made here, concerns the famous 'n-tier'[5] problem, where the question of how to securely marry identities existing in separate security domains is important. To understand this, we need to look at federated authentication and Shibboleth, which at first glance appears to magnify the n-tier problem, but through use of other middleware can provide a solution.

B. Shibboleth

The Grid and recent incarnation to Cloud Computing presents a major user management challenge. In an environment which may comprise thousands of users from locations scattered around the globe, a simple Access Control List (ACL) can quickly become unmanageable. And even if management is possible, centralised user lists can only be accurate for a short period of time. Over longer spells, users may leave their respective institutions, privileges may be revoked - all of which mean central lists maintained by resource providers are no longer authoritative sources of user information.

Consider the situation where a new member of staff or student begins work at a new university, company or other institution. Typically a student or member of staff will register *in person* with a registrar or human resources department, and this information will be used or delegated to

create a campus identity for the user. If the user subsequently leaves the institution, through fair means or foul, the identity can be quickly disabled and the user's rights on campus will be revoked. The choice for user management therefore comes down to two options, either create a system of in-person registration that can match the reliability of the institutional method, some of which may have been in place for centuries, or somehow leverage the information provided by the institution directly. The former is geographically and not financially viable for arbitrary systems, as it may involve great expense and travel to perform the face-to-face validations. The latter model is the focus of great effort in defining standards and protocols to achieve federated access management.

Shibboleth[2] provides an architecture and protocols for transport of security information between institutions and providers. This security information is implemented in the Security Assertion Markup Language (SAML), the first version of which[6] is supported in the UK until 2010. This standard describes how information about the identity and privileges of users may be transported between services without compromising confidentiality or reliability. SAML describes specific entities within an organisation known as a *federation*, which is a collection of sites that have agreed to trust the information they give and receive. SAML defines an Identity Provider (IdP), which is a source of user information, and a Service Provider (SP) which is a consumer of user information. SAML also defines a core set of user attributes, which are extra pieces of information which may be used by an SP to make further access control decisions, e.g. based upon Role-Based Access Control (RBAC). Shibboleth also provides a single sign-on solution for portals, allowing a user to navigate between federation resources without having to input their credentials with each move. Or more precisely a user is able to access a Shibboleth protected portal/web resource, and if their privileges allow access other portal/web resources without the need for further re-authentication.

In the UK, the UK Access Management Federation [7] performs the job of registering SPs and IdPs within the trust framework, and is responsible for publishing metadata giving a snapshot of all the trusting entities in the community. Using national-level federations makes configuration of Shibboleth a lot easier, as one of the hardest parts of installing Shibboleth is the creation of the metadata which describes how a system interacts with the others in the federation. It may become necessary soon to conduct investigations into integrating national federations to allow continental-scale access to services, an example being the EuroDSD project (www.eurodsd.eu) which provides data services for collaborators in seven countries across Europe.

C. Shintau

Shibboleth is a complex, but stable way to distribute reliable assertions that the person who is attempting to access a resource is the person they say they are. But this is only one half of the security story, as establishing a user's identity doesn't in itself tell you anything about that user's permissions on a resource. As mentioned above, SAML can

be used to also transmit user *attributes* for access control, which will tell the system what the user is allowed to do on that system. Currently, Shibboleth transports this information directly from the IdP in a signed assertion, meaning that the IdP needs to have all user attributes present either in the IdP's own database, or aggregated automatically by the IdP for this to be of any use. A user could feasibly accrue hundreds of attributes for access to remote services, which makes managing attributes at the IdP an administrative nightmare. Surveys [8] have looked into various ways in which access control via multiple Attribute Authorities may be achieved without compromising security or usability.

Shintau [9] is a proposed extension to the Shibboleth infrastructure, where instead of user attributes being directly asserted by the IdP, they are aggregated by the *user* during the authentication process, and these attributes are asserted by the IdP on the user's behalf. This allows an IdP to remain an accountable source of user authentication, but the responsibility for collecting the correct attribute(s) required for access at the SP becomes the responsibility of the user. Shintau defines the concept of a SAML Linking Service (LS), where a user registers the remote IdPs (or Attribute Authorities) which they have registered with, and links it with their main IdP. When logging into their home IdP, the linking service automatically forwards the attributes from the remote IdPs and presents them within the main IdP's SAML attribute assertion.

