

# Reducing system toil in a university library

Reducing  
system toil

Emily E. Boss and Nathan Gerth  
*University of Nevada, Reno, Nevada, USA and  
University Libraries, Reno, Nevada, USA*

173

## Abstract

**Purpose** – The purpose of this paper is to expand the information technology concept of toil by applying it to system analysis for academic libraries. System mapping and toil definitions are used to measure costs and benefits of maintaining library information systems. The concept of toil provides a vehicle to shape the decisions made by managers when refining or building new library system infrastructure.

**Design/methodology/approach** – Using a system map of all library systems, the definition of toil was used to identify highly manual or high toil systems. This mapping was used to illuminate the need to migrate or eliminate certain library systems to managers in order to reduce the amount of toil undertaken by library employees.

**Findings** – Two library systems were identified for migration to more automated systems and two library systems were recommended for elimination in order to reduce the overall amount of toil present in the library systems ecosystem.

**Practical implications** – Using system mapping and toil definitions are helpful in identifying and relaying the cost of outdated systems to library managers.

**Originality/value** – This paper provides library managers an easy way to assess the ecosystem and efficiency of library systems without the need for technical expertise. It also expands the use of information technology concepts into the field of academic libraries.

**Keywords** Information technology, Academic libraries, System management, System mapping, System migrations, Toil

**Paper type** Case study

Received 9 November 2019  
Revised 29 February 2020  
Accepted 29 February 2020

## Introduction

In 2016, the University of Nevada, Reno Libraries acquired the papers of Senator Harry Reid. Consisting of the records generated by the Senator's offices over the course of his 30-year career in the US Congress, the collection encompassed over 1,300 linear feet of physical materials and 12 terabytes of electronic records. As such, the Reid Collection was the largest acquisition ever made by the University Libraries. The extent of the collection quickly raised questions about the libraries' information systems and their ability to scale to meet the demands of the collection, resulting in a deep assessment of existing systems infrastructure.

While case studies of implementation and migration efforts are widely discussed in library literature, there are few models that help library managers evaluate the processes that tie library information systems together (Buljung and Johnson, 2017; Mayo and Bowers, 2017; Neatrou *et al.*, 2017; Tuyl *et al.*, 2018). As a result, a library working group at the University of Nevada, Reno Libraries turned to models from information technology to measure the negative or positive effects of adopting platforms and their associated integrations. In particular, the team began to measure these changes in terms of "toil."

Proposed by Google as a means to improve the morale and effectiveness of its site reliability engineers (SRE), the model of "toil" has since become an influential approach to measuring the effect of systems on the people who maintain them (Rau, 2016; Catchpoint, 2018, 2019). It offers a clear set of parameters for measuring whether or not a system support action constitutes "toil" and therefore should be eliminated. Although there is no evidence of it being directly applied in library settings, the application of this model does closely parallel adoption of other informational technology models within academic libraries, such as agile



Library Management  
Vol. 41 No. 4/5, 2020  
pp. 173-181  
© Emerald Publishing Limited  
0143-5124  
DOI:10.1108/LM-11-2019-0081

software development and technical debt (Clair, 2016; Critchlow *et al.*, 2010; Lindström and Malmsten, 2008).

The purpose of this paper is to expand use of the concept of toil by applying it to system analysis for academic libraries. Emphasis will be placed on the use of system mapping and toil to measure costs and benefits of changes for the employees who maintain library information systems. In this manner, toil, as a concept, provides a vehicle to shape the decisions made by managers when refining or building new library system infrastructure. The extension of this model to academic libraries also shows the further blurring of lines between the concepts used to support information technologies and those found in libraries and archives. By thinking about systems in an information technology sense and adopting a concept like toil, library leaders can approach the systems they manage in a comprehensive manner, while also bolstering library staff members who are weighed down with manual system processes. This paper begins by outlining the nature of the toil model, moving on to consider its application at the University of Nevada, Reno Libraries, and then by considering the outcomes of its use.

