Improved Corporate Search Engine for the National (Australian) Spatial Information Management System: Case Study

Richard Crawford
BAE Systems
richard.crawford2@baesystems.com

Dr Sergey Nesterov
University of South Australia
sergey.nesterov@unisa.edu.au

Abstract

Paper presents case study of the evolution of the National Spatial Information Management System (NSIMS). It also describes trade-off analysis approach used to compare different Search Engines for NSIMS. Paper addresses design considerations of the Improved Corporate Search Engine (ICSE) recently developed in Australia. The architecture of the NSIMS has been described, as well as existing search engine capabilities and limitations. Two key areas for search engine improvements have been identified: (1) search context information usage for data relevance criteria application in a search algorithm; (2) user interfaces for entering search requests and query results presentations. Paper presents technical details of ICSE implementation and discussion of characteristics of the proposed improvements in the identified areas. Practical evaluation technique of the search engine efficiency has been proposed and results of trade-off analysis discussed.

Paper sections:

1. National Spatial Information System Overview
2. Existing Search Engine Capabilities and NSIMS Limitations
3. The Research Problem and Solution Approach
4. System and Data Analysis
5. Trade-off Study between Existing and Proposed Search Engine
6. Conclusion

National Spatial Information System Overview

The National Spatial Information System (NSIMS) is a system developed by BAE Systems for the Defence Support Group (DSG). DSG is responsible for the provision of the working, training and living environment required by the Australian Defence Organisation (ADO). The purpose of NSIMS is to manage the use of spatial data across DSG staff and provide ready access to information about spatial data and the ability to view, visually analyse, download and re-submit a spatial dataset that has been edited. A key component of this is the ability to quickly and easily find spatial datasets that have been loaded into the system.

NSIMS includes the following software and data components:

- A browser based spatial data viewer that enables viewing, querying and 'marking up' of Geographic Information System (GIS) and Computer Aided Design (CAD) generated spatial data;
- A spatial data metadata for searching and locating spatial and non-spatial data;
- A spatial data repository for storing spatial data;
- A gazetteer for defining and performing searches on geographic extents of a named feature

NSIMS is a web-based intranet application installed on the Defence Restricted Network (DRN) written in ASP.NET framework with an Oracle database. All datasets loaded into the system contain one or more actual spatial files and a metadata file, which is an xml based text file describing spatial data. The metadata files are stored using Oracle XML DB technology. The system is expected to store approximately 1,000,000 datasets in future, and the current number is approximately 220,000.

A high level diagram representing the NSIMS architecture, supported distributed application, is shown in Figure 1. An NSIMS instance is deployed at each of the main communication nodes on the DRN. Each of the NSIMS instances have similar configuration: a Web Server, MapGuide Server (note: these can be combined on the one physical computer if necessary), Database Server and local File Storage.

The Web Server runs the main web application code, including the search logic modules, and is the only server that external users interface with. The MapGuide Server is responsible for generating the map images that are transferred to the users' web browser. The Database Server hosts an Oracle database, which persists all of the applications information, including the metadata for the datasets. The databases are configured using master to master replication architecture. The local File Storage stores all of the spatial files loaded into the system.

