WISERD

Wales Institute of Social & Economic
Research, Data & Methods
Sefydliad Ymchwil Gymdeithasol ac
Economaidd, Data a Dulliau Cymru

WISERD DATA RESOURCES

International Access to Restricted Data – A Principles-Based Standards Approach

WISERD/WDR/003

Felix Ritchie

November 2011

SEFYDLIAD YMCHWIL GYMDEITHASOL AC ECONOMAIDD, DATA A DULLIAU CYMRU
WALES INSTITUTE OF SOCIAL & ECONOMIC RESEARCH, DATA & METHODS WALES INSTITUTE OF SOCIAL & ECONOMIC RESEARCH, DATA & METHODS WALES INSTITUTE OF SOCIAL & ECONOMIC RESEARCH, DATA
SEFYDLIAD YMCHWILL GYMDEITHASOL AC ECONOMAID, DATA A DULLIAU CYMRU SEFYDLIAD YMCHWILL GYMDEITHASOL AC ECONOMAID, DATA A DULLIAU CYMRU SEFYDLIAD

Authors
Felix Ritchie

Address for Correspondence:
Microdata Analysis and User Support
Office for National Statistics
Cardiff Road
Newport
South Wales
NP10 8XG

Email: felix.ritchie@ons.gsi.gov.uk or felix.ritchie@virgin.net

WISERD Hub Contact:

Cardiff University
46 Park Place
Cardiff
CF10 3BB

Tel: 02920879338

Email: wiserd@cardiff.ac.uk

Abstract

Access to restricted microdata for research is increasingly part of the data dissemination strategy within countries, made possible by improvements in technology and changes in the risk-benefit perceptions of NSIs. For international data sharing, relatively little progress has been made. Recent developments in Germany, the Netherlands and the US are notable as exceptions. This paper argues that the situation is made more complex by the lack of a general coherent risk-assessment framework. Discussions about whether something should be done become sidetracked into discussions about how procedural issues would constrain implementation. International data sharing negotiations quickly become bilateral, often dataset-specific, and of limited general value.

One way forward is to decouple implementation from principles. A principles-based risk-assessment framework could be designed to address the multiple-component data security models which are increasingly seen as best practice. Such a framework allows decisions about access to focus on legal-procedural issues; similarly, secure facilities could be developed to standards independent of dataset-specific negotiations. In an international context, proposals for classification systems are easier to agree than specific multilateral implementations. The paper concludes with examples from the UK and cross-European projects to show how such principles-based standards could work in practice.

Acknowledgements and reservations

This paper is based on presentations given to the 3rd Workshop on Data Access, and IASSIST 2010, as well as submissions made to ESSNet (2010) with Richard Welpton of the Secure Data Service. I am grateful to all workshop and conference participants for helpful comments.

The views expressed in this paper are those of the author and do not represent a policy statement by the Office for National Statistics.

1. Introduction

In recent years, significant and widespread improvements have been made in the availability for research purposes of confidential microdata from government sources. Whilst implementations differ widely between countries, greater awareness of the value of such data and an increasing confidence in the ability of National Statistics Institutes (NSIs) or other bodies to manage the risks of access means that use of confidential microdata for research is becoming the norm in developed countries.

The same cannot be said for sharing of confidential data internationally. Traditionally, only highly aggregated data has been made available across national borders. Researchers wanting to carry out multi-country studies have either had to restrict their analysis to these aggregates, or work with collaborators performing isolated local analyses. This limits the inferences that can be drawn. For example, the identification of common influences across national healthcare systems on life expectancy cannot be statistically tested without allowing for the international interaction of variables.

There are two main barriers to this, creating a chicken-and-egg problem:

(a) Internationally agreed ‘secure’ technical solutions are not developed because the legal framework is not in place to allow this technology to be exploited

(b) Legal arguments over what is allowable are not being considered because there is no agreement on how the security of a proposal could be technically ensured

In addition, NSIs have very limited incentives to invest resources in exploring these issues. As Ritchie (2010b) notes, traditional risk/reward perspectives skew the incentives for NSIs to release confidential data even within countries. Across countries, the risks borne by and effort needed from the NSI are larger and the benefits to the NSI much less direct, attenuating the risk/reward dichotomy.