### Model

Google defined the concept of toil when it began its site reliability engineer (SRE) program in 2003. At Google, SREs are tasked with the maintenance of running production systems as well as engineering the code needed to move applications forward. SRE came about because of the tension between wanting to maintain reliable services for users and wanting to deliver new innovative features. Google defines toil as, “work tied to running a production service that tends to be manual, repetitive, automatable, tactical, devoid of enduring value, and that scales linearly as a service grows” (Rau, 2016). Since toil limits the potential of systems and the time of those who run them, its definition is helpful in identifying the difference between worthwhile administrative tasks that may be tedious compared to tasks that are hands-on, but purely operational. Tasks considered toil will depend on the size and shape of your production service. Rau (2016) points out that the more attributes of toil a task has, the more likely it is to be toil, but worthwhile tasks may also share some toil attributes. He goes on to define the attributes of toil further:

- (1) Manual: hands-on time needed to perform the task.
- (2) Repetitive: work happening over and over.
- (3) Automatable: human judgment is not needed.
- (4) Tactical: work that is reactive rather than proactive.
- (5) No enduring value: the task does not improve the service.
- (6) Scalability: work needed to perform the task grows as the service grows.

Google has used this definition to keep SRE’s tasks under a 50% toil threshold so that the focus is on proactive engineering such as improving reliability, performance or utilization instead of simply running the service (Rau, 2016). Without the focus on defining and reducing toil, organizations like Google are at risk of inactivity where their systems and employees are overwhelmed. The consequences of a high toil environment can be depreciation on morale, confusion in employee roles or work outcomes, the stagnation of progress, and, ultimately, attrition (Rau, 2016).

Toil can affect any organization running a production service but especially organizations running various tools not designed to work together. Incidences of this are high in libraries where the acquired tools often perform a specific function but data or other outputs must be exchanged manually between many tools in the ecosystem. The more distinct tools are

acquired, the more work library staff members undergo attempting to create a cohesive ecosystem of library tools. Adding to this the often high cost or availability of migrating these distinctive systems to more current technology, library staff members also end up handling library tools from many different technology eras. Given these factors, evaluation of toil as it relates to the work done between library tools can reveal to library decision-makers the amount of toil undergone by staff members and which tools contribute the most to toil. This evaluation can help libraries attempt to avoid the consequences of high toil.

## Method

The application and evaluation of the toil model took place in three stages: systems mapping, applying the model and quantitative analysis of the model's impact. The Libraries were already invested in one large-scale tool migration at the time. The migration involved moving from a local CONTENTdm instance to a hosted Islandora instance. This process took place over the course of six months in 2018. During this migration, a library working group from several units including Metadata and Cataloging, Digital Services, Discovery Services, Access Services, Library Information Technology and Special Collections convened to begin evaluating the impact of this migration on the wider array of systems in the Libraries' portfolio.

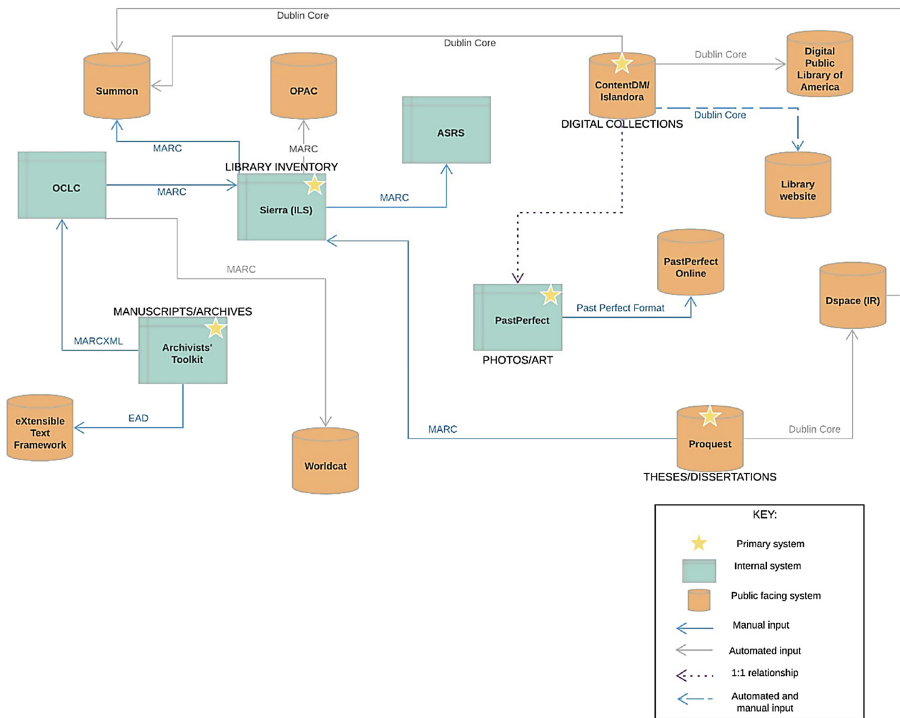
In April of 2018, seven library employees identified as key stakeholders from the aforementioned units met to draft and explore the creation of a library systems map. Informally, members of the working group began discussing specific library systems integral to their day-to-day work. Once the systems were mapped, the group laid out the data pathways between the systems specifically relying on those whose jobs intersected most directly with the systems in question.