Whilst Shintau and Shibboleth can both be used for delivery of attribute information, ultimately a resource will itself need to use this information to make local authorisation decisions on access and usage. One way that this can be achieved is through the RBAC-based Privilege and Role Management Infrastructure Standards Validation (PERMIS) [3] technology.

D. PERMIS

PERMIS is a generic authorisation infrastructure which issues roles and privileges to users using X.509 Attribute Certificates. PERMIS provides a plug-in authorisation enforcement point (PEP), and tools with which to issue ACs and write local security policies.

A security policy is a document written typically in XML which details the precise access control requirements of a resource. In PERMIS, policy definition and enforcement use an XML triple comprising a Role, a Action and an Target. For many purposes, the Role is contained within an X.509 Attribute Certificate (AC)[10] which the user provides directly or through extraction of the Distinguished Name (DN) and subsequent LDAP lookup. The Target represents the URI of the Grid/Web service which the user is attempting to access, and the Action is the individual method that the user is attempting to invoke on this target. The XML policy dictates which combinations of Action and Target are permitted based on the Role that the user presents in their AC. Additional rules about which PKI keypairs are recognised by the PERMIS PEP, which certificate DNs are permitted and validity time can be expressed in the XML policy. All objects in PERMIS are digitally signed, ensuring

the information hasn't been tampered with, and allowing the issuer of the certificate/policy to be confirmed.

Direct usage of PERMIS is non-trivial and requires that someone at the resource provider becomes proficient at creating and managing Public Key Infrastructures (PKI), operation of LDAP servers, and then the PERMIS tools themselves (Attribute Certificate Manager, Policy Editor). One of the NeSC projects (SPAM-GP, described in section 3) has provided some assistance in running this middleware, but the complexity of the surrounding infrastructure is one of the key issues with PERMIS.

E. Globus Toolkit and MyProxy

The Globus Toolkit is a suite of tools which enable Grids and especially computational Grids to be built. The suite consists of a vast array of open-source tools and applications, however these may be installed all as one package, or only parts of the kit may be installed depending on the application. For Grid Service security, Globus provides Grid Security Infrastructure (GSI) tools that allow to pass user credentials on to services. Typically this is through creation and use of proxy certificates. Using the UK e-Science Certification Authority for these certificates provides compatibility with UK Grid resources like the NGS. Utilising the MyProxy[11] tool, a degree of automation may be achieved with the handling of proxy certificates.

Building upon these technologies we have focused upon developing tools that simplify the creation and use of user-oriented, security-driven portal research environments.

III. INTEGRATION TOOLS

The SPAM-GP project was commissioned by OMII-UK to provide a suite of tools to allow Shibboleth and portals to be integrated and configured. Three areas of infrastructure configuration were proposed and delivered, each providing a different level of security control.

A. SCAMP

A Shibboleth SP can receive SAML attributes from any IdP in the federation it subscribes to. The rules about what form of attributes may be accepted and which locations these attributes may originate from is expressed by the Attribute Acceptance Policy (AAP). This is an XML document which must be edited by hand to reflect the desired security rules. Editing raw XML can be daunting at the best of times, but when this policy expresses access control rules, then extra care must be taken to ensure that the edited policy does not compromise the security of the resource. The default AAP which ships with Shibboleth has fairly lax rules in place which should be edited prior to first deployment.

The Scoped Attribute Manager Portlet (SCAMP) is a JSR-168 compliant portlet which allows the Shibboleth AAP to be edited correctly. The portlet reads in configuration files from the Shibboleth installation, building a list of the federation IdPs and the current state of the AAP. The portal administrator is then presented with the current policy in a JSP page, which may be added to or edited as the

resource requires, safe in the knowledge that the XML produced will be valid. Since the Shibboleth AAP dictates the attribute set that *every* service on the resource will see, the tool is intended to be used by resource owners only, so the tool is normally deployed in the portal container so it can only be accessed by the portal administrator.