This description is of course a simplification. In practice, there a number of international situations where barrier (a), developing specific technical solutions without an overarching legal framework, is being overcome. Current examples include

- The IPUMS project www.ipums.org to harmonise and share confidential but anonymised Census microdata (McCaa et al, 2011)

- 'Mesodata' models, as used in Airaksinen et al (2008) for example, where semi-aggregated microdata are created from country-specific disclosive microdata with a view to a particular type of analysis; knowing in advance what models are to be run allows non-disclosive aggregates to be used to generate interactions which may not be possible from independent country studies
- IAB's 'RDC-in-RDC' model (RDC is a research data centre), using remote access technology to allow researchers in the US to access German microdata (Bender, 2010)
- The Dutch statistical office (CBS) use of contractual arrangements within the umbrella of European law to allow Italian researchers to have live access to Dutch data
- The Lissy remote job submission system allowing researchers around the world to run queries on confidential earnings data (Kruten, 2008)

This paper is concerned with the second barrier: can legal or procedural frameworks be developed without fixing on specific solutions? The examples above are specific case studies, and are the results of single or repeated bi-lateral agreements. Committed individuals are making progress by solving specific cases which can then be used to set precedents. Is it possible to develop a general framework for taking forward international data sharing, ideally without going into specifics?

This paper argues that this is possible by the creation of a coherent comprehensive frame of reference based upon common standards decoupled from implementation. In this, avoiding specifics is not just desirable; it is essential. To develop a consistent frame of reference for dealing with security issues, that framework should focus on the purpose of risk assessments, not on how those purposes are fulfilled. The framework is positive, not normative: a way of describing outcomes, not a specification for 'the' way to do things. That may come, but the first stage is to agree the language for describing concepts.

The next section examines the 'decoupling' argument in more detail. Section 3 imagines a specific framework and then shows how it could be used to both describe and recommend alternative security frameworks. Section 4 provides a long-term vision of how these proposals could be used as the basis for a true data sharing network, and examines some of the practical difficulties that would arise. It also looks at how such a framework copes with new technological developments such as the 'cloud' or grid computing. Section 5 concludes, offering some thoughts on planning horizons and whether the journey has value even if the aim is ultimately too ambitious.

A companion paper, Ritchie and Welpton (2011) investigates the network model further and explores the implementation of the framework elaborated here; considers how developments such as thin-client models affect the practicality of solutions; and discusses the implications for competition and innovation.

The issues discussed are relevant to all an interest in sharing data internationally. However, for simplicity of exposition this paper assumes NSIs (national statistical institutes) are the data owners interested in providing access.

2. Decoupling principle from practice

There is no universal view on the appropriate technologies for sharing data internationally. This should not stop agreements about principle, but in practice discussions about data sharing quickly become discussions about implementation. Decisions about ‘what’ and ‘why’ turn into discussions about ‘how’.

But as Ritchie (2011) notes, there are many different ways to provide secure access to confidential data, ranging from anonymised public-use files to remote-access research facilities and synthetic data. With the possible exception of the last, these are all well-understood and implemented in different countries in different ways. An NSI wishing to acquire a secure solution can almost pull one ‘off the shelf’. More importantly, an NSI can focus on its own specific legal or procedural requirements comfortable in the knowledge that a solution can be implemented, somehow.

This is the basis for decoupling: decisions are made on principle, based upon what the NSI aims to achieve. The wide variety of ways to achieve the same end relieves the NSI from concerns about implementation. An analogy is with computer networks. A standard specifies how computers talk to each other over the internet; the actual devices used for this are irrelevant.

For bilateral negotiations, it could be argued that this argument over whether principles or solutions are at stake is just hair-splitting, as the two parties will have to agree on a specific solution. But the aim of this paper is provide a basis for multilateral agreement and unilateral action; one outcome should be improved efficiency of bilateral discussions. Again taking the example of computer networks, some organisations do have dedicated links to partners. However, for most companies it is cheaper and easier to connect to the internet using open standards.

The focus on principles works because it reflects the thinking of the NSI. Although prospective collaborators may ask questions about technical security of solutions, the real underlying interest is, for example, "can my data be transferred to the internet?" Decoupling makes these questions the basis for any agreement, not any specific technology.

To build a general-purpose framework also requires an acknowledgement that the questions NSIs ask will be different. Some will be more concerned about ensuring that the data itself is confidentialised; others that any statistical outputs are non-disclosive; others that the system cannot be hacked into. Hence, a useful (and thus credible) set of standards needs to have multiple dimensions and multiple levels of compliance within these dimensions.

