Source Selection for Image Retrieval in Peer-to-Peer Networks

Daniel Blank

Media Informatics Group, University of Bamberg
Feldkirchenstr. 21, 96052 Bamberg, Germany

<http://www.uni-bamberg.de/minf>
daniel.blank@uni-bamberg.de

Abstract

With the emergence of web albums such as *Flickr.com* or *Picasa.com*, the amount of personal image collections administered on the web has increased dramatically. As a consequence, efficient storing, indexing and retrieval techniques are needed. Peer-to-peer (P2P) networks are an interesting solution to maintain large image collections. When performing a query on certain types of P2P networks, source selection is very important. In our scenario, compact summaries of each peer's image collection, which are known to other peers, are used to determine the most promising peers for a given query. These summaries have to address (1) date and time information, (2) textual information, (3) geolocations and (4) content-based image features. The present paper outlines a large-scale image retrieval system, relying on data summaries and source selection strategies. While our scenario is based on a P2P system, we also describe how results can be transferred to other application domains such as distributed information retrieval (distributed IR) or tree-based index structures.

Keywords: Source Selection, Peer-to-Peer IR, Image Retrieval, Distributed IR, CBIR

1. INTRODUCTION

Web albums such as *Flickr.com* or *Picasa.com*, which offer storage capabilities for personal photo collections online, have become very popular in the last years. People upload their photos in order to share them with friends and to interact with each other e.g. by collaboratively tagging the photos.

Different criteria for image retrieval can be identified in such a scenario: mainly (1) date and time information, (2) tags and textual descriptions, (3) the geographic footprint, and (4) content-based image features describing e.g. colour or texture.

Our work addresses image search employing these criteria in peer-to-peer (P2P) networks. P2P scenarios for the administration of large image collections are attractive for multiple reasons. Photo collections can be stored locally on people's individual PCs. No expensive infrastructure has to be maintained by applying a scalable P2P protocol such as Rumorama [21] and remote computing power can be used to maintain the image collection. Users can decide which image features to publish in order for the corresponding images to become searchable without any need for crawling activities.

Our approach is based on—but not restricted to—Rumorama [21], a scalable P2P protocol building hierarchies of networks that are accessible by an efficient multicast. Its leaf networks behave like PlanetP networks [8]. In PlanetP, randomised rumour spreading assures that every peer knows summaries of all other peers' data in the network. The summaries provide the basis for source selection decisions, i.e. which peers to contact during query processing. While we examine peer summaries and source selection strategies for image retrieval in PlanetP-like middle-sized networks, we can easily extend this to large-scale Rumorama-like P2P networks.

This paper is organised as follows. Section 2 gives a brief overview on related work. In section 3, the more general applicability of peer summaries in different application scenarios is discussed.

Section 4 describes our main approach, different summary-types and source selection strategies. In section 5, we conclude with an outline of challenging research questions for future work in order to design a large-scale, distributed image retrieval system based on source selection.

2. RELATED WORK

To our best knowledge, there are no multi-feature P2P information retrieval (IR) systems that allow for text-based and content-based image retrieval (CBIR) additionally employing temporal and geographic metadata. In our scenario, in order to support image retrieval based on these criteria, a peer will have to maintain and distribute at least four different summary types. For the summarisation of linear time and date information, we assume that our summaries presented in section 4 are directly applicable. Traditional histogram techniques (cf. [17]) as well as techniques used for aggregating sensor data (cf. e.g. [24]) can also be applied. Summarising the textual annotations and descriptions for image retrieval will be part of future work. Within this paper in section 4, we focus on the summarisation of multi- and high-dimensional data.

In general, P2P IR systems can be classified into several groups (cf. [6])¹. Systems of the first group follow a semantic query routing approach based on peer summaries. Routing Indices [9] are among the first approaches presented in literature belonging to this group. Based on summary information of neighboring peers that is aggregated along multiple hops, a peer routes queries towards the direction of peers potentially containing relevant documents w.r.t. the query. In order to restrict the size of peer summaries, topics are indexed rather than individual terms.