B. ACP

The PERMIS infrastructure provides several tools for issuing ACs, which normally run as separate components invoked from the command line. The Attribute Certificate Portlet (ACP) is a JSR-168 compliant portlet that can be deployed in the portal along with other application specific portlets (interfaces to services). The ACP allows a privileged user to issue X.509 ACs for access to potentially remote backend services. The roles that may populate the issued AC may be completely user-defined, or they may be restricted to particular attributes which have been asserted by Shibboleth. Provided the PKI Source-of-Authority is recognised by the PERMIS policy of external services, these ACs may be used to enforce RBAC.

C. CCP

The Content Configuration Portlet (CCP) is an extension to the GridSphere framework which provides an additional login module that can build user login sessions from information provided by Shibboleth via SAML. Since this module makes changes to the core framework, it can't really be published as a JSR-168 compliant portlet. The CCP is based on the MAMS Shibboleth GridSphere model[12], but with extensions to allow any asserted SAML attribute to be utilised for access control, and stored in the GridSphere user/role database.

Other portal frameworks such as Sakai and LifeRay support functionality similar to CCP. The Sakai workspace concept for example allows to define new user groups that users can be mapped to when they log in to the portal. Indeed GridSphere also supports this, however the roles and groups are very limited (admin, guest etc).

To understand how these portlets and the underlying technologies can be used to support security-oriented collaborations we describe their application in two projects: SEE-GEO and DAMES.

IV. APPLICATIONS

A. SEE-GEO

The Secure access to Geospatial Services (SEE-GEO), funded under the JISC Grid OGC Collision programme, aimed to investigate how to make geospatial data accessible utilising 'Grid' technologies. The project produced a GridSphere portal-based linking service, joining geospatial boundary data hosted at EDINA (edina.ac.uk) with census data sets provided by MIMAS (mimas.ac.uk). Both of these data sets have security requirements associated with them, e.g. the EDINA data comes from Ordnance Survey and thus has strict licensing terms and conditions which users must

abide by. NeSC was tasked with implementing the associated security infrastructure and demonstrating practise in how to secure these types of data in the future.

One of the challenges in building security onto existing data services is the fact that these resources usually have their own completely customised authorisation requirements. A centre hosting a production-level data service will typically not be a willing candidate to deploy and test new middleware, no matter how much these tools meet their own requirements. They are by nature conservative and are tasked at support production level services to user communities. With this in mind, the focus of the project was on delivering an accessor service to the EDINA-hosted Web Feature Set (WFS) [13] service. This was written in Java and deployed using Globus Toolkit 4, offering method-level access to three different geographical feature sets (England 1910, England 2001, Scotland 2001). The accessor service was designed to be as lightweight as possible. The service itself supported several methods which were protected, i.e. required authorisation. These methods receive queries generated by a portal, and if sufficient authorisation information was provided, they would subsequently allow access to each individual feature set. PERMIS is well suited to protecting methods of a GT4 service, so the authorisation for the accessor service was tailored to require that the user hold a valid X.509 AC containing the correct role for that specific data set to be visible.

Figure 1 shows the interactions of the services in SEE-GEO. A WFS rendering portlet was deployed in GridSphere, to which access was restricted via the SCAMP tool to only access attributes of types *wfs*. The CCP was used to restrict the portlet view to only those people holding that attribute. The portal administrator was given access (through CCP) to the ACP which was used to issue the necessary ACs required to query the WFS service. Based on the feature set selected by the user, a GSI call to the accessor service was made, using a proxy credential loaded by the user. PERMIS utilised the DN of this certificate to successfully retrieve the correct AC issued for the user by the administrator. Based on this AC, access was either denied, or if allowed, the method was invoked, the query executed and data returned.

Figure 1: The GT4 'Accessor Service' model (marked by PEP) adopted for the SEE-GEO project. SPAM-GP tools are in green.

Since the accessor service is itself independent of location, it was possible for development of the accessor service to take place without disruption to the production service. Ideally the accessor service would be deployed at the data service location, acting as a secure gateway to the data. In fact, it should always be the case that the responsibility for defining the access policy, deploying the service, and issuing the required credentials is performed by the resource providers. In this work the GT4 service was itself hosted at NeSC in Glasgow. It is hoped the work done in implementing this demonstrator will encourage data providers to adopt a more generic standard that sacrifices none of their security needs, but allows a much more distributed and automated level of access.

