

# STRATEGIES FOR OPTIMIZING QUERYING THIRD PARTY RESOURCES IN SEMANTIC WEB APPLICATIONS

Albert Weichselbraun

*Institute for Information Business, Vienna University of Economics and Business Administration, Vienna, Austria*

**Keywords:** Search-test-stop, query optimization, Web services.

**Abstract:** One key property of the Semantic Web is its support for interoperability. Combining knowledge sources from different authors and locations yields refined and better results. Current Semantic Web applications only use a limited amount of particularly useful and popular information providers like Swoogle, geonames, etc. for their queries. As more and more applications facilitating Semantic Web technologies emerge, the load caused by these applications is expected to grow, requiring more efficient ways for querying external resources. This research suggests an approach for query optimization based on ideas originally proposed by McQueen for optimal stopping in business economics. Applications querying external resources are modeled as decision makers looking for optimal action/answer sets, facing search costs for acquiring information, test costs for checking the acquired information, and receiving a reward depending on the usefulness of the proposed solution. Applying these concepts to the information system domain yields strategies for optimizing queries to external services. An extensive evaluation compares these strategies to a conventional coverage based approach, based on real world response times taken from three different popular Web services.

## 1 INTRODUCTION

Semantic Web applications provide, integrate and process data from multiple data sources including third party providers. Combining information from locations and services is one of the key benefits of semantic applications.

Current approaches usually limit their queries to a number of particularly useful and popular services like for instance Swoogle, geonames, or Dbpedia. Research on automated web service discovery and matching (Gupta et al., 2007) focuses on enhanced applications, capable of identifying and interfacing relevant resources in real time. Future implementations, therefore, could theoretically issue queries spawning vast collections of different data sources, providing even more enhanced and accurate information. Obviously, such query strategies - if applied by a large enough number of clients - impose a considerable load on the affected services, even if only small pieces of information are requested. The World Wide Web Consortium's (W3C) struggle against excessive document type definition (DTD) traffic provides a recent example of the impact a large number of clients

achieves. Ted Guild pointed out <sup>1</sup> that the W3C receives up to 130 million requests per day from broken clients, fetching popular DTD's over and over again, leading to a sustained bandwidth consumption of approximately 350 Mbps.

Service provider like Google restrict the number of queries processed on a per IP/user base to prevent excessive use of their Web services. From a client's perspective overloaded Web services lead to higher response times and therefore higher cost in terms of processing times and service outages.

Grass and Zilberstein suggest applying value driven information gathering (VDIG) for considering the cost of information in query planning (Grass and Zilberstein, 2000). VDIG focuses on the query selection problem in terms of the trade off between response time and the value of the retrieved information. In contrast approaches addressing only the coverage problem put their emphasis solely on maximizing precision and recall.

Optimizing value under scarce resources is a classical problem from economics and highly related to decision theory. In this research we apply the search-

---

<sup>1</sup>[p.semanticlab.net/w3dtd](http://p.semanticlab.net/w3dtd)

Table 1: Response times of some popular Web services.

Service	Protocol	$\bar{t}_r$	$\tilde{t}_r$	$t_r^{min}$	$t_r^{max}$	$\sigma_{t_r}^2$
Amazon	REST	0.8	0.3	0.2	663.5	150.2
Dbpedia	SPARQL	0.9	0.5	0.1	301.2	42.7
Del.icio.us	REST	0.6	0.4	0.1	24.3	0.5
Geo	REST	1.8	0.1	0.0	1160.4	771.4
Google	Web	0.3	0.2	0.1	10.3	0.2
Swoogle	Web	35.8	1.6	0.2	101022.2	1762682.4
Wikipedia	Web	0.4	0.2	0.1	60.9	1.3

test-stop (STS) model to applications leveraging third party resources. The STS model considers the user's preferences between accuracy and processing time, maximizing the total utility in regard to these two measures. In contrast to the approach described in (Grass and Zilberstein, 2000) the STS model adds support for a testing step, designed to obtain more information about the accuracy of the obtained results, aiding the decision algorithm in its decision whether to acquire additional information or act based on the current answer set. Similar to (Ipeirotis et al., 2007) the resulting query strategy might lead to less accurate results than a "brute force" approach, but nevertheless optimizes the balance between accuracy and costs. This paper's results are within the field of AI research facilitating techniques from decision theory to address problems of agent decision making (Horvitz et al., 1988).