As opposed to Routing Indices, which follow a multi-hop semantic routing approach, PlanetP [8] and its scalable extension Rumorama [21] apply single-hop semantic routing. Therefore, summaries are sent to all peers in a (sub-)network. Summaries and source selection strategies for text data based on Bloom filters are analysed in [8, 11].

The single-hop semantic query routing approach originally comes from distributed IR (for references see e.g. [22, 25]). Many of the approaches proposed in literature assume that it is feasible to base routing decisions on term frequency or term weighting information which is available for all terms of a particular resource. But, summaries within P2P IR systems need to be more space efficient than the ones designed for distributed IR, because of limited bandwidth capacities and the frequent joining and leaving of peers, often with updated document collections. Therefore, most of the traditional distributed IR approaches are not directly applicable.

The second group of P2P IR systems are semantic overlay networks (e.g. [18]) where the content of a peer's data defines its place within the network topology. Peers are organised by semantic clusters and within query execution the query is routed to the most promising cluster(s). Here, the indexing of multiple feature types, e.g. textual information as well as image content, would require the definition of a similarity between peers combining textual and image content information. Alternatively, several overlays might be maintained inducing a higher maintenance effort.

A third class of P2P IR systems is represented by distributed indexing structures with distributed hashtables (DHTs) as its most prominent class member. Minerva [2] has been designed for the administration of text documents, where term statistics are indexed in a DHT. Every peer is responsible for a certain set of terms. Novak et al. have presented a large-scale CBIR architecture [23] based on a DHT. Within DHTs, indexing data of a peer's content is transferred to remote peers with every peer being responsible for a certain range of the feature domain of an individual feature. Presumably, for example, correlations between geographic information and image content are difficult to exploit. If we e.g. assume an image from the Sahara Desert with shades of beige sand and blue sky, different peers might be responsible for indexing the geographic and the image content information. Therefore, when distributing the indexing data of the Sahara image, querying for it, or removing it from the network, two different peers (at least)

¹In the following, we only describe *true* P2P IR systems. Super-peer scenarios, where some peers are chosen to perform certain tasks because of increased capabilities of their resources, are not analysed in more detail. Nevertheless some of the strategies presented in section 4 can also be applied in a super-peer scenario (for a brief description see section 3).

have to be contacted. Even with only one feature type being indexed (e.g. text in the case of Minerva), the frequent joining and leaving of peers leads to an increase in network traffic as term statistics are transferred to or removed from remote peers.

In order to summarise collections of personal photos for distributed CBIR, other approaches than the ones presented in section 4 are also possible. Earlier work in distributed CBIR follows a clustering approach. Chang et al. [7] create summaries of remote databases based on image templates, i.e. feature vectors of reference images, sampled from remote databases. Hierarchical clustering is applied in order to obtain a set of cluster centroids. The authors present two approaches to source selection, the first based on statistical information and the second based on histograms. The latter is similar to our approach, being different for example w.r.t. the sampling phase used for obtaining the centroids, the clustering process itself, the way how histograms are computed, and especially w.r.t. the number of centroids used for computing the summaries.

Berretti et al. [3] apply a special form of hierarchical clustering to the image features of a remote database. With the use of a threshold similarity it is possible to adjust the number of centroids, i.e. the granularity and size of the resource descriptions. The resource descriptions consist of the centroids themselves. In [14] we also applied a local clustering technique using mixtures of Gaussians. For two-dimensional geographic data, cluster hulls have been proposed in order to summarise sets of geographic coordinates of an individual resource by several convex hulls [16].

In general, we expect approaches explicitly transferring centroids to be less space efficient than our approach presented in section 4, especially for high-dimensional feature vectors. Histogram-techniques implicitly using centroid information seem to be more promising.