### *Mapping*

The library working group followed a "data governance" model to mapping the existing systems. Current literature defines data governance as, "a companywide framework for assigning decision-related rights and duties in order to be able to adequately handle data as a company asset" (Otto, 2011b, p. 47). The growing interest in commercial and academic institutions to create governance structures reflects growing ubiquity of data and transition from systems that manage physical materials to those that manage information assets (Otto, 2011a, p. 241). The actual system mapping involved identifying which of the platforms served as "systems of record" or the "authoritative data source for a given data element or piece of information" and delineating data pathways from one platform to the next (VoltDB, 2017). It is important to note that the working group purposefully excluded platforms from the map if they did not fall under the full control of the library. The team captured this information in a high-level flow chart using a standardized visual vocabulary (Figure 1).

### *Applying the model*

Applying the model involved analyzing the data flows between platforms. The library working group categorized each interaction as being either manual, automated or automated and manual input. Specific steps in the systems interaction process were interrogated according to Rau's (2016) six attributes of toil. Specific questions were asked of those who worked in the systems about their tasks including if tasks could reasonably be done a different way that might mitigate manual actions such as activating a batch publishing process instead of individual publishing of data. This step allowed the group to begin interrogating each interaction to determine whether or not each constituted, by definition, toil. The group then added visual markers for these interactions to the systems map. As a result,



**Figure 1.**  
System map for the  
University of Nevada,  
Reno Libraries

the team was able to calculate the number of data interactions in the legacy systems whose maintenance constituted toil.

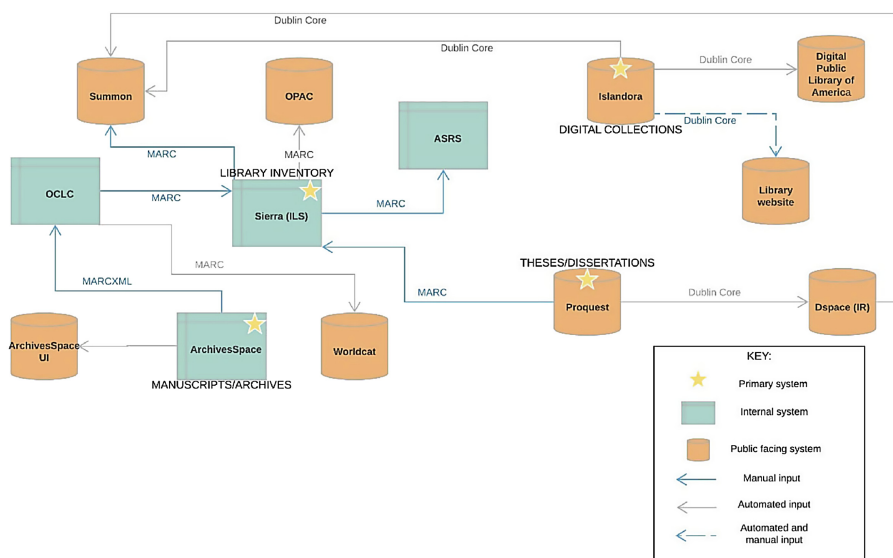
*Quantitative analysis*

After completion of the CONTENTdm to Islandora migration, the library working group then remapped the platforms and their interactions. By remapping the system, the group was able to measure the change over time in both visual and quantitative manners. Again, the majority of the group’s attention turned to the data interactions and the degree to which the associated maintenance actions constituted toil. The end product of this process was then compared against the results from the initial evaluation. In doing so, the team was able to measure the net gain or reduction of toil in the Libraries’ systems (Figure 2).

**Findings**

The initial system mapping before the Islandora migration project gave each data exchange point a net value of automated or manual depending on how many attributes of toil were exhibited in the exchange. Given that 7 of 15 data change points were determined to be manual, an approximate 63% overall toil rate was calculated for library systems. One exchange was classified as a worthwhile administrative task due to its processes being equally manual and automated and therefore was not included in the findings.