The article is organized as follows. Section 2 presents known query limits and response times of some popular Web services. Section 3 provides the theoretical background for the search-test-stop model, and presents its extension to discrete probability functions. Afterwards the application of this method to applications utilizing external resources is outlined in Section 4 and an evaluation of this technique is presented in Section 5. This paper closes with an outlook and conclusions drawn in Section 6.

## 2 PERFORMANCE AND SCALABILITY

As more and more applications facilitating external data repositories emerge, strategies for a responsible use of these resources gain in importance.

Extensive queries to external resources increases their share of the program's execution time and may lead to longer response times, requiring its operators to impose limits on the service's use.

Even commercial providers like Google or Amazon restrict the number of accesses to their services. For instance, Google's Web API only allows 1000 re-

quests a day, with exceptions for research projects. Workarounds like the use of Google's public Web interface may lead to blacklisting of the client's IP address<sup>2</sup>. Google's geo coding service imposes a limit of 15,000 queries per day and IP address. Amazon limits clients to 20 queries per second, but restrictions vary between the offered services and might change over time<sup>3</sup>. Other popular resources like geonames.org and Swoogle to our knowledge currently do not impose such limits.

A Web service timing application issuing five different queries to popular Web resources in 30 min intervals over a time period of five weeks yielded Table 1 listing the services' average response time ( $\bar{t}_r$ ), the response time's median ( $\tilde{t}_r$ ), its minimum and maximum values ( $t_r^{min}$ ,  $t_r^{max}$ ), and variance ( $\sigma_{t_r}^2$ ). These response times vary, depending on the client's Internet connectivity and location, but adequate values can be easily obtained by probing the service's response times from the client's location.

Table 1 suggests that Google provides a fast and quite reliable service ( $\sigma_{t_r}^2 = 0.2$ ) with only small variations in the response times. This result is not very surprising considering the global and highly reliable infrastructure Google employs.

Smaller information providers which cannot afford this kind of infrastructure in general provide good response times (due to fewer requests), but they are more sensitive to sudden peaks in the number of clients accessing their services. Table 1 reflects these spikes in terms of higher variances and  $t_r^{max}$  values.

Our experiments suggest (see Section 5) that especially clients querying services with high variances benefit from implementing the search-test-stop model.

Another strategy from the client's perspective is avoiding external resources at all. Many community projects like Wikipedia or geonames.org provide database dumps which might be used to install a local copy of the service. These dumps are usually rather large (a current Wikipedia dump including all

<sup>2</sup>p.semanticalab.net/gooso

<sup>3</sup>developer.amazonwebservices.com

pages, discussions, and the edit history comprises approximately 6.4 GB<sup>4</sup>) and often outdated (Wikipedia dumps are sometimes even more than one month old, other services like geonames update their records very frequently).

The import of this data requires customized tools (like mwddumper<sup>5</sup>) or hacks and rarely processes without major hassles. In some cases the provided files do not contain all available data (geonames.org for instance does not publish the `relatedTo` information) so that querying the service cannot be avoided at all.

### 3 THE SEARCH-TEST-STOP MODEL

This section outlines the basic principles of the STS model as found in decision theory. For a detailed description of the model please refer to MacQueen (MacQueen, 1964) and Hartmann (Hartmann, 1985).

MacQueen (MacQueen, 1964) describes the idea of the STS model as follows: A decision maker (a person or an agent) searches through a population of possible actions, sequentially discovering sets of actions ( $S_A$ ), paying a certain cost each time a new set of actions is revealed (the search cost  $c_{s_i}$ ). On the first encounter with a set of possible actions, the person obtains some preliminary information ( $x_0$ ) about its utility ( $u$ ), based on which he can

1. continue looking for another set of possible actions (paying search cost  $c_{s_{i+1}}$ ),
2. test the retrieved set of actions, to obtain ( $x_1$ ) - a better estimation of the actions value - paying the test cost ( $c_{t_i}$ ) and based on this extended information continue with option 1 or finish the process with option 3, or
3. accept the current set of answers (and gain the utility  $u$ ).