3. APPLICABILITY OF SUMMARIES AND SOURCE SELECTION STRATEGIES

We analyse summaries for PlanetP-like P2P systems (cf. section 4). Nevertheless, these summaries are not restricted to PlanetP and related protocols that make use of single-hop semantic routing. Within multi-hop semantic routing networks, summaries will have to be aggregated along multiple hops. Summarisation of peer content is also needed in semantic overlay networks in order to derive a peer's place within the network topology. In this case, it is necessary to define a similarity between peer summaries so that peers with similar summaries can be grouped together into "clusters of interest".

Many P2P IR protocols rely on super-peers. Typically, they are characterised by increased storage or bandwidth capabilities w.r.t. "normal" peers. Super-peers also tend to stay in the system for most of the time. Query routing is usually performed amongst super-peers and "normal" peers transfer their indexing data to responsible super-peers. Only if a document is transferred from a "normal" peer to another, this is done at the peer level without involving super-peers. Therefore, our summaries can also be used in a super-peer scenario for content summarisation of "normal" peers and as the basis for query routing amongst super-peers.

We believe that our work can also be beneficial for the design of tree-based index structures. Summaries in the P2P context with their enforced space limitations are similar to approximations maintained in inner nodes of a tree. Indexing structures e.g. relying on minimum bounding rectangles such as the R-tree [15] and its variants can provide the basis for summary construction of geographic data [5]. Signatures / Bloom filters are used to summarise textual data (cf. e.g. [8, 10]).

During the last years, local, content-based image features like e.g. SIFT [19] have become popular. In the case of SIFT, an individual image is characterised by several hundred 128-dimensional feature vectors. This poses increased challenges on indexing techniques. We believe that our summarisation and source selection techniques presented in section 4 can be adapted to summarise and index SIFT features, both on a per-image as well as on a per-collection basis.

After more than a decade of P2P research, application scenarios, where P2P IR technology is successfully used, are still missing. At the moment, cloud computing—with some respect

oppositional to the P2P concept—seems to be appealing. Large computer infrastructures with free capacities are used in order to provide software as a service as well as storage capacities online. Following this trend, it might become reasonable for companies and individuals to use several different storage services in parallel because of pricing and availability reasons. Therefore, source selection based on compact data summaries might become important.

Summarisation of documents as well as source selection strategies are of course also important in traditional client/server applications. If we think of market situations with many buyers and sellers, source selection strategies based on summarisations of indexing data might provide a benefit for all participants. There is for example a huge amount of photo agencies and services providing images that are subjective to licence conditions and charge. For clients searching via traditional web search, it is difficult to identify the most promising service providers in order to browse through their sites. Also the design of a meta search engine might be difficult since textual metadata is often stored in content management systems, largely hidden for traditional web crawlers. The extraction of content-based image features is hampered as images are often only available in a very small resolution, with many of them being modified by watermarks. In such a scenario, providing indexing data for a centralised indexing broker might be beneficial for service providers to gain attention. Additionally, users might benefit from this service, as automatic source selection would prevent them from browsing too many irrelevant sites. Brokers might gain revenue—similar to traditional search engines—through advertisements and/or online auctions with service providers bidding for high ranks.

Personal metasearch is a novel application domain of distributed IR [25]. All of a user's online resources are summarised and metasearch is provided based on summaries integrating heterogeneous resources that largely vary in size. This is oppositional to many traditional distributed IR scenarios. It is therefore important to make use of source selection strategies which—in the case of varying resource sizes—do not prefer to contact large collections as this is not always a good choice. Additionally, personalised online resources consist of different types of data with high update frequencies (email accounts, Web sites, photo and video sharing communities, local databases etc.). We therefore believe that our summaries designed for P2P networks can also be applied for personal metasearch, as we require them to be selective (also being able to identify promising small peers administering few documents) and space efficient (because of the dynamic nature of P2P systems). Additionally, we target on multiple summary types for temporal, geographic, textual and content-based (meta)data useful for the summarisation of textual, image, audio and video content.