Visually, it was obvious to the library working group which applications were contributing the most to the high toil rate. The first was the use of eXtensible Text Framework (XTF) for finding aids. XTF is an open-source software developed by the



**Figure 2.** System map for the University of Nevada, Reno Libraries after system changes

California Digital Library to provide access to digital content. The Library was running XTF locally in order to display EAD finding aids that were input into Archivists' Toolkit, the finding aid repository software. XTF was a tool adopted to get around the fact that Archivists' Toolkit did not provide a public facing display. However, routing the EAD finding aids from Archivists' Toolkit to XTF was a fully manual process that required staff members to export files and then reupload them every time a change was made to a finding aid. This meant that as the number of finding aids grew, the number of changes that needed to be pushed manually to XTF also grew, which required more and more manual hands-on staff time and thus resulted in more toil. The second high toil system was PastPerfect, a museum software that was adopted in order to accession and describe photographs that lived in Special Collections. Unlike Archivists' Toolkit, PastPerfect came with an online display (PastPerfect Online); however, in order to push content into PastPerfect Online, each record needed to be manually published when it was finished. This was not an insurmountable hurdle and in general did not create a lot of undue toil. However, the Libraries' decision to maintain matching metadata between the photograph description in PastPerfect as well as the digital surrogate in the Digital Archive Management software, CONTENTdm, was isolating. Figure 1 illustrates how PastPerfect functioned as an island of content because of its inability to pass description metadata to any other system without manual intervention. This resulted in constant data upkeep between the two systems that was fundamentally manual and repetitive.

### Discussion

Using the information gleaned from the system mapping, the library working group was able to explain to library decision-makers the toil issues exposed, especially in terms of XTF and PastPerfect. Adopting Google's toil definition as a framework, it became obvious the ways in which the high rate of toil tied to these systems was limiting the work done or needing to be done. For example, because of the attempt to maintain a one-to-one metadata relationship between PastPerfect and CONTENTdm, metadata clean-up efforts were stalled or shut down

because of the amount of manual work involved. This resulted in collections where data was stale and sometimes outdated. Similarly, one staff member was in charge of pushing updated EAD findings aids to XTF and often was overwhelmed with keeping track of what changes needed to be made and where. High toil processes were taking their toll in being able to proactively make improvements to data and overworking employees with reactive tasks.

Knowing that the Reid Collection was about to upset these processes even more by forcing large amounts of data through these pathways, University of Nevada, Reno's Library Administration supported efforts to reduce the toil by adopting new systems and sunsetting others. Although the Reid Collection remains closed to the public, the processing team working on the records handled a significant amount of data during the first year of arranging and describing the collections' physical and digital components. One subseries alone included records for over 63,000 photographs. Recognizing the tension created by this influx of data, the working group began to explore options for replacing systems that would not depend on toil-heavy pathways. As a result of the group's systems appraisal, the Libraries eventually undertook four concurrent system changes (Figure 3).

In particular, collections with both photograph and manuscript pieces benefited heavily from the introduction of new systems and sunsetting of overly complex systems. Previously, this data pathway held the most toil. Manuscript pieces of the collection would be accessioned in Archivists' Toolkit with a separation note for the photographs while the photograph pieces were accessioned in PastPerfect. The accession records in both systems would lead to description of the collection in two separate systems and then the manual push of that metadata through two separate publication pathways (XTF and PastPerfect Online) by a staff member. This disparate data model was especially toil laden if any of the collection was digitized and therefore had digital surrogates and description of those surrogates stored in yet another system (CONTENTdm). From the system map, it was obvious to note that data was hitting manual pathways for each system combined with the fact that these systems were heavily removed from one another even though they were attempting to describe and facilitate access to the same materials.

In order to remove the manual pathways, Library Administration approved sunsetting of PastPerfect and its online component as well as the discontinuation of XTF in favor of ArchivesSpace and its native public interface. Removal of these barriers meant that description of collections could be stored in one "system of record" instead of two, thus negating the need for a one-to-one relationship between digital surrogate description and physical item description, since it was understood that the system of record held the original and most reliable data. Additionally, ArchivesSpace allows for linking between digital objects and the associated finding aids, which enables the user to quickly move between the

### 2018 System Changes




Figure 3.  
University of Nevada,  
Reno Libraries' system  
changes in 2018

---

two systems without having to rely on manual interventions by library staff members. By eliminating PastPerfect and XTF in favor of more automation friendly systems, the Libraries were able to reduce the overall system toil rate from 63% to 41%.