The challenge is combining these three options so that the total outcome is optimized by keeping the search ( $c_{s_i}$ ) and test ( $c_{t_i}$ ) costs low ( $\sum_{i=1}^m c_{s_i} + \sum_{i=1}^n c_{t_i}$ ) without jeopardizing the obtained utility  $u$ .

Introducing the transformation  $r = E(u|x_0)$  yields the following description for a policy *without testing*:

$$v = vF(v) + \int_v^{+\infty} rf(r)dr - c_s \quad (1)$$

with the solution  $v = v_0$ .  $F(r)$  represent the cumulative distribution function of the expected utility and  $f(r)$  its probability mass function. The constant  $c_s$

refers to search cost and  $v$  (better  $v_0$ ) to the utility obtained by the solution of this equation.

Extending Equation 1 to testing yields Equation 2:

$$v = vF(r_D) + \quad (2)$$

$$\int_{r_D}^{r_A} T(v, r)f(r)dr + \int_{r_A}^{+\infty} rf(r)dr - c_s \quad \text{and}$$

$$T(v, r_D) = v \quad (3)$$

$$T(v, r_A) = r_A \quad (4)$$

$T(v, r)$  refers to the utility gained by testing,  $r_D$  to the value below which the discovered action set ( $S_A$ ) will be dropped, and  $r_A$  to the minimal utility required for accepting  $S_A$ . A rational decision maker will only resort to testing, if the utility gained outweighs its costs and therefore the condition  $T(v_0, v_0) > v_0$  holds which is the case in the interval  $[r_D, r_A]$ .

In the next two sections we will (i) describe the preconditions for applying this model to a real world use case, and (ii) present a solution for discrete data.

#### 3.1 Preconditions

MacQueen (MacQueen, 1964) defines a number of preconditions required for the application of the STS model. Hartmann (Hartmann, 1985) eases some of these restrictions yielding the following set of requirements for the application of the model:

1. a common probability mass function  $h(x_0, x_1, u)$  exists.
2. The expected value of  $u$  given a known realization  $x_0$  ( $z = E(U|x_0, y_0)$ ) exists and is finite.
3.  $F(z|x_0)$  is stochastically increasing in  $x_0$ . For the concept of stochastically increasing variables please refer to (Lehmann and Romano, 2005, p75).

#### 3.2 The Discrete Search-Test-Stop Model

This research deals with discrete service response time distributions and therefore applies the discrete STS methodology. Hartmann transferred MacQueen's approach to discrete models. The following section summarizes the most important points of his work (Hartmann, 1985).

Hartmann starts with a triple  $(x_0, x_1, u)$  of *discrete* probability variables, described by a common probability function  $h(x_0, x_1, u)$ . From  $h$  Hartmann derives

1. the conditional probability function  $f(u|x_0, x_1)$  and the expected value  $Z = E(u|x_0, x_1)$ ,

<sup>4</sup>download.wikipedia.org; 2008-03-15

<sup>5</sup>www.mediawiki.org/wiki/MWDDumper

1. Austria/Carinthia/Spittal/Heiligenblut/Grossglockner (mountain)
2. Austria/Carinthia/Spittal/Heiligenblut (village)
3. Austria/Carinthia/Spittal (district)
4. Austria/National Park Hohe Tauern (national park)
5. Austria/Carinthia (state)
6. Austria/Salzburg (Neighbor) (state)
7. Austria/Tyrol (Neighbor) (state)
8. Austria (country)

Figure 1: Ranking of possible “correct” results for geo-tagging an article covering the “Grossglockner”.

2. the probability function of  $r$ ,  $f(r|x_0)$  and  $F(r|x_0)$ ,
3. the probability of  $x_0$ ,  $f(x_0)$  and  $F(x_0)$ .

Provided that the conditions described in Section 3.1 are fulfilled only five possible optimal policies are possible - (i) always test, (ii) never test, (iii) test if  $u > u_t$ , (iv) if  $u < u_t$ , or (v) if  $u_t < u < u'_t$ .