One aspect of this case study was understanding how user identities are propagated and confirmed through n-tier systems. The accessor service identifies users by the proxy certificate that is presented by GSI, the DN of which is extracted to allow PERMIS to locate the correct ACs for the user. The question of how this proxy credential is generated with no user input is still an open one. We currently use a MyProxy server to which we have designed a portal interface that allows a user to load and issue themselves (with interaction) with a proxy certificate based on their UK e-Science X.509 credential. Although this method works and is secure, it is by no means a standards-based approach, and better methods based on MyProxy could be envisaged. For example, a MyProxy server that is able to consume valid SAML assertions to issue proxy certs in real time is one possibility. Other methods include a MyProxy IdP which could pass proxy certificates along with the Shintau-aggregated attribute set, or utilisation of the future SAML Holder-of-key profile (possibly consumed by MyProxy).

B. DAMES

The ESRC-funded Data Management through e-Social Science (DAMES) project is a three-year project looking at data management activities relating to occupation, education, ethnicity and clinical/e-Health and wider social care data sets. Data management challenges faced in the social sciences are numerous: investigation of trends in data over decades (longitudinal studies) where different coding systems and classifications and categorizations are used, e.g. regional boundaries of local authorities change over time and understanding this when dealing with changing population dynamics or the change/impact upon health policy is essential to guide policy or understand research questions impacting upon society more generally.

In DAMES, NeSC has been extending the model developed in SEE-GEO to provide linkage into live Census data hosted by MIMAS and linkage with other clinical resources, specifically to understand research into self-harm and depression. When a user wishes to gain access to any census data sets, they visit the Census Registration Service hosted at the University of Essex, and make an online declaration for each data set that they will abide by the terms and conditions for access. Once approved, the user is taken to an online Census Service (casweb.mimas.ac.uk) where the user can subsequently download the census data of interest.

These interactions are done via Shibboleth inside web browsers. One of the challenges with regard to security and dealing with access to distributed data however, is that there is no programmatic API or service through which the Census data can be accessed. Rather it is a web form that a user is offered for direct data download. Ideally for authorised access, a service should be defined through which access control should be enforced.

In the first instance we have directly downloaded the Census data to NeSC in Glasgow. A GT4 service has been implemented that has security (authorisation) requirements that must be fulfilled before this data can be accessed. We have developed two specific portlets that provide different query interfaces to this Census data. These allow researchers to select subsets of the Census variables related to health and well-being.

In addition to Census data, the researchers we work with also require seamless access to clinical data. The Scottish Morbidity Records cover clinical data from over 30-years across Scotland. This includes all hospital admissions (SMR01), mental health/psychosis (SMR04), cancer registrations (SMR06) and death registrations (SMR99). It is planned through the recently funded Scottish Health Informatics Platform for Research (SHIP) project (www.scot-hip.org.uk), that we will have direct access to these data sets hosted by the NHS. However, until now we have been given a pseudonymised subset of these data sets (apprx. 4 million records in total) that allow for basic data explorations and linkage to be supported. For these different services we have developed GT4 services that require authorisation.

Given that the Census have established a user registration list of who has registered for which particular data sets we would like to leverage these efforts, since ultimately they are the source of authority of who has/has not signed up to which license conditions on which data. A proposed architecture for this is shown in Figure 2. We note how the fine-grained user attributes are only transmitted between the data source and the registration service, and also how the registration service may require the DN of the user's X.509 certificate to make the link when the portal makes a request to the PERMIS-protected census data service.

Figure 2: The Accessor Service model proposed for DAMES, showing interaction with existing Registration Service.

We note that we have also investigated several newer portal frameworks, including Sakai and Liferay, but our focus is on portal-independence, and indeed some of our services may utilise non-portal based access via WebDAV and iRODS.

V. CONCLUSIONS

A security model for federated data services has been demonstrated conceptually, and in practise, as a possible generic solution for robust access control. The model assumes a federated authentication system based on Shibboleth, with extra user information delivered through Shibboleth from remote IdPs. We are also working on exploitation of aggregated ACs from several attribute authorities using the Shintau Linking Service acting on behalf of the user.