### *Limitations*

Since the purpose of this exercise was focused on library systems and data pathways, no quantitative data was gathered from individuals about their processes in terms of toil. Though the exercise yielded many quantitative comments from library staff members about the draining nature of their day-to-day work within the mapped systems, toil discussions were largely focused on the toil output of specific systems onto employees rather than the amount of toil work undergone by employees working in those systems. Thus, metrics were calculated using the net value of data pathways between systems and not measuring the individual toil attributes per system. This approach allowed administrators who had no experience or stake in the systems to visualize the need for system changes instead of simply workflow changes.

### *Future directions*

One of the systems identified as the biggest toil contributor during the system mapping was Sierra, the integrated library system (ILS). Due to the comprehensive and costly expense of attempting to migrate this system, it was ruled out as a reasonable change to make quickly. However, after the mapping was presented to Library administrators, a push was made to investigate how newer systems might better meet the Library's needs in regards to data exchange. The University of Nevada, Reno Libraries submitted a formal request for proposal to replace the current ILS in March of 2019 with the hope that more manual data pathways can be remediated with a new system.

### **Conclusion**

This paper has worked to expand use of the concept of toil from the field of information technology to that of academic libraries. The importance of toil in the field of information technology emerged because it gave SREs a clear mechanism to identify and address tasks that otherwise placed a negative burden on their workload. This paper showed a close alignment between the application of the concept by engineers and its use by library decision-makers to analyze systems at the University of Nevada, Reno Libraries. It demonstrated the utility of the toil concept as part of a larger effort to map library systems and derive a quantitative measurement of the costs and benefits associated with platforms and their given pathways. In this way, toil provides an invaluable means for analyzing the positive or negative impact of individual systems.

The strength of toil, as a concept, derives from its ability to guide decisions when refining or building new infrastructure. Toil, with its six criteria, acts as a concrete and flexible framework that can be applied to a variety of systems. In fact, it offers a means to evaluate more than just software, but also other large-scale technical processes in academic libraries. For example, it could be used to compare the value and sustainability of different descriptive processes. It also could be applied more directly to the impact of systems on librarians and archivists. In this sense, toil could be used to ask fundamental questions about the negative or positive effects of a particular process on the individuals tasked with carrying it out.

The extension of this model to academic libraries also shows the further blurring of lines between the concepts used to support information technologies and those found in libraries and archives. While information technology concepts do not map perfectly onto the needs of

library administrators, they do provide a starting point for evaluating how libraries administer their increasingly technical portfolio of tools. Similarly, using the concept of toil to eliminate unsustainable and highly manual task pathways in a library setting has shown the value and promise of the concept for library administration.

Technology turnover in libraries has often led to hasty adoption of single solution systems without thought about the library's overall system ecosystem. Creating a system map and visually plotting manual and automated data pathways can reveal to library staff members and administration the overall toil levels being experienced. Thinking about systems in terms of their toil can also expose those that are no longer meeting the needs of the library or its users and can encourage investigation of newer or more robust systems that can be adopted instead. By thinking about systems in an information technology sense and adopting applicable vocabulary, libraries and their leaders can begin to endorse a comprehensive approach to library systems and bolster library staff members who are weighed down with manual system processes.