The expected utility equals to

1.  $E(u|x_0)$  for accepting without testing,
2.  $T(r, v)$  with testing, and
3.  $v_0$  if the action is dropped and a new set ( $S_A$ ) is selected according to the optimal policy.

## 4 METHOD

This section focuses on the application of the STS model to Web services. At first we describe heuristics for estimating cost functions ( $c_s, c_t$ ), and the common probability mass function  $h(x_0, x_1, u)$ . Afterwards the process of applying search-test-stop to tagging applications is elaborated.

### 4.1 Cost Functions

In the conventional STS model costs refer to the investment in terms of time and money for gathering information. By applying this idea to software, costs comprise all expenses in terms of CPU-time, bandwidth and storage cost necessary to search for or test certain answers.

Large scale Semantic Web projects, like the ID-IOM media watch on climate change (Scharl et al., 2007), process hundred thousands of pages a week. Querying geonames for geo-tagging such numbers of documents would add *days* of processing time to the IDIOM architecture.

This research focuses solely on costs in terms of response time, because they are the limiting factor in our current research projects. Other applications might require extending this approach to additional factors like CPU-time, bandwidth, etc.

### 4.2 Utility Distributions

Applying the STS model to economic problems yields cash deposits and payments. Transferring this idea to information science is a little bit more subtle, because the utility is highly dependent on the application and its user’s preferences. Even within one domain the notion of an answer set’s ( $S_A$ ) value might not be clear. For instance in a geo context the “correct” answer for a certain problem may be a particular mountain in Austria, but the geo-tagger might not identify the mountain but the surrounding region or at least the state in which it is located (compare Figure 1). Assigning concrete utility values to these alternatives is not possible without detailed information regarding the application and user preferences. Approaches for evaluating the set’s value might therefore vary from binary methods (full score for correct answers; no points for incomplete/incorrect answers) to complex ontology based approaches, evaluating the grade of correctness and severe of deviations.

### 4.3 Application

This work has been motivated by performance issues in a geo-tagging application facilitating resources from geonames.org and WordNet for improving tagging accuracy. Based on the experience garnered during the evaluation of STS models, this section will present a heuristic for determining the cost functions ( $c_s, c_t$ ) and the common probability mass function  $h(x_0, x_1, u)$ .

#### 4.3.1 Cost Functions

Searching leads to external queries and therefore costs. Measuring a service’s performance over a certain time period allows estimating the average response time and variance.

STS fits best for situations, where the query cost  $c_s$  is in the same order as the average utility retrieved ( $O(c_s) = O(\bar{u})$ ). In settings with  $O(c_s) \ll O(\bar{u})$  the search costs have no significant impact on the utility and if  $O(c_s) \gg O(\bar{u})$  no searching will take place at




Figure 2: Database schema of a simple tagger.

all (because the involved costs are much higher than the possible benefit).

In real world situations the translation from search times to costs is highly user dependent. To simplify the comparison of the results, this research applies a linear translation function  $c_s = \lambda \cdot t_s$  with  $\lambda = const = 1/\tilde{t}_s$  yielding costs of  $O(c_s) = 1$ . To reduce the influence of service outages the median of the response times  $\tilde{t}_s$  has been selected and a timeout of 60 seconds for any search operation is implemented.

### 4.3.2 Utility Distribution

The discrete common probability mass function  $h$  is composed of three components: The probability mass function of (i) the utility  $u$ , (ii) the random variable  $x_0$  providing an estimate of the utility and, (iii) the random variable  $x_1$  containing a refined estimate of the answer's utility.

In general a utility function assuming linearly independent utility values might look like Equation 5.

$$u = \sum_{S_A} \lambda(i) f_{eval}(i) \quad (5)$$

The utility equals to the sum of the utility gained by each answer set  $S_A$ , which is evaluated using an evaluation function  $f_{eval}$ , and weighted with a factor  $\lambda(i)$ . To simplify the computation of the utility we consider only correct answers as useful (Equation 6) and apply the same weight ( $\lambda(i) = const = 1$ ) to all answers.

$$f_{eval}(i) = \begin{cases} 0 & \text{if } a_i \text{ incorrect;} \\ 1 & \text{if } a_i \text{ correct.} \end{cases} \quad (6)$$

Geo-tagging identifies geographic entities based on a knowledge base as for instance a gazetteer or a trained artificial intelligence algorithm.