Using this framework to describe implementations may be interesting but ultimately is of limited value. The ultimate worth comes from using the framework as an accreditation standard. An NSI wishing to deposit data can consider what security questions it wants answering – not the specific implementation. Similarly, a research data centre (RDC) could, for example, identify itself as being compliant to a given standard in a particular dimension. The key point is that the standard becomes the focus for discussion, not the implementation.

3. The framework in practice: examples of standards and use

This section details an example framework to show how useful standards could be defined and used. These standards arise from the author's experience in the UK. It could be argued that there are too few or too many dimensions, for example, or that the risk assessments are wrong. These are valid points but not relevant here, where the purpose is to illustrate the way forward using a familiar model. Section 3.4 considers the lessons this simple model holds.

3.1 Defining the dimensions of the framework

The dimensions of this model are taken from the VML Security Model, a common framework for defining security as a set of 'safe' characteristics, which are assessed independently and jointly; see Ritchie (2009). This model recognises that, for example, greater trust being placed in 'safe people' means that the NSI does not need to rely entirely upon restrictive IT systems.

For the purposes of this exercise, the standard dimensions of projects, people, data, settings and outputs need some refinement, as shown in Table 1.

Table 1 Dimensions of risk in the extended VML security model

Dimension	**Subcategory**	**Meaning**
Safe projects		Project meets legal/ethical/process requirements
Safe people(1)	Knowledge	Researchers have sufficient knowledge to use data safely
Safe people(2)	Incentives	Researchers have sufficient incentives to use data safely
Safe data		There is protection in the data itself
Safe setting(1)	Access	Connection to the data is secure
Safe setting(2)	Networks	Opportunity to move data to other networks/media is limited
Safe outputs		Statistical outputs are checked for confidentiality breaches

More dimensions and subdimensions could be added; for example, ‘safe outputs’ could be broken down into both the number of outputs that are checked and the standards of checking that are applied; or ‘safe settings’ could include some system specification such as ISO27001. But these serve for illustrative purposes.

3.2 Defining the criteria

In Table 2, in each of the security dimensions a level of protection has been indicated. These are then allocated a score from 0 (implying no protection) to 4 (implying the best level of protection available). So, to take the top row, a system which does no checks on researchers or projects bar the necessary administrative processes scores 0; a system whereby all projects are read and reviewed by an expert gains the maximum score. Moving down, the ‘safe data’ scores directly reflects the identification possibility. Finally, along the bottom row outputs are scored from 0 (where no checking is done on outputs to see whether they breach confidentiality) to 4 (where nothing is released until it has been scrutinised by the NSI.

How does this relate to the principles of access? Consider the ‘networks’ option. The scores can be rephrased as:

- “There need be no further restrictions on where the data can be transferred to”
- “I don’t want users to be able to easily share data with the internet”
- “I don’t want users able to transfer data to any mobile devices”
- “I want users to work in a restricted area of the network”
- “I want users physically isolated from all other systems”

Table 2 Illustrative standards

Safe…	**Level of Protection**				
	0 **No protection**	**1**	**2**	**3**	**4** **Strong protection**
Projects	Administrative processes only	Check researcher background	Check use is for statistical purpose	Review by support officers able to critically assess feasibility and need for data	Review by support officers able to critically assess impact of research
People (1) knowledge	Administrative processes only	Check researcher background	Written assent to conditions of access	Passive training	Active training
People (2) incentives	No effective sanctions	Procedural sanctions only	Mix of civil, criminal or procedural sanctions	Civil, criminal and procedural sanctions	Civil, criminal, procedural and institutional sanctions
Data	No data protection	Removal of direct identifiers	Identification within RDC environment unlikely	Identification outside RDC unlikely	Public use microdata
Setting (1) Access	No restrictions	Access only from limited sites with no supervision	Access from secure networks with no supervision	Access from secure networks with occasional supervision	Access from secure networks with continual supervision
Setting (2) Networks	No restrictions on data transfer	No internet access	No internet, local/mobile storage or printers	No access to other parts of network	No network connection, no mobile storage
Outputs	No checks	Random checks	Random plus targeted partial checking	Full checking except for 'experienced' researchers	Full checking

This table can be refined further. For example, what is meant by 'active training'? In the world of UK RDCs, this means compulsory attendance at one of the 'safe researcher training' programmes run by ONS or the UK Data Archive. This does not make the criteria overly prescriptive; both of these training programmes are being aligned to best practice standards identified in Brandt et al (2010), which again defines 'what needs to be known' rather than 'how this is done'. Hence it is possible to refine these criteria substantially without specifying particular implementations.