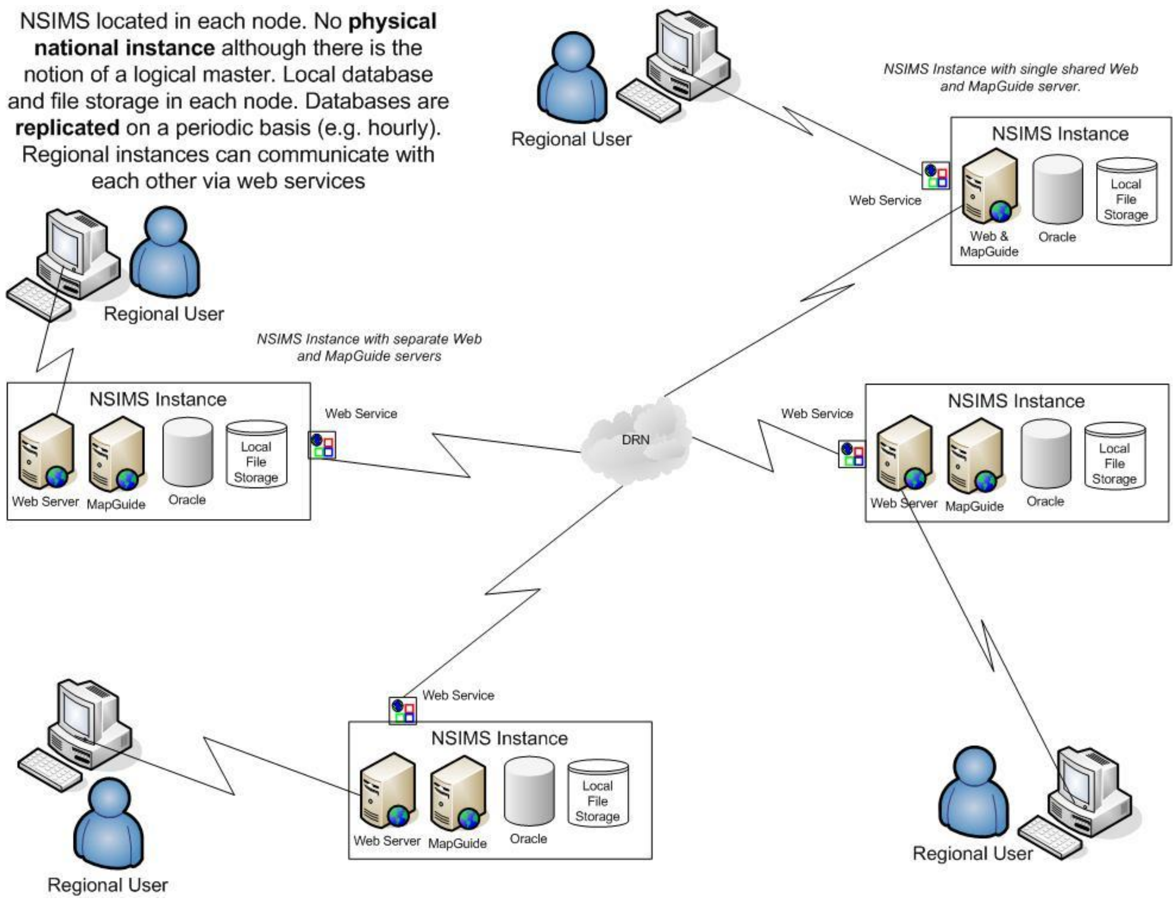

Figure 1 - NSIMS Architectural Diagram

Spatial Data Characteristics

Spatial data is information that identifies the geographic location of natural or constructed features and boundaries on the earth including GIS, CAD and Imagery data. DSG Spatial Information may be divided into two categories geo-referenced information and geo-coded information.

Geo-referenced Information is any information that is aligned to a known coordinate system so it can be viewed, queried, and analysed with other geographic data. Examples of geo-referenced spatial data used by DSG include:

- Master site plans for Defence Properties.
- Aerial photography and Satellite imagery.
- Co-ordinate locations e.g. fauna sightings, contamination locations.
- Environmental mapping e.g. vegetation communities, fire management areas.

Geo-coded Information is any information that can be assigned to a geographic feature. Examples of geo-coded spatial data include:

- Drawing files and plans that are not geo-referenced e.g. floor plans, as constructed, works as executed, underground and above ground services.
- Word and PDF documents, excel spreadsheets e.g. reports and manuals.

- Electronic images of scanned paper plans.
- Hand-held imagery e.g. photos of buildings.

Metadata is 'data about data', usually in a text format. NSIMS allows searching of the metadata files associated with its spatial information, just as popular WEB search engines Google or Lookle search the metadata, associated with web pages. However, the spatial metadata loaded into NSIMS must conform to the DSG Metadata Profile which is derived from the Australian Defence Organisation Metadata Application Profile, which is inturn derived from the ANZLIC 3 metadata profile of ISO 19115. The DSG Metadata Profile contains approximately 200 different fields many of which are allowed to have multiple occurrences.

Existing Search Engine Capabilities and NSIMS Limitations

The current NSIMS search engine is essentially a data retrieval implementation as opposed to an information retrieval system. A data retrieval system requires an exact match between the search criteria and the data fields for the record to be returned. This is very similar paradigm to an SQL query on a database. An information retrieval system uses a less deterministic and more intelligent approach and may incorporate data from other sources, synonyms or allow partial matches.

The difference between an information retrieval and data retrieval system is briefly discussed in (Larson, n.d.), with the following definition of Geographic Information Retrieval:
"Geographic Information Retrieval, as we define it here, is an applied research area that combines aspects of DBMS research, User Interface Research, GIS research, and Information Retrieval research, and is concerned with indexing, searching, retrieving and browsing of geo-referenced information sources, and the design of systems to accomplish these tasks effectively and efficiently..."(Larson, n.d.)

In the context of NSIMS the existing search engine characteristics closely align with the attributes of a data retrieval system, as illustrated in Figure 2. NSIMS uses a deterministic data retrieval model whereby if a dataset satisfies the search criteria entered by the user then the dataset is considered relevant and subsequently returned. Conversely if the dataset fails to meet any of the search criteria then it's considered not relevant and is not returned in the search results. The search results are presented to the user in an arbitrary order with no concept of relevance used.

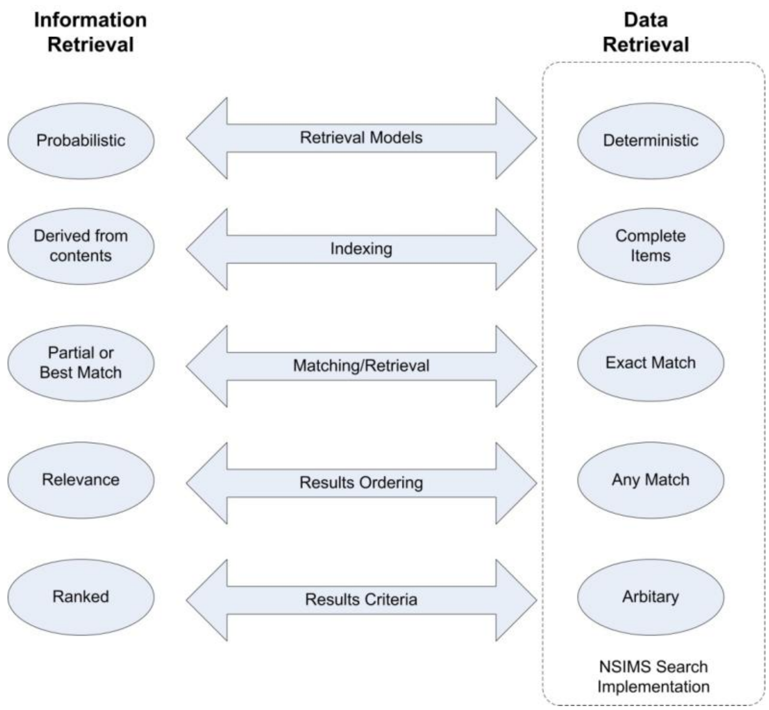

Figure 2 – Information vs. Data Retrieval Systems

The NSIMS has a comprehensive queries collection mechanism to gather all information of the user data requests. This logging information is used for security reports and analysis to guarantee only authorised access to the spatial data. However, this information can be used to analyse users preferences and patterns in accessing relevant data sets.

Limitations of the Metadata

Each dataset contains a number of files and a DSG Metadata Profile compliant metadata record that describes the files. The vast majority of the spatial data to be loaded into NSIMS had no existing metadata. To analyse feasible solutions and create metadata records for all of these files a special data collection project was undertaken. The tight timeframes and the sheer quantity of data meant that the only practical method of creating the metadata was to use an automatic process.