After searching the number of identified entries ( $|S_a| = x_0$ ) provides a good estimation of the expected value of the answers utility. Applying a focus algorithm (e.g. (Amitay et al., 2004)) yields a refined evaluation of the entity set ( $|S'_a| = x_1$ ) resolving geo ambiguities.  $S'_a$  might still contain incorrect answers due to errors in the geo disambiguation or due to ambiguous terms not resolved by the focus algorithm (e.g. turkey/bird versus Turkey/country). Based on

the probabilities of a particular answer  $a_i \in S_a/a'_i \in S'_a$  of being incorrect  $P_{incorr}(a_i)/P_{incorr}(a'_i)$  the expected value  $u$  for a given combination of  $x_0, x_1$  is determined. Evaluating historical error rates yields estimations for  $P_{incorr}(a_i)$  and  $P_{incorr}(a'_i)$ .

If no historical data is available heuristics based on the number of ambiguous geo-entries are useful for providing an educated guess of the probabilities.

A tagger recognizes patterns based on a pattern database table. The relation `hasPattern` translates these patterns to `TaggingEntities` as for instance spatial locations, persons, and organizations. Figure 2 visualizes a possible database layout for such a tagger. Unfortunately, the `hasPattern` table often does not provide a unique mapping between patterns and entities - names as for instance Vienna may refer to multiple entities (Vienna/Austria versus Vienna/Virgina/US). On the other side many entities have multiple patterns associated with them (e.g. Wien, Vienna, Vienne, Bech, etc.). Based on the database schema above,  $P_{incorr}(a_i)$  for such a tagger is estimated using the following heuristic:

$$n_{Entities} = |TaggingEntity| \quad (7)$$

$$n_{Mappings} = |hasPattern| \quad (8)$$

$$n_{ambiguous} = |\sigma_{[isAmbiguous='true']}(TaggingEntry * hasPattern)| \quad (9)$$

$$P_{incorr} = 1 - \frac{n_{Entities}}{n_{Mappings} + n_{ambiguous}} \quad (10)$$

Extending the database schema visualized in Figure 2 to non geo entries using WordNet and applying Equations 7-10 yields  $P_{incorr}(a'_i)$ .

## 5 EVALUATION

For evaluating the STS model's efficiency in real world applications a simulation framework, supporting (i) item a solely coverage based decision logic and the search-test-stop model, (ii) artificial (normal distribution) and measured (compare Section 2) distributions of network response times, and (iii) common probability mass functions  $h(x_0, x_1, u)$  composed