4. SOURCE SELECTION STRATEGIES FOR IMAGE RETRIEVAL

We have analysed cluster histograms for summarising collections of real-valued feature vectors [12]. In order to compute cluster histograms, all peers need to know a unique set of cluster centroids in feature space. A peer joining the network will try to obtain the centroids from peers already present. If peer p has obtained the set of centroids $C = \{c_i | 1 \leq i \leq \kappa\}$, peer p assigns every feature vector of its local collection to the closest centroid c_i , according to a given distance measure. The same distance measure is used by all peers in the system. So, peer p computes a histogram that assigns to every centroid c_i the number of feature vectors closest to c_i . The cluster histogram summarises the data collection of peer p . Peers publish their summaries by randomised rumour spreading.

Obviously, the combination of summaries and source selection strategies is crucial for the performance of query processing. In [12], three source selection strategies have been evaluated. The most promising strategy sorts the cluster centroids c_i in a list L in ascending order according to their distance to the query. The first element out of L corresponds to the centroid of the cluster that is closest to the query. Peers with many documents inside this so called query cluster are ranked higher than peers with less documents in the query cluster. If two peers share the same amount of documents in the currently analysed cluster, the next element out of L is chosen and the two peers are recursively ranked w.r.t. the number of documents within the current cluster.

As an example, let us assume three centroids ($\kappa = 3$), two peers p_a and p_b with corresponding summaries $s_a := (5, 8, 3)$ and $s_b := (7, 8, 2)$. We assume that centroid c_2 is the closest and c_1 is the second closest centroid w.r.t. query feature vector q . Both peers have assigned 8 image feature vectors to the cluster represented by centroid c_2 . But, out of peer p_b 's image feature vectors 7 are closest to c_1 , compared to 5 feature vectors that are closest to c_1 out of peer p_a 's image collection. As $7 > 5$, peer p_b is ranked higher than peer p_a and contacted before p_a during query processing.

An important finding in [12] is that a distributed clustering for computing the set of centroids might be dispensable. A random selection of feature vectors out of the data collection that is administered in the P2P system can be used as centroids. Their usage results in a minimal decrease in retrieval performance at the same time making distributed clustering obsolete. In internal experiments also other variants of distributed clustering like fuzzy k -Means clustering and self-organising maps could not improve retrieval performance.

In [13] we analysed the performance of cluster histograms for a large number of distance measures and image features. It became clear that our source selection strategy is influenced by the curse of dimensionality.

The creation of cluster histograms presented so far needs some global knowledge. The peers present in the system must agree on the set of centroids. This does not affect the usability of cluster histograms, but it restricts the adaptivity of our approach, as an update of the centroids becomes expensive. Approaches where peers do not have to agree on the set of centroids are therefore desirable. Local clustering techniques and Gaussian mixture models (GMMs), that do not rely on global knowledge when computing the summaries, are analysed in [14]. When performing local clustering, a peer computes a small number of local clusters and their centroids are published as summaries. GMMs model the point density distribution as superposition of Gaussians with different means and covariances. GMMs outperform local clustering. But, in our experiments relatively small cluster histograms (256 bins or clusters) with globally distributed cluster centroids are superior to GMMs in terms of retrieval performance. Therefore we optimised the former strategy as explained in the following.

Our new summaries—called highly fine summaries (HFS)—evolve from the cluster histograms described earlier by varying κ , the number of cluster centroids [4]. We increase the number of centroids from 256 to e.g. 16,384 or even more. This offers several benefits for our scenario:

- Retrieval performance is improved since the data space is partitioned in a more fine-grained way.
- At the same time, the costs for distributing the summaries only increase moderately as we compress the summaries. We use runlength encoding which allows us to substantially reduce summary sizes. This is possible since with large numbers of centroids, many of the small peers (i.e. peers administering few images) will compute summaries with some histogram entries set to very small values (often 1), but most of the histogram values will stay 0 as no image is assigned to the corresponding centroid.
- Administration overhead for distributing the centroids even decreases as we distribute the set of centroids with software updates. Within our experiments we showed that if we choose the centroids from a different, disjoint collection within the same application domain (we use images from *Flickr.com*), average retrieval performance is not affected.