The DSG Metadata Profile potentially offers opportunity for a very detailed description of the data. The reality is that the majority of the original metadata has been created using simple automatic methods based on only a small set of information, namely the name of the file and the directory it's located in. As a result the original metadata contains only a minimal set of useful information fields with a large number of fields just having default values that will be the same for large groups of datasets. For example, the abstract is identical for all floor plans.

Limited value of original metadata restricts the number of useful fields that can be searched because if we searched on the abstract field the results will either return all or none of the floor plan datasets. The primary purpose of a building isn't captured in the metadata but relies on an obscure primary key link from an external system. For example, a search for "RAAF Tindal Base Medical Wards" will not return any datasets.

Typically the only information that is available to populate the original metadata is the Defence Asset (e.g. property, structure) it relates to, the name of the file and the location of the file within a directory in the File Storage.

This information is represented in the metadata in the following fields:

- Citation Title (name of the dataset) – derived from the file name
- File type – from the file extension
- Spatial location – extracted from the file for geo-referenced files
- Theme keywords – from the file type and directory location

User Interface

A screenshot showing the default NSIMS search interface is shown in Figure 3. Each one of the links on the left hand side shows the user additional fields that they can search on. As can been seen from this screenshot the user is presented with a significant number of search options, particularly when compared to popular search engines such as Google.

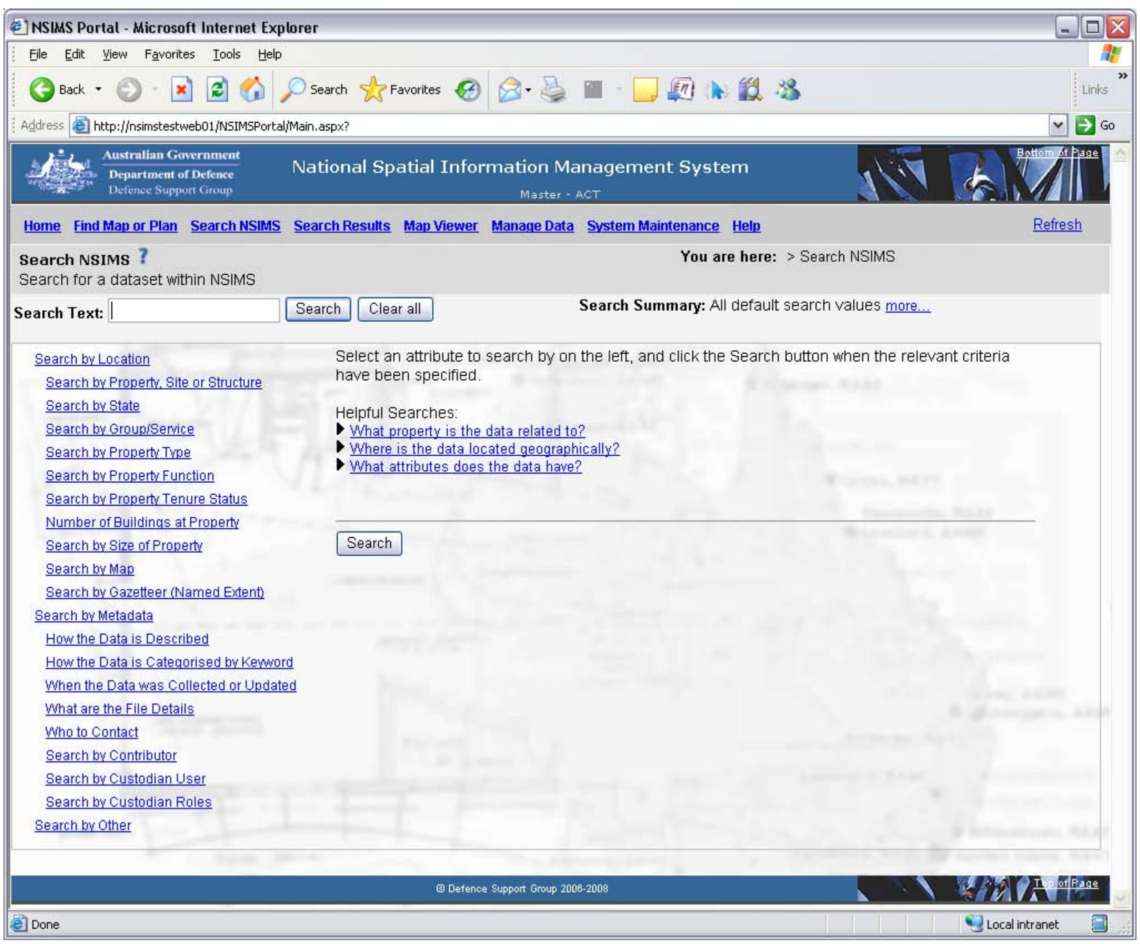

Figure 3 - Existing NSIMS search engine user interface to make a query

A screenshot showing sample results of a search is shown in Figure 4.