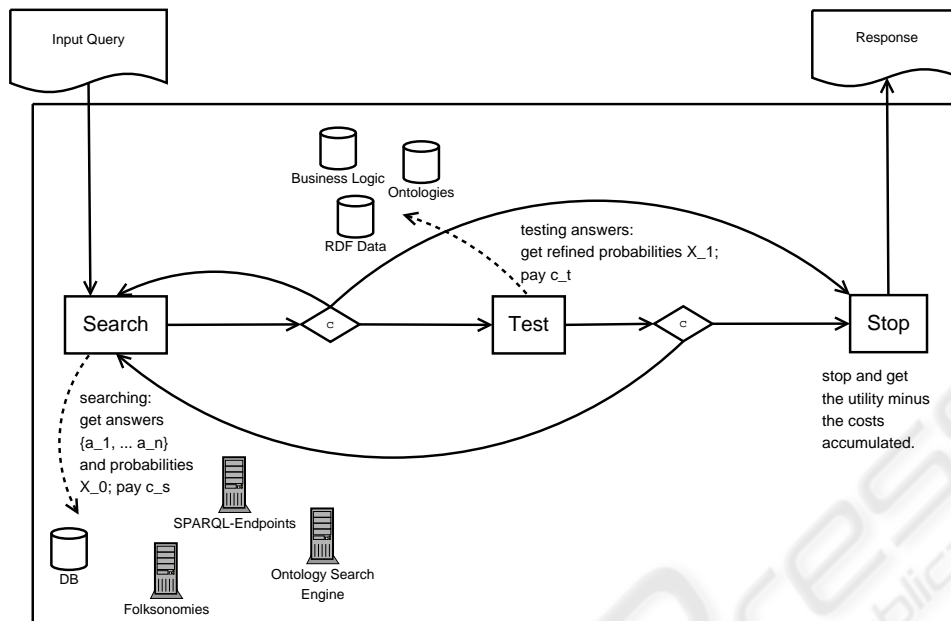


Figure 3: The search-test-stop approach.

from user defined  $P_{incorr}(a_i)$  and  $P_{incorr}(a'_i)$  settings has been programmed.

Integration of the python numarray library<sup>6</sup> enables efficient processing of matrix operations as required for computing decisions based on the search-test-stop model.

To prevent the coverage based decision logic from delivering large amounts of low quality answers, the simulation controller only accepts answers with an expected utility above a certain threshold ( $u_{min}$ ). In contrast the search-test-stop algorithm computes  $u_{min} = r_D$  on the fly, based on the current responsiveness of the external service and the user's preferences.

### 5.1 Performance

Comparing the two approaches at different minimum quality levels ( $u_{min}$ ), and service response time distributions approximated by a normal distribution  $N(\bar{r}, \sigma_r^2)$  yields Table 2. The common probability mass functions has been composed with  $P_{incorr}(a_i) = 0.3$ ,  $P_{incorr}(a'_i) = 0.1$ . The parameters for the normal distribution are  $c_s = N(2, 1.9)$  for high search costs,  $c_s = N(1, 0.9)$  for medium search costs, and  $c_s = N(0.5, 0.4)$  for low search costs.

Table 2 evaluates the search strategies according to two criteria: (i) the quality  $\bar{u}$ , the average utility of an answer set ( $S_A$ ) retrieved by the strategy, and (ii) the quantity  $\frac{\Delta u}{\Delta t}$  - the rate at which the number of correct answers (and therefore the total utility ( $u$ )) grows.

<sup>6</sup>sourceforge.net/projects/numpy

High  $\bar{u}$  values correspond to accepting only high quality results, with a lot of correct answers, and dropping low quality answer sets (at the cost of a lower quantity).

The conventional coverage based approach (Conv) delivers the highest quantity for small  $u_{min}$  values because virtually all answers are accepted and contribute to the total utility. This greedy approach comes at the cost of a lower answer quality and therefore low average utility  $\bar{u}$  per answer. Increasing  $u_{min}$  yields a better answer quality, but lower quantity values. At high search costs this strategy's performance is particularly unsatisfactory, because it doesn't consider the costs of the search operation.

In contrast to the conventional approach STS maximizes answer quality and quantity based on the *current* search cost adjusting queries to the responsiveness of the service and the user's preferences. These preferences formalize the trade off between quality and quantity by specifying a transformation function between search cost and search times.

STS therefore optimizes the agent's behavior in terms of user utility. This does *not* mean that STS minimizes resource usage. Instead STS dynamically adjusts the resource utilization based on the cost of searching ( $c_s$ ) and testing ( $c_t$ ), providing the user with optimal results in terms of accuracy *and* response times.

Enforcing a minimal utility  $u_{min}$  boosts the average utility  $\bar{u}$  of the non STS service, but at the cost of a higher resource utilization, independent from