The SPAM-GP project has implemented administrator tools at the portal level to control the user environment based on their held privileges, and also a tool to allow signed X.509 credentials to be used for final fine-grained access control to the constituent back-end services. The SEE-GEO project has demonstrated that PERMIS can be used to protect a GT4 accessor service which acts as a proxy to the data source that can enforce access control via PERMIS. The adoption of standards-based solutions for each technology challenge (SAML, X.509, JSR-168/286) ensures broad coverage of most scenarios, and a degree of future-proofing not possible with the many bespoke solutions on offer today. Once the final challenge of the secure generation of X.509 proxy certificates for users has been met, we believe this framework offers a scalable method to enforce multi-layered security across administrative and organisational domains.

REFERENCES

- [1] R. Housley et al., "RFC 2459: Internet X.509 Public Key Infrastructure Certificate and CRL Profile," IETF, 1999, <http://www.ietf.org/rfc/rfc2459>
- [2] S.Cantor et al., "Shibboleth Architecture: Protocols and Profiles", Internet2-MACE (Document ID: internet2-mace-shibboleth-archprotocols-200509) <http://shibboleth.internet2.edu/docs/internet2-mace-shibboleth-arch-protocols-latest.pdf>
- [3] D.W.Chadwick and A.Otenko, "The PERMIS X.509 Role Based Privilege Management Infrastructure", Future Generation Computer Systems, 19(2) (Elsevier Science BV), 2002, pp. 277-289.
- [4] JSR-168: Portlet Specification, <http://jcp.org/en/jsr/detail?id=168>
- [5] C. La Joie, "Trusted Delegation of Privileges in an N-Tier Environment", http://middleware.internet2.edu/webiso/docs/draft-lajoie-trust_and_delegation-02.html
- [6] E. Maler et al., "Assertions and Protocols for the OASIS Security Assertion Markup Language", OASIS (Document ID: oasis-sstc-saml-core-1.1), <http://www.oasis-open.org/committees/security/>
- [7] UK Access Management Federation for Education and Research, <http://www.ukfederation.org.uk/>
- [8] G.Inman, D.Chadwick and N.Klingenstein, "Authorisation using Attributes from Multiple Authorities – A Study of Requirements", Presented at HCSIT Summit - ePortfolio International Conference, 16-19 October 2007, Maastricht, Netherlands
- [9] Shib-Grid Integrated Authorisation (Shintau) Project website, <http://sec.cs.kent.ac.uk/shintau/>
- [10] D.W.Chadwick, A.Otenko, E.Ball, "Role-Based Access Control with X.509 Attribute Certificates", IEEE Internet Computing, Mar-Apr 2003, pp. 62-69
- [11] J.Novotny, S.Tuecke, V.Welch, "An Online Credential Repository for the Grid: MyProxy", Proc. of the Tenth International Symposium on High Performance Distributed Computing (HPDC-10), IEEE Computer Society Press, Aug 2001
- [12] MAMS Shibbolized GridSphere 3, <http://www.federation.org.au/software/shibbolized-gridsphere-3.0.5.zip>
- [13] The Open Geospatial Consortium (OGC) Web Feature Service (WFS) standard, <http://www.opengeospatial.org/standards/wfs>

Minerva Access is the Institutional Repository of The University of Melbourne

Author/s:

Watt, J;Sinnott, R;Jiang, J;Doherty, T;Higgins, C;Koutroumpas, M

Title:

Tool Support for Security-oriented Virtual Research Collaborations

Date:

2009-01-01

Citation:

Watt, J., Sinnott, R., Jiang, J., Doherty, T., Higgins, C. & Koutroumpas, M. (2009). Tool Support for Security-oriented Virtual Research Collaborations. Liao, XF (Ed.) Jin, H (Ed.) Zheng, R (Ed.) Zou, DQ (Ed.) 2009 IEEE INTERNATIONAL SYMPOSIUM ON PARALLEL AND DISTRIBUTED PROCESSING WITH APPLICATIONS, PROCEEDINGS, pp.419-424. IEEE COMPUTER SOC. <https://doi.org/10.1109/ISPA.2009.49>.

Publication Status:

Published

Persistent Link:

<http://hdl.handle.net/11343/28803>