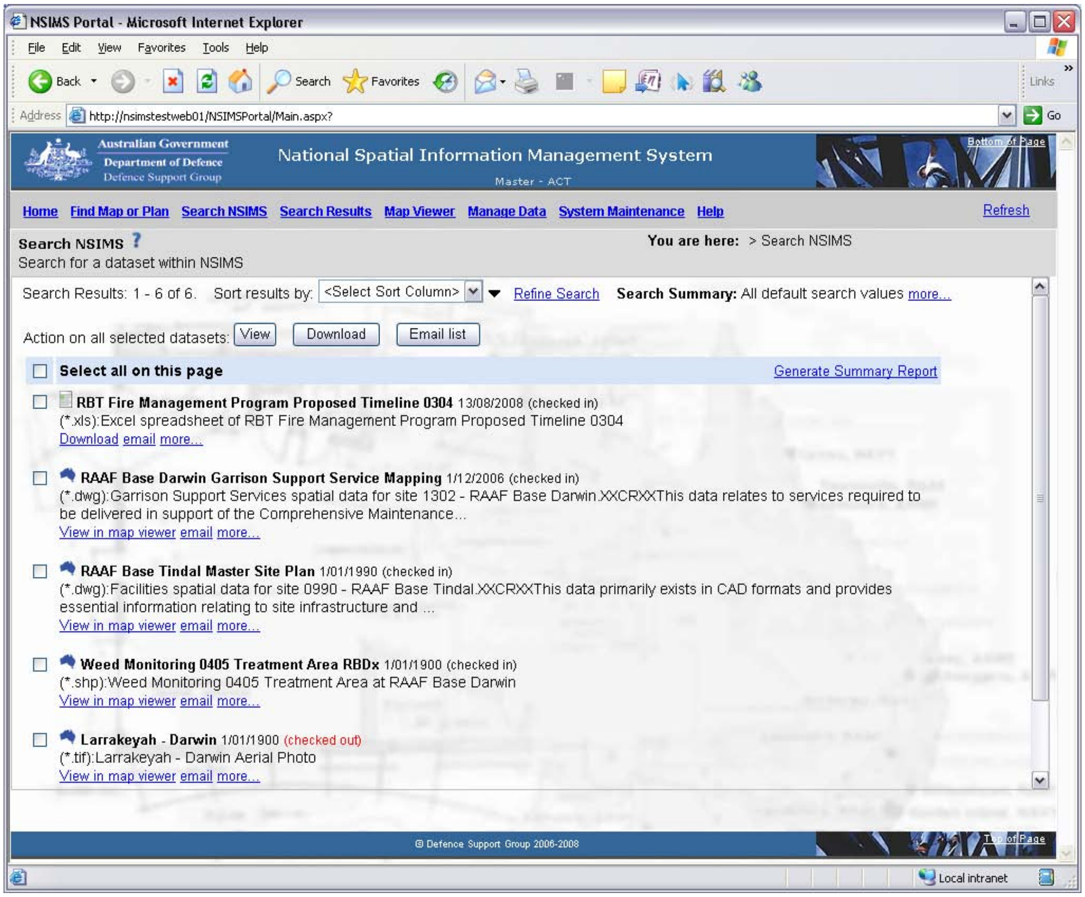

Figure 4 - Example of search results presentation in the existing NSIMS Search Engine

Conclusion

The key issues identified with the existing NSIMS implementation are limited metadata value, complicated, as a result, search interface and a data matching search engine. The combination of these issues makes it difficult for the users to find relevant datasets, which could then have the following negative consequences:

- The most relevant datasets not being used to assist the decision making process;
- Duplicate and therefore redundant datasets being purchased because the original ones weren't found;
- Degradation of system performance due to users having to do many searches to find the datasets that they are interested in;
- Manual and expensive re-populating of the metadata records to provide more detailed information to perform subsequent searches on;
- NSIMS not being used as intended and instead local experts being tasked to retrieve the relevant datasets;
- Users becoming disillusioned about using the system because they cannot find the datasets that they are after; and
- Lost DSG staff productivity.

The Problem Statement and Solution Approach

Statement of the Problem

To overcome limitations of the existing search engine we have asked the question:
How to design and implement the more efficient intranet search engine for NSIMS?

Hypotheses

To provide a solution we stated key hypothesis and assumptions which form the solution approach framework.

1. NSIMS is a corporate intranet application that is a special case and needs to be handled differently from a web search. This means that many of the techniques that work well in a web search will not be directly applicable to NSIMS and conversely that NSIMS offers the opportunity to exploit information which is not possible for a web search. A summary of the list of differences is contained in Table 1.
2. The datasets loaded into NSIMS can be characterised by a limited set of document classes.
3. The NSIMS user community can be represented by limited number of groups and each group has specific document usage patterns.
4. The concept of document relevancy to user request and system usage indicators can be used as a measure of search engine efficiency.
5. Existing system log data can be used for search engine efficiency valuation. Although it may be necessary to identify additional data to be collected for the ongoing improvement of the search engine.
6. A new search engine prototype implementation will be able to provide evidence of a more efficient solution through data collection and interviews.
7. The limitations in the metadata will result in a large percentage of searches failing to return relevant datasets.
8. The search interface is overly complex and presents the user with too many options to search on. The reality is that there are probably only a handful of criteria that will form the vast majority of all searches.

Characteristics	NSIMS	Web Search Engine	Notes
Users	Limited number of professionals	Vast variety	Different
Information system purpose	Disseminate Information	Attracting and holding attention	Different
Content & System management	Limited number of organisations	Variety of organisations and individuals	Different
Search result success criteria	The right answer: unique document or set of unique documents	Most relevant documents	Different
User Interface	Web browser layout	Web browser layout	Similar
Search Criteria	Textual patterns, Geo positions	Textual patterns	Similar and Different
Search Relevance	Geo closeness,	Most popular pages,	Different

Characteristics	NSIMS	Web Search Engine	Notes
Criteria	Corporate specific	Generic	
Response time	Quick	Best effort	Different
Information type	Spatial	Variety	Different
Information storage	Corporate distributed database in closed network	WEB servers world wide	Different
Search space	Metadata	WEB pages and associated metadata	Different
Index file volume	Hundreds of Thousands of Records	Hundreds of Millions of Records	Different
Data access policy	Strict	N/A	Different

Table 1 - Characteristics of Web and Corporate Spatial Information System

Sub-problems

To answer the key research question we addressed the following sub-problems to implement the solution approach:

1. System and data analysis to identify user groups, document classes and usage patterns.
2. Identification of the common search criteria specified by the users
3. Concept of document relevancy to user request and indicators. Analysis of the corporate environment to identify influential factors on search engine.
4. Search engine improved algorithm to count relevancy in request processing and in results presentation.
5. Evaluation of the new search engine efficiency based on preliminary results.