HFS are designed for high-dimensional image feature vectors. In [5] we analyse the applicability of modified HFS—now called ultra fine summaries (UFS)—in the context of summarising sets of geographic coordinates. Instead of cluster histograms, we use bit vectors, where bit i is set, if for any image out of a peer's collection, centroid c_i is the closest. Otherwise, the corresponding bit remains zero. We have evaluated the usage of UFS for point queries against summarising a peer's geographic footprint by either a minimum bounding rectangle or a grid-based, binary index. UFS show the best performance in terms of selectivity. Therefore we will evaluate source selection strategies for k -nearest-neighbour queries based on UFS in future work.

5. CHALLENGES AND FUTURE DIRECTIONS

HFS/UFS seem promising for summarising content-based image features as well as sets of geographic coordinates. Other important criteria for image retrieval are time and date information and textual information. We believe that HFS/UFS might also be promising for summarising text data. Therefore, we will compare HFS/UFS with traditional Bloom filter approaches. Spectral Bloom filters might be employed to encode term frequency information [11] or impact information [1] into the summaries. In order to do so, a large, distributed test collection is needed that offers a realistic distribution of text documents to peers. As our main application scenario is image retrieval, we might use image collections from Flickr.com, being crawled together with textual descriptions and tags.

When compressing HFS (i.e. histograms of integer values) of big peers, the summaries might become large. In the future, we will analyse the differences between UFS and HFS in more detail. Retrieval performance might decrease as UFS do no longer maintain frequency information about how many of a peer's images are closest to a certain centroid. At the same time, UFS allow us to use more centroids as only a single bit is needed to encode, if any of a peer's images is closest to a certain centroid. Additionally, UFS might be beneficial for compression as bit vectors with many bits set to 1 are still suited for compression.

In order to restrict the summary size of big peers, it is necessary that a peer is granted a maximum amount of space to encode its summary. The size of a summary in the case of HFS might therefore be chosen depending on n , i.e. the number of documents that a peer administers, e.g. by multiplying a basic summary size with $1+\log(n)$ or similar factors. Another approach for restricting the size of the summaries might be to hierarchically partition the data space. A peer can then choose the number of centroids according to this space partitioning so that the overall summary size does not exceed a given upper bound.

We have to derive a better stopping criterion determining if it is promising to contact further peers or to stop query processing. Currently we stop after having contacted a certain fraction of peers. This value is determined through experiments. In future, we will analyse and adapt solutions originally proposed for textual data [8, 20] and design new ones if necessary.

Within our approach, we use a secondary collection from which the centroids are chosen. Currently we choose them randomly. But, applying clustering techniques or specialised selection techniques might be beneficial in order to gain in retrieval performance. A distance matrix D with pairwise distances between centroids may provide additional benefits. The ordering of indexes within HFS/UFS depending on D or based on some type of space filling curve might be beneficial for compression. Furthermore, using D in addition with the triangular inequality might be helpful in order to derive algorithms for a centralised indexing structure based on our HFS/UFS approach.

The secondary collection might also be used in order to apply dimensionality reduction. In [13], we have seen that the quality of the source selection decreases with increasing dimensionality of the underlying feature space. Therefore, the secondary collection and the centroids might provide a basis for local dimensionality reduction. We will therefore compare the effects of distributed PCA (PCA: principal component analysis) with local PCA based on the secondary collection as well as other local techniques for dimensionality reduction.

As discussed earlier in section 3, using modified versions of HFS/UFS in combination with adapted source selection strategies seems also promising in order to index local feature descriptors with several hundred feature vectors per image – both on a per-image as well as on a per-collection basis.

REFERENCES

- [1] Anh V. N. and Moffat A. (2005) *Simplified similarity scoring using term ranks*. Proc. of 28th Int. ACM SIGIR Conf. on Research and Development in Information Retrieval, pp. 226-233, Salvador, Brazil.