3.3 Applying the standards

Table 3 applies the above criteria to a number of different solutions currently existing; note that these scores are the author's perspective and do not necessarily reflect the view of the service providers

Table 3 Applying the standards

Safe…	VML	SDS	RADL	LISSY	IPUMS	UKDA EUL	Intern et
Projects	4	3	3	3	3	0	0
People (1) knowledge	4	4	2	2	2	1	0
People (2) incentives	3	4	2	2	4	2	0
Data	1	2	3	3	3	3	4
Setting (1) Access	3	2	1	1	0	0	0
Setting (2) Networks	3	3	4	4	0	0	0
Outputs	4	4	1	4	0	0	0

VML=ONS Virtual Microdata Laboratory (UK); SDS=Secure Data service (UK); RADL=Remote Access Data Laboratory (Aus/NZ); LISSY=UK/Lux remote job submission; IPUMS=anonymised Census microdata (US); UKDA-EUL=UK Data Archive End User Licence (UK)

Thus it can be seen that the SDS has stronger incentives for people to act safely compared to the VML. In contrast, the VML has more detailed data and more direct supervision of researchers and takes more interest in the non-statistical aspects of the project. Alternatively, IPUMS relies heavily upon the inherent safety in the data and on its institutional agreements to ensure good practice.

So far, this identifies how different models solve the security issue, allowing an NSI to start considering which are the dimensions that are most relevant for the data release under consideration. Note that an NSI can use multiple channels: in the UK, ONS supplies VML, SDS, IPUMS, and the UK Data Archive with data (see Ritchie, 2009).

The next obvious question is how this can be used actively to design secure systems. For example, ESSNet (2010) recommended, in the short term, using the existing RDC infrastructure in Europe to trial pan-European data access. The framework specified here could be used as the basis for defining perhaps three levels of European ‘safe centre’; see Table 4.

Table 4 Specific standards for European-accredited secure RDCs

Safe…	‘minimum’	‘best practice’	‘maximum security’
Projects	2	3	4
People (1) knowledge	3	4	4
People (2) incentives	3	4	4
Data	1	1	1
Setting (1) Access	3	3	4
Setting (2) Networks	3	3	3
Outputs	3	4	4

Note that ESSNet (2010) envisaged using the existing RDC infrastructure but the point of the standards model is that it is infrastructure-independent. If a remote-job model such as Lissy or RADL can meet the requirements, then this approach says that it is a valid alternative to an RDC as far as security is concerned.

3.4 Is the model extensible?

[Ref] distinguishes between 'systems' and 'networks' in access to data for research. Broadly, a system brings researchers together via a monolithic solution designed, and possibly implemented, by a single authority; this implies ownership and a central architecture. In contrast, a network focuses on gateways and communication protocols; what goes on behind those gateways is of no concern to the network.

The focus on the network rather than system characteristics decentralises the decision making process and is designed to encourage innovation in solutions. For example, the world-wide web was developed without the need to change the basic operation of the internet. In the principles-based description of data access, not specifying particular solutions should encourages alternatives to be explored.

For example, one area of current interest to data owners is the emerging field of 'cloud computing', buying 'live' computing services from third parties. Beecher and Leclere (2010) give an example of using a cloud model to provide access to confidential data. This can be accommodated in the above framework, with data owners being able to delineate the difference between, for example, using cloud services and building their own solution. More importantly perhaps, the standards could be turned on its head: the framework could be used to define the service levels expected by the data owners.

Two other concepts suggesting opportunities for data owners are 'grid computing' (using multiple linked computers to carry out major processing tasks in parallel) and distributed storage (where the data files are kept separately and are combined dynamically at the time of processing). The latter in particular is appealing for international data access, because it suggest the source data can always be stored in the country of origin.

In both these examples, having a common frame of reference is valuable because it allows developers to consider when they might be crossing boundaries between 'acceptable' and 'not acceptable'. In a world of virtual or remote processing, the concept of where data actually 'is' becomes a real issue. The framework can help to identify which are the relevant points of concern.