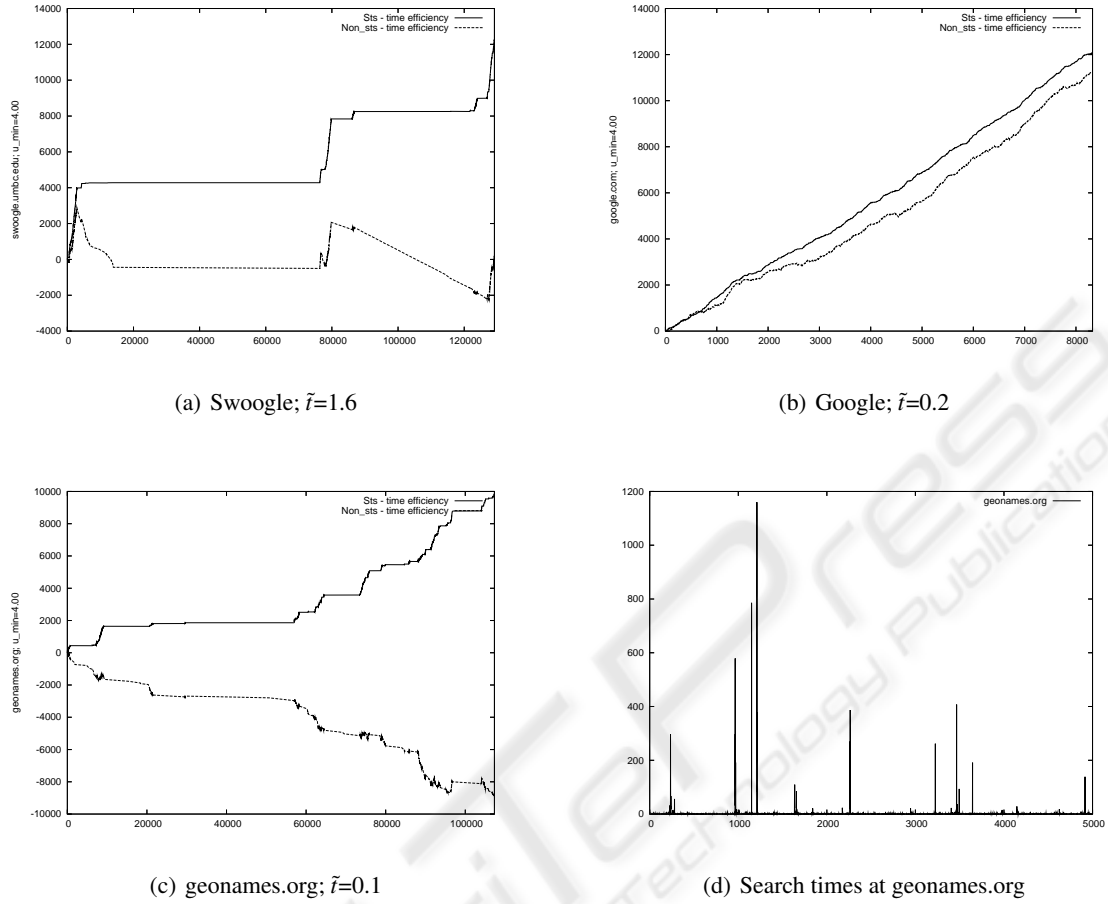


Figure 4: Search-test-stop (STS) versus conventional (NON-ST) decision logic.

Table 2: Tagging performance.

Search Cost ( $c_s$ )	$u_{min}$	Quality ( $\bar{u}$ )		Quantity ( $\frac{\Delta u}{\Delta t}$ )	
		STS	Conv	STS	Conv
low	2	6.62	5.58	3.47	7.79
low	4	6.64	6.13	3.56	6.93
low	6	6.69	6.55	3.57	5.95
low	8	6.66	6.39	3.55	2.75
medium	2	4.99	4.84	1.88	3.22
medium	4	5.02	5.15	1.92	2.76
medium	6	5.01	5.32	1.89	2.27
medium	8	5.00	3.86	1.87	0.79
high	2	2.81	3.20	0.78	1.05
high	4	2.75	3.25	0.76	0.88
high	6	2.84	2.81	0.80	0.59
high	8	2.81	-0.91	0.76	-0.09

the server's load (leading to extremely high response times during high load conditions). Static limits also do not consider additional queries at idle servers, leading to lower utilities under low load conditions.

In contrast to the conventional approach STS (i) utilizes dormant resources of idle servers, and (ii) spares resources of busy servers, maximizing utility according to the user's preferences.

## 5.2 Web Services

In this section we will simulate the effect of STS on the performance of real world Web services, using search costs as measured during the Web service timing (compare Section 2). Figure 4 visualizes the application of the search-test-stop model to Web services. The simulation facilitates the cost and common probability mass functions from Section 5.