### References

- Buljung, B. and Johnson, C. (2017), "Up against the clock: migrating to LibGuides v2 on a tight timeline", *Information Technology and Libraries*, Vol. 36 No. 2, pp. 68-86.
- Catchpoint (2018), "2018 SRE report", [online], available at: <https://pages.catchpoint.com/2018-SRE-Report> (accessed 25 July 2019).
- Catchpoint (2019), "Catchpoint: site reliability engineers say incident resolution, related stress and excessive manual tasks key concerns", *Manufacturing Close - up*, [online], available at: <http://unr.idm.oclc.org/login?url=https://search.proquest.com/docview/2203604787?accountid=452> (accessed 25 July 2019).
- Clair, K. (2016), "Technical debt as an indicator of library metadata quality", *D-lib Magazine*, Vol. 22, [online], available at: <http://www.dlib.org/dlib/november16/clair/11clair.html> (accessed 25 July 2019).
- Critchlow, M., Friedman, L. and Suchy, D. (2010), "Using an agile-based approach to develop a library mobile website", *Code4Lib Journal*, No. 12, [online], available at: <https://journal.code4lib.org/articles/4642> (accessed 25 July 2019).
- Lindström, H. and Malmsten, M. (2008), "User-centred design and agile development: rebuilding the Swedish national union catalogue", *Code4Lib Journal*, No. 5, [online], available at: <https://journal.code4lib.org/articles/561> (accessed 25 July 2019).
- Mayo, D. and Bowers, K. (2017), "The devil's shoehorn: a case study of EAD to ArchivesSpace migration at a large university", *Code4Lib Journal*, No. 35 [online], available at: <https://journal.code4lib.org/articles/12239> (accessed 25 July 2019).
- Neatrou, A., Myntti, J., Brunsvik, M., Maringanti, H., McBride, B. and Witkowski, A. (2017), "A clean sweep: the tools and processes of a successful metadata migration", *Journal of Web Librarianship*, Vol. 11, pp. 3-4, [online], doi: [10.1080/19322909.2017.1360167](https://doi.org/10.1080/19322909.2017.1360167) (accessed 25 July 2019).
- Otto, B. (2011), "Data governance", in *Business and Information Systems Engineering*, Vol. 3 No. 4, pp. 241-244, [online], doi: [10.1007/s12599-011-0162-8](https://doi.org/10.1007/s12599-011-0162-8) (accessed 25 July 2019).
- Otto, B. (2011), "Organizing data governance: findings from the telecommunications industry and consequences for large service providers", *Communications of the Association for Information Systems*, Vol. 29, [online], doi: [10.17705/1CAIS.02903](https://doi.org/10.17705/1CAIS.02903) (accessed 25 July 2019).
- Rau, V. (2016), "Eliminating toil", *Site Reliability Engineering: How Google Runs Production Systems*, O'Reilly Media, available at: <https://landing.google.com/sre/sre-book/chapters/eliminating-toil/> (accessed 14 June 2019).
- Tuyl, S.V., Gum, J., Mellinger, M., Ramirez, G.L., Straley, B., Wick, R. and Zhang, H. (2018), "Are we still working on this? A meta-retrospective of a digital repository migration in the form of a



---

classic Greek Tragedy (in extreme violation of Aristotelian Unity of Time)", *Code4Lib Journal*, No. 41, [online], available at: <https://journal.code4lib.org/articles/13581> (accessed 25 July 2019).  
VoltDB (2017), "A single source of truth: is it the mythical creature of data management?", [online], available at: [https://www.voltdb.com/wp-content/uploads/2017/08/Voltdb\\_SSOT\\_Whitepaper.pdf](https://www.voltdb.com/wp-content/uploads/2017/08/Voltdb_SSOT_Whitepaper.pdf) (accessed 20 September 2019).

### About the authors

Emily E. Boss is the Head of Metadata and Cataloging at the University of Nevada, Reno Libraries, where she oversees a variety of projects related to metadata and metadata infrastructures. Previously she worked in cataloging and electronic resources evaluating data pathways and migrations. Her research and expertise focuses on the intersection of technology and libraries as well as the changing nature of traditional cataloging. Recently she has been an active proponent of applying ethical principles to authority records as part of the Cataloging Ethics Steering Committee Authority Working Group. Emily E. Boss is the corresponding author and can be contacted at: [eboss@unr.edu](mailto:eboss@unr.edu)

Nathan Gerth is the Head of Digital Services and Digital Archivist at the University of Nevada, Reno Libraries. Splitting his time between application support and digital preservation, he is the primary custodian of the electronic records from the papers of Senator Harry Reid. Outside of UNR he has been an active participant in the congressional papers community, serving as the chair of the Congressional Papers Section and as a member of the Association of Centers for the Study of Congress CSS Data Task Force.

---

For instructions on how to order reprints of this article, please visit our website:

[www.emeraldgroupublishing.com/licensing/reprints.htm](http://www.emeraldgroupublishing.com/licensing/reprints.htm)

Or contact us for further details: [permissions@emeraldinsight.com](mailto:permissions@emeraldinsight.com)

Reproduced with permission of copyright owner. Further reproduction prohibited without permission.