Of course, the downside of setting gateways and protocols is that they themselves can become blocks to progress, however well-intentioned the original plan. This is particularly the case where network protocols cannot be changed without massive expenditure or disruption. By focusing on an abstract level, some of the risk of this may be avoided, but it is very possible that a framework devised now might be quite unsuitable a few years hence.

However, this is a risk to be managed. The decentralisation may help this process by providing the incentives for competition. There is a positive example in the case of the internet: the enormous success of the web, built upon an network designed in the 1970s for a very different purpose, has led to the redevelopment of the internet protocols in ways which the original web designers never envisaged.

4. Issues of implementation

The framework described above is a simple back-of-the-envelope model designed to illustrate the point. Assuming, however, that this was felt to be a useful approach to take, how could it be implemented?

The most effective way to kill any development would be to insist on a framework being adhered to. The point of focusing on the principles of security to build a frame of reference is that agreeing a framework can be done without any explicit commitment to meet any appropriate 'standard' This makes it easier to get agreement.

The framework is only valuable is it is found to be useful. The most likely way for it to be useful is to save time or effort, or to improve credibility. For example, both the VML and SDS are aiming to adhere to the SDC standards laid down in Brand et al (2010). Both RDCs already operate to a similar standard, but by describing themselves with reference to an independent one they do not need to write their own statement of operating standards; moreover, they can compare how they operate against other countries who have quite different technical systems.

As noted above, the framework can also serve as a way of specifying standards to be achieved when designing new systems. One possible three-stage development route then is

1. **Definition**: a principles-based reference framework is defined

2. **Retrospective adoption**: data owners/infrastructure providers begin using the framework to describe their systems
3. **Prospective adoption**: the framework is used as a design criterion:
    a. Data owners begin using the framework to define their requirements
    b. Infrastructure builders begin using the framework to define their systems

Most importantly, each of these stages is voluntary.

This approach by itself does not solve all the issues of international data sharing. In particular, it has nothing to say about the legality of data sharing. It is not intended to do this; it is meant as part of a suite of methods for separating complex issues into more digestible fragments.

NSIs take decisions on whether to invest in research infrastructure on the basis of cost, benefits and the security risks involved. The principles-based framework is designed to take the latter out of the equation, by making clear that all levels of security are available, albeit it at different cost. The framework is implementation-independent by design, and it requires NSIs to focus on what they want, rather than how the aim to achieve.

In this, this paper complements Ritchie (2010a), which argued that strategic decisions about access should be taken without direct reference to law, technology or risk; these are *enablers*, factors which can allow or constrain objectives but do not define those objectives. That paper sought to break the link between objectives and solutions; this paper reinforces that point by shifting the focus to the characteristics of solutions rather than any particular (and country-specific) implementation.

5. Summary and thoughts on future development

The framework defined above is necessarily basic. Even if the framework were agreed as it stands there are many refinements needed. For example, ‘active training’ is marked as giving the most security, but what does this consist of? Should there be an exam and certification? Do all countries want the same training? How active is ‘active’? On outputs, is automatic checking of all outputs safer than manual checking of only some? Most importantly, can a linear model be devised – should there be, for example, several different ways of scoring top marks in a dimension?

Nevertheless, even in outline this approach may help to break through the tangled concepts of international data access, by separating out the security components which are important from those which are not.

Just as important as the concept is the scope for agreement. Countries are more likely to support a framework which focuses on describing effectively, rather than one which seeks to impose a standard. The lack of prescription for favoured methods or technology, and the explicit recognition that a variety of approaches have value in achieving access outcomes, makes acceptance more easy. Even the long term aim, of such an approach being used to specify acceptable implementations, assumes that countries will want to do this because they find it beneficial, not because it is required.

From the NSI’s internal perspective, the framework can also bring gains. Descriptions of what an NSI’s security policy aims to achieve are easier to align with corporate goals, compared to a more implementation-specific viewpoint. In addition, setting security goals in terms of standards allows for multiple implementations to be covered by the same corporate policy, and for those implementations to be changed as the situation arises without the need to change the policy.