System and Data Analysis

Data Collection

The data for this system usage or user behaviour research was collected after NSIMS had been in operational use for about 10 weeks. The majority of the information came directly from NSIMS with the exception being the work location and the job role of all users that had accessed NSIMS, which were retrieved from their contact details on the DRN. The information collected from NSIMS included the following:

- Search logs – these detailed all of the searches that had been performed on the system, what search criteria was specified and the username of the user performing the search;
- User activity details – these specified all datasets that users had accessed (e.g. viewed or downloaded);
- Dataset Metadata – provide a detailed description of each dataset;
- Mapping of datasets to Defence Properties – includes all datasets that expert users have identified as being the most important for a Defence Property and are most likely to meet the needs of the majority of users' requests;

- NSIMS Gazetteer – provides relationships between location names and geographic positions;
- DEMS Common Name – the common name used to describe a Defence Asset.

Data Analysis

The data analysis process involved data normalisation and analysis the data itself. The two major areas were user searches and dataset access history. The user searches were analysed to try and determine if the search history could be used as an input into updating the search user interface. The dataset access history was examined to discover the most popular types of datasets accessed and whether the job role of the user had an impact on the types of datasets that they accessed.

Search Logs

The search logs were examined to determine the most popular criteria that users searched on. Combining the most commonly used search criteria with the perceived intent of the text searches shows that the user typically framed their search either one or a combination of the following criteria:

- DEMS identifier or location name
- Name/Citation title of the dataset
- DSG Region
- Keywords
- Datasets that the user performing the search contributed or is the custodian of

The searches were compared with the user access history to determine whether a search resulted in a dataset being accessed. This was done in an attempt to uncover any patterns between searches that resulted in datasets being accessed compared with those that didn't. This highlighted that approximately 50% of text searches that resulted in no datasets being accessed contained words from the DEMS common name, which isn't contained in the metadata. This seems to support the hypothesis that the limitations of the metadata will impact on the existing search engines ability to return the relevant datasets for a search. This indicates that the DEMS common name, even though it's not contained in the metadata, should be included as part of the search. The approach of using context information from databases to improve intranet search engines is discussed by Mangold, Schwarz and Mitschang (2005). A method that could be used to apply this to NSIMS with the least impact to the existing search engine implementation is to automatically include the DEMS common name as a type of meta-metadata. This has similarities to an approach proposed in „Using Annotations in Enterprise Search' (Dmitriev, Fontoura, Eiron, & Shekita 2006) where annotations are added to web pages to assist searches. We would follow a similar approach except that the annotations are automatically retrieved from DEMS and these are related to a dataset instead of a web page.

Dataset Relevance Interpretation

The dataset access history was analysed to determine whether there were any characteristics of the datasets accessed by user groups that could be used as an input into calculating relevance. The steps followed were:

1. All NSIMS users were grouped based on their job category, which was determined from their contact details on the DRN. The percentage of datasets accessed for each group of users was calculated from the user activity details.
2. The datasets metadata was used to classify all datasets loaded into the system into a discrete set of classes. The classes were based on the file type, theme keywords and whether the dataset was geo-referenced or not. Note that the enforced metadata structure of the datasets stored in NSIMS enables us to group them into a relatively small number of classes, which wouldn't be feasible in a generic search environment.
3. The usage pattern for all of the NSIMS user groups for each of the datasets classes was calculated.

The results of this showed that different categories of dataset were accessed more frequently than others and that this access history varied amongst the different user groups. The concept of analysing the usage patterns of social networks and exploiting this information to improve the search experience for users belonging to certain social groups is discussed by Mukherjee and Verity (2004). It also showed that some categories of datasets were rarely accessed and yet accounted for a large percentage of the datasets loaded into the NSIMS and vice-versa where some categories of datasets were accessed frequently and yet only accounted for a small percentage of datasets loaded into NSIMS. This lead to the following approach of calculating a relevance weighting based on the datasets class and frequency of access:

1. There are k document groups or classes (1..k) . Each class j has Nj documents in it, and N= SUM(of Nj) – number of documents in the system.
2. The dataset access (viewed or downloaded) is calculated per the particular class. So, M=SUM(of Aj) – total number of accesses to the documents and Aj – number of accesses to the particular class.
3. The share of access per class Sj is calculated as Sj=Aj/M
4. The document interest Ij based on a particular class if we assume a random nature of dataset access is calculated as Ij=Sj/Nj – Interest per document belonging to the class j.
   Note: the document interest is calculated for all users as well as each of the user groups

The document interest measure is used as a surrogate measure of document relevance in prioritizing document order in responses.

Distance from User to Location in Dataset Accessed

It was expected that most users would be more interested and therefore likely to access datasets that were at or near their current work location. The user activity reports were analysed to determine the distance between the centre point of the users work location and dataset accessed location. The results of this are shown in Figure 5, which indicates that there is a steep increase in the percentage of datasets accessed within 40km of the users' location and after this the line is reasonably linear, which indicates distance isn't a factor for datasets greater than 40km away (e.g. a user is just as likely to want to view a dataset that is 80km or 3000km away).