Figure 4 compares the tagger's performance for three different Web services (Swoogle, Google, geonames) with  $u_{min} = 4$ . The fourth figure visualizes geoname's response times over the observation period of five weeks. In all three use cases STS performs well, because the search times are adjusted according to the service's responsiveness. Geonames and

Swoogle experience the highest performance boost, due to high variances in the search cost, leading to negative utility for the conventional query strategy. Services with low variances ( $\sigma_{r_i}^2$ ) in their response times as for instance Google, del.icio.us and Wikipedia benefit least from the application of the STS model, because static strategies perform better under these conditions.

## 6 OUTLOOK AND CONCLUSIONS

This work presents an approach for optimizing access to third party remote resources. Optimizing the clients resource access strategy yields higher query performance and spares remote resources by preventing unnecessary queries.

The main contributions of this paper are (i) applying the search-test-stop model to value driven information gathering, extending its usefulness to domains where one or more testings steps allow refining the estimated utility of the answer set; (ii) demonstrating the use of this approach to semantic tagging, and (iii) evaluating how the search-test-stop model performs in comparison to a solely value based approach.

The experiments show that search-test-stop and value driven information gathering perform especially well in domains with highly variable search cost.

In this work we only use one level testing, nevertheless, as Hartmann has shown (Hartmann, 1985) extending STS to n-levels of testing is a straight forward task. Future research will transfer these techniques and results to more complex use cases integrating multiple data sources as for instance semi automatic ontology extension (Liu et al., 2005). The development of utility functions considering partially correct answers and user preferences will allow a more fine grained control over the process's performance yielding highly accurate querying strategies and therefore better results.

## ACKNOWLEDGEMENTS

The author wishes to thank Prof. Wolfgang Janko for his valuable feedback and suggestions. The project results have been developed in the IDIOM (Information Diffusion across Interactive Online Media; [www.idiom.at](http://www.idiom.at)) project funded by the Austrian Ministry of Transport, Innovation & Technology (BMVIT) and the Austrian Research Promotion Agency (FFG).

## REFERENCES

- Amitay, E., Har'El, N., Sivan, R., and Soffer, A. (2004). Web-a-where: geotagging web content. In *SIGIR '04: Proceedings of the 27th annual international ACM SIGIR conference on Research and development in information retrieval*, pages 273–280, New York, NY, USA. ACM.
- Grass, J. and Zilberstein, S. (March 2000). A value-driven system for autonomous information gathering. *Journal of Intelligent Information Systems*, 14:5–27(23).
- Gupta, C., Bhowmik, R., Head, M. R., Govindaraju, M., and Meng, W. (2007). Improving performance of web services query matchmaking with automated knowledge acquisition. In *Web Intelligence*, pages 559–563. IEEE Computer Society.
- Hartmann, J. (1985). *Wirtschaftliche Alternativensuche mit Informationsbeschaffung unter Unsicherheit*. PhD thesis, Universität Fridericiana Karlsruhe.
- Horvitz, E. J., Breese, J. S., and Henrion, M. (1988). Decision theory in expert systems and artificial intelligence. *International Journal of Approximate Reasoning*, 2:247–302.
- Ipeirotis, P. G., Agichtein, E., Jain, P., and Gravano, L. (2007). Towards a query optimizer for text-centric tasks. *ACM Trans. Database Syst.*, 32(4):21.
- Lehmann, E. L. and Romano, J. P. (2005). *Testing Statistical Hypotheses*. Springer, New York, 3rd edition edition.
- Liu, W., Weichselbraun, A., Scharl, A., and Chang, E. (2005). Semi-automatic ontology extension using spreading activation. *Journal of Universal Knowledge Management*, 0(1):50–58.
- MacQueen, J. (1964). Optimal policies for a class of search and evaluation problems. *Management Science*, 10(4):746–759.
- Scharl, A., Weichselbraun, A., and Liu, W. (2007). Tracking and modelling information diffusion across interactive online media. *International Journal of Meta-data, Semantics and Ontologies*, 2(2):136–145.