As noted, this standards-based approach only addresses some of the problems surrounding international data sharing, and omits to consider, for example, legal implications. Also, the standard is necessarily voluntary. Hence, the vision of international data sharing discussing focusing solely on principles may be unrealisable. Is the journey then worth it? This paper argues that it is, for three reasons.

- A common way of describing disparate solutions has merit in itself. For example, the multiple access paths in ONS’ ‘data access spectrum’ (see Ritchie, 2009) can all be placed in the above framework, allowing one to easily see where the differences between options arise.
- The framework focuses on the principles of security, rather than the technology. A key recent development has been the recognition that security should be an outcome, not a solution. This framework makes that focus explicit.
- Everything known about networks shows that clear, effective standards are essential for their use and development; and that standards encourage innovation by lowering the cost of connecting to the network. There are many uncertainties about the

development of technology, as well as approaches to risk; clarifying the network gateways allows innovation to flourish without needing to redesign systems.

Overall, these suggest that the journey may be worth the effort even if the ultimate end is unattainable. This is a long-term proposal. Even if the above framework were accepted now as it stood – which would require a massive change in approach by NSIs – the move from retrospective to prospective adoption is likely to take some years. Moreover, the shift from implementation to objectives may have unexpected effects. For example, a natural implication of this approach is a stronger case for the outsourcing of data management.

At present, international data sharing is a long way from any prospect of general arrangement or approach. However, over the last ten years access to microdata within many countries has changed beyond recognition – both in technology and in the approaches to risk and security. Will there be a similar shift in international data sharing? It is not possible to say, but this paper has tried to suggest a way in which the discussion may be usefully taken forward.

References

Airaksinen A., de Panizza A., Bartelsman E., Hagsten E., van Leeuwen G., Franklin M., Maliranta M., Kotnik P., Stam P., Rouvinen P., Farooqui S., Quantin S., Svanberg S., Clayton T. & Barbesol Y. (2008) *Information Society: ICT impact assessment by linking data from different sources; Final Report*, Eurostat http://epp.eurostat.ec.europa.eu/portal/page/portal/information_society/documents/Tab/ICT_IMPACTS_FINAL_REPORT_V2.pdf

Beecher B. & Leclere F. (2010) *Exploring new methods for protecting and distributing confidential research data*, http://www.slideshare.net/bryanbeecher/exploring-new-methods-for-protecting-and-distributing-confidential-research-data

Bender S. (2010) *The need for a RDC network and for sharing data in an RDC approach*, http://www.safe-centre.eu/sitefiles/07_Day%201_Session%203_Stefan%20Bender.pdf

Brandt M., Franconi L., Guerke C., Hundepool A., Lucarelli M., Mol J., Ritchie F., Seri G. and Welpton R. (2010), *Guidelines for the checking of output based on microdata research*, Final report of ESSnet sub-group on output SDC http://neon.vb.cbs.nl/casc/ESSnet/guidelines_on_outputchecking.pdf

ESSNet (2010) *Decentralised access to European Microdata*, http://www.safe-centre.eu/index.php?s=3

Kruten T. (2008) *A remote access solution: The Lissy System at LIS Asbl* http://epp.eurostat.ec.europa.eu/portal/page/portal/research_methodology/documents/7_2_LIS.pdf

McCaa R., Ruggles S., Sobek, M. & Thomas W. (2011) *IPUMS-International: Free, Worldwide Microdata Access Now for Censuses of 62 Countries--80 by 2015,* 58th International Statistical Institute, Dublin, Ireland, 21-26 August, 2011 http://www.hist.umn.edu/~rmccaa/sts065_ipums_international_future_microdata_access.pdf

Ritchie (2009) “UK Release Practices for Official Microdata”, *Stat. J. of the IAOS*, v26:3-4 pp103-111 http://iospress.metapress.com/content/08km250w1n536w85/?p=db53a55c3325436c89c4632a67edc846&pi=6

Ritchie F (2010a) “Access to sensitive data: satisfying objectives rather than constraints “, mimeo, Office for National Statistics

Ritchie F. (2010b) “Risk assessment for research access to confidential microdata”; presentation to John Deustch Insitute and WDA http://www.felixritchie.co.uk/publications/wda_risks_v1.ppt

Ritchie F. (2011) “Methods for analytical access to confidential data”, paper prepared for OECD Paris Microdata Group

Ritchie F. and Welpton R. (2011) “Access without boundaries”, IASSIST 2011 presentation