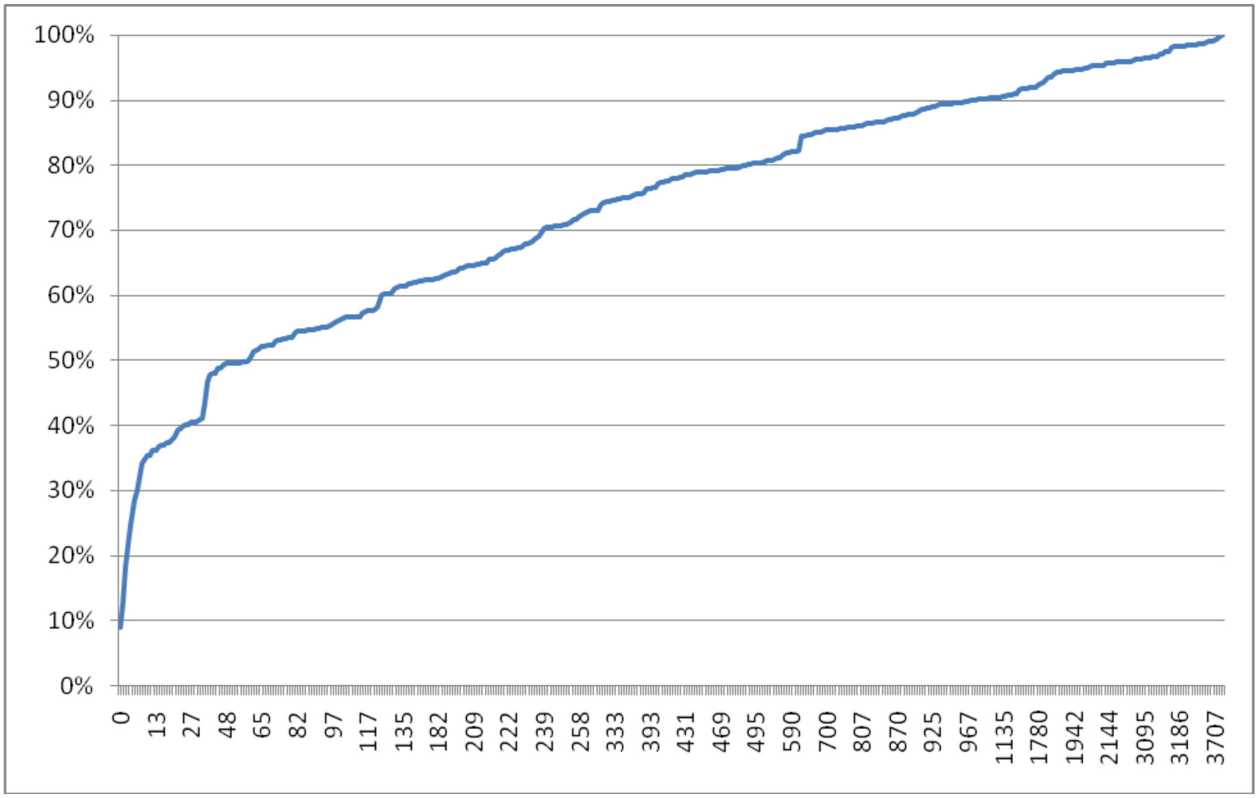

Figure 5 – Distance from User to Dataset Accessed

Summary on search criteria

The analysis of the search logs and user access history provides a key input into the design changes for improving the NSIMS search engine. The search logs indicate that the combination of five criteria make up the vast majority of searches performed on the system and that the users frequently use terms not contained within the metadata but that are contained in an external system (DEMS). The user access history shows that the class of dataset and job role of the user performing the search can be used for calculating the document interest which could then be used as a surrogate measure for datasets relevance. The distance between the user and dataset accessed shows that there is a higher percentage of datasets accessed within 40km of the users work location.

Trade-off Study between Existing and Proposed Search Engine

To enable the evaluation of the proposed search engine changes a trade-off study was conducted using the existing and proposed search engine. This required the following steps:

1. Determining the criteria to evaluate the changes impact and effects;
2. Building a prototype solution incorporating the proposed changes;
3. Developing a set of user scenarios to exercise the searches;
4. Setting up a test environment which contained two NSIMS instances in a replicated environment. One of the NSIMS instances had the original search engine and the other the prototype;
5. Run search scenarios on both the NSIMS instances and record the results; and
6. Evaluate the results.

Evaluation Criteria

We suggest that search engine performance has to be considered as a composite index of key indicators, addressing different characteristics of the search process, and quality of retrieved results (information retrieval process).

We propose to assess search engines basic technical performance characteristics; two widely accepted standard search quality characteristics: precision and recall; and user perception characteristics, listed below:

- Technical characteristics:
  - Response time - consider user time to enter the query + the time for system to perform a query,
  - Number of searches required to find the relevant dataset,
  - Number of relevant documents on the first search page, and
  - Relevance of the first dataset returned.
- Standard Search metrics, quality (of relevance filtering criteria) search characteristics:
  - Precision – ratio of relevant documents returned to the total number of returned documents, and
  - Recall – ratio of relevant documents returned to the total number of relevant documents in the database.
- User related characteristic:
  - Implicit (average number of "clicks" (or requests) to find a document, and
  - Explicit –User "satisfaction index", data obtained from questionnaires and interviews.

Search Engine Prototype

A prototype solution based on the proposed ICSE design was built to enable evaluation of this design. The prototype solution included the changes to the search engine to display the datasets in order of relevance and include additional details from DEMS when performing the search. The setting of all of the dataset relevance weights was done by running a database script and included none of the feedback mechanism to dynamically update these weights in response to the system usage.

A search user interface developed as part of the prototype is shown in Figure 6 This layout is based on a number of popular search engines (e.g. Google, Live Search, Dogpile) and is designed to include the main search criteria identified in the search logs.

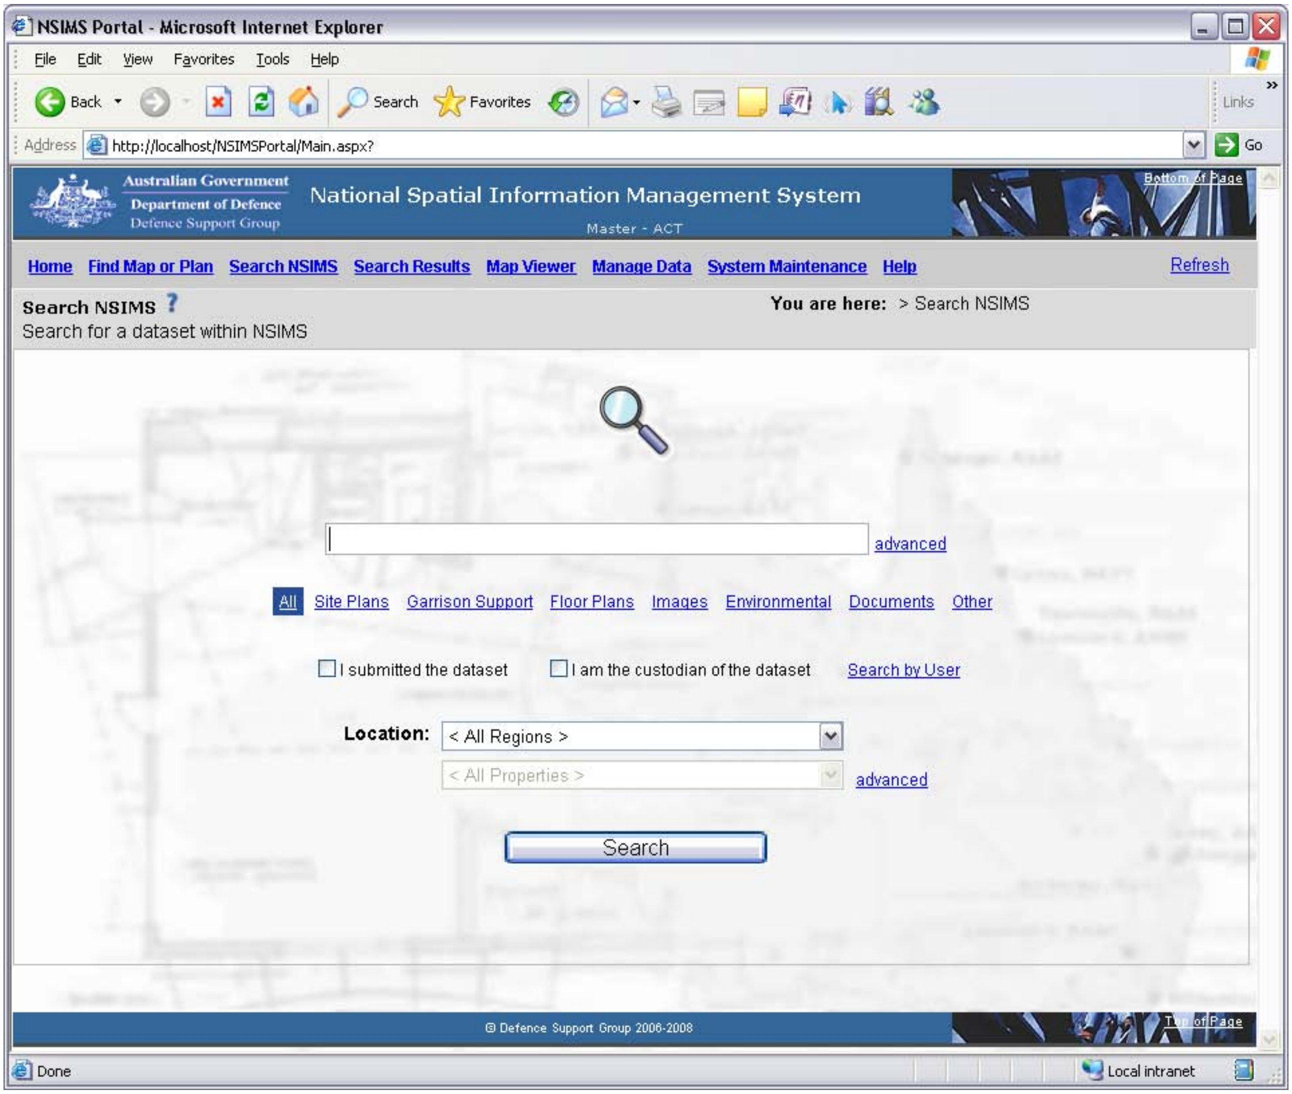

Figure 6 – Proposed Search Interface

Trade-off Study

Search Engine Evaluation Questionnaire

To enable evaluation of the prototype ICSE compared to the original search implementation a questionnaire was developed that required a group of users to conduct some common searches on both systems. The questionnaire recorded the following:

- Steps followed to perform the search,
- Time taken from when the user commenced the task until they identified the correct dataset(s),
- The number of search pages accessed to enter the criteria,
- The number of documents returned,
- The number of relevant documents on the first page, and
- The relevance of the first document returned.

The questionnaire was run with five different users of varying experience in using NSIMS. Two of the users had never used NSIMS before. For User #5 (who had never used NSIMS before) the search scenarios were done in the opposite order with the ICSE prototype used first.

Technical Characteristics

A summary of the results for the search scenarios are detailed in Table 2. In every category there is a significant improvement with the ICSE prototype.

Category	Original Search	Prototype Search
Average time taken	75.03 sec	30.75 sec
Average number of searches performed per scenario	2.23	1.28
Percentage of times the relevant datasets were found	85%	100%

Table 2 - Search Scenario Summary

Standard Search Metrics

The basic measures used for text retrieval are the precision and recall. The precision is a measure of the percentage of retrieved documents that are relevant to the query. The recall is the percentage of documents that were relevant to the query and were in fact retrieved. The goal is to balance these two measures, which is typically done by coming up with the harmonic mean (Han & Kamber, 2006). For example if a search in NSIMS returned all datasets in the system then the recall would be 100% as we are guaranteed that all relevant documents are retrieved. However the precision would be very low as there would be many non-relevant datasets returned.

In the context of NSIMS it is only when the user searches using a text string that precision and recall can be considered. The search scenarios had a number of questions where the users only entered text as the search criteria. These queries have been used as a basis to measure the precision and recall characteristics of the two search engines. The results are contained in Table 3.

Original Search		Prototype Search	
Precision	**Recall**	**Precision**	**Recall**
100%	57%	100%	100%

Table 3 - Precision and Recall Comparison

The characteristics of the original NSIMS search engine, which uses an SQL query on a database, result in only relevant documents being returned. Therefore the precision will always be 100% if any datasets are returned. The prototype search engine also uses an SQL query but includes the common name for the asset from DEMS as well as the metadata in the search. This potentially could result in non-relevant documents being returned although this didn't occur during the test scenarios. The full implementation of the ICSE will include gazetteer names and because a tiling system is used to improve the performance it's likely that non-relevant data could be returned from a near-neighbouring area.

Neither of the search engines will cater for spelling mistakes or synonyms in either the search text or the metadata. As more data is loaded into the system with metadata manually created it's likely that this could become an issue and will affect the precision and recall measures. Due to the way the search scenarios were conducted, where the user was given a task to perform in written form, this issue didn't eventuate. If however the user was told verbally the datasets to find this would have been more likely to occur. For example many people may have spelt „Tindal' or „sergeant' incorrectly if they didn't have something to copy from.

User Response

The implicit user related characteristics measured were the average number of mouse clicks and average number of search pages required to perform a search. The prototype search showed significant improvement in both these measures with the results detailed in Table 4.

Category	Original Search	Prototype Search
Average number of mouse clicks	8.05	3.8
Average number of search pages	3.25	1.35

Table 4 - Average User Actions to Find Dataset(s)

The summary of the useability feedback for using the original and prototype search interfaces are detailed in Figure 7. This shows that the prototype ICSE received higher ratings in all of the categories. The comments from the users also indicated that they felt the prototype ICSE was a significant improvement.

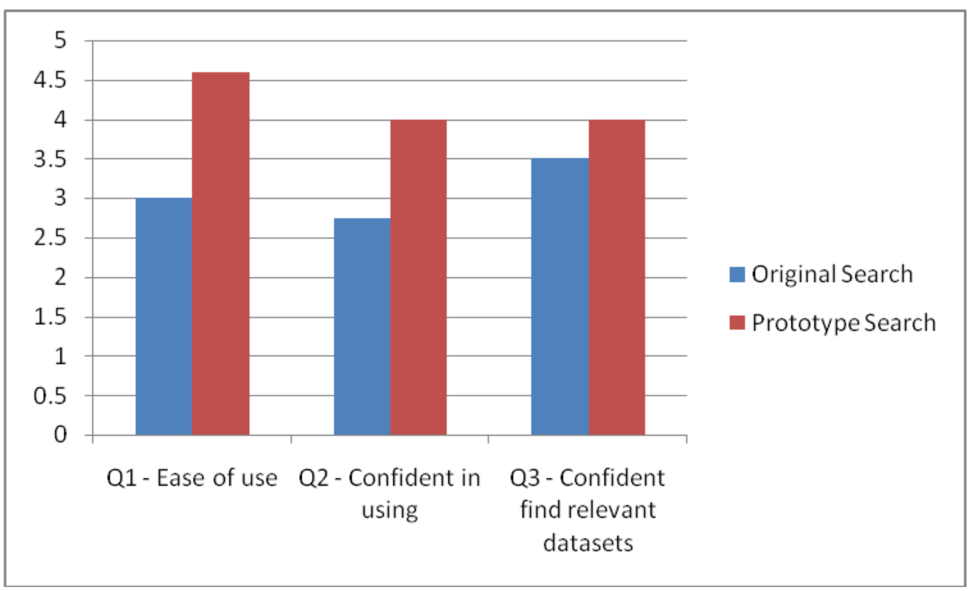

Figure 7 – Search User Feedback

Conclusion

Paper presented design considerations of the Improved Corporate Search Engine (ICSE) for the National Spatial Information Management System (NSIMS) recently developed in Australia. The architecture of the NSIMS has been described, as well as existing search engine capabilities and limitations. Two key areas for search engine improvements have been identified: (1) search context information usage for data relevance criteria application in a search algorithm; (2) user interfaces for entering search requests and query results presentations. Paper presented technical details of ICSE implementation and discussion of characteristics of the proposed improvements in the identified areas.

Paper highlighted differences between corporate search engine characteristics and search principles versus generic WEB based search engines, dealing with vast variety of users and indexing variety of WEB pages.

Practical evaluation technique of the search engine efficiency has been proposed and evaluated performing a trade study with a prototype ICSE solution. In all evaluation criteria the prototype demonstrated significant improvement.

References

Dmitriev, P. A., Fontoura, M., Eiron, N., & Shekita, E. (2006). *Using Annotations in Enterprise Search*. WWW 2006. Edinburgh: ACM.

Han, J., & Kamber, M. (2006). *Data Mining Concepts and Techniques (2nd Edition)*. San Francisco: Morgan Kaufmann Publishers.

Larson, R. (n.d.). *Geographic Information Retrieval and Spatial Browsing*. Retrieved October 25, 2008, from https://sherlock.ischool.berkley.edu/geo_ir/PART1.html

Mangold, C., Schwarz, H., & Mitschang, B. (2005). *Improving Intranet Search-Engines Using Context Information from Databases*. CIKM 05. Bremen: ACM.

Mukherjee, R., & Verity, J. M. (2004). *Enterprise Search Tough Stuff*. ACM Queue.

Biography

Richard Crawford is a senior software engineer with BAE Systems Australia. He has worked on a number of information system projects for the DoD in a number of different capacities, including system architect, system engineer and software developer. On the NSIMS project he filled the system architect and systems engineering roles and additionally was a member of the development team. His other qualifications include an Advance Diploma in Information Technology, an Associate Diploma in Cartography, a Certificate in Surveying and a Certificate in Air Camera Operation. Prior to working at BAE Systems, Richard has spent 14 years working for the Australian Defence Force.

Dr Sergey Nesterov has PhD in Computer Science & Mathematics. He is a Senior Lecturer within Defence And Systems Institute in University of South Australia. He has sound work experience in Telecommunication, IT, and Defence Industries. Research areas cover Systems Engineering and Systems Safety, Statistical Analysis, Systems Modelling and Simulation.