- [2] Bender M., Michel S., Weikum G. and Zimmer C. (2005) *The MINERVA Project: Database Selection in the Context of P2P Search*. GI-Conf. on Database Systems in Business, Technology and Web, pp. 125-144, Karlsruhe, Germany.
- [3] Berretti S., Del Bimbo A. and Pala P. (2004) *Merging Results for Distributed Content Based Image Retrieval*. Multimedia Tools and Applications, 24(3), pp. 215-232, Kluwer.
- [4] Blank D., El Allali S., Müller W. and Henrich A. (2007) *Sample-based Creation of Peer Summaries for Efficient Similarity Search in Scalable Peer-to-Peer Networks*. 9th ACM SIGMM Int. Workshop on Multimedia Information Retrieval. pp. 143-152, Augsburg, Germany.
- [5] Blank D. and Henrich A. (2009) *Summarizing Georeferenced Photo Collections for Image Retrieval in P2P Networks*. Int. Workshop on Geographic Information on the Internet, Toulouse, France.
- [6] Blank D., Müller W. and Henrich A. (2008) *Designing Benchmarks for the Evaluation of Peer-to-Peer Information Retrieval Systems*. In: Baumeister J., Atzmüller M. (Edts.), GI Workshop: Lernen, Wissen & Adaptivität, LWA 2008, Würzburg, Germany.
- [7] Chang W., Sheikholeslami G., Wang J. and Zhang A. (1998) *Data Resource Selection in Distributed Visual Information Systems*. IEEE Trans. on Knowl. and Data Eng., 10(6):926-946.
- [8] Cuenca-Acuna F. M., Peery C., Martin R. P. and Nguyen T. D. (2003) *PlanetP: Using Gossiping to Build Content Addressable Peer-to-Peer Information Sharing Communities*. Proc of the 12th IEEE Int. Symp. on High Performance Distributed Computing. pp. 236-246, Seattle, WA, USA.
- [9] Crespo A. and Garcia-Molina H. (2002) *Routing Indices For Peer-to-Peer Systems*. Proc. of 22nd IEEE Int. Conf. on Distributed Computing Systems, pp. 23-32, Vienna, Austria.
- [10] Deppisch U. (1986) *S-tree: a dynamic balanced signature index for office retrieval*. Proc. of 9th Int. ACM SIGIR Conf. on Research and Development in Information Retrieval, pp. 77-87, Pisa, Italy.
- [11] Eisenhardt M., Müller W. and Henrich A. (2004) *Spektrale Bloom-Filter für Peer-to-Peer Information Retrieval*. Lecture Notes in Informatics, No. P-51, pp. 44-48, Ulm, Germany – ISBN 3-88579-380-6.
- [12] Eisenhardt M., Müller W., Henrich A., El Allali S. and Blank D. (2006) *Clustering-based source-selection in summary-based P2P networks*. Proc. of 8th IEEE Int. Symp. on Multimedia, pp. 823-830, San Diego, CA, USA.
- [13] El Allali S., Blank D., Eisenhardt M., Henrich A. and Müller W. (2007) *Untersuchung des Einflusses verschiedener Bild-Features und Distanzmaße im inhaltsbasierten P2P Information Retrieval*. GI-Conf. on Database Systems in Business, Technology and Web, pp. 382-396, Aachen, Germany.
- [14] El Allali S., Blank D., Müller W. and Henrich A. (2008) *Image Data Source Selection Using Gaussian Mixture Models*. Proc. of 5th Int. Workshop on Adaptive Multimedia Retrieval, pp. 170-181, Paris, France, Springer.
- [15] Guttman A. (1984) *R-Trees: A Dynamic Index Structure for Spatial Searching*. Proc. of ACM SIGMOD, pp. 47-57, Boston, MA, USA.
- [16] Hershberger J., Shrivastava N. and Suri S. (2008) *Summarizing spatial data streams using ClusterHulls*. ACM Journal of Experimental Algorithmics, 13 (2.4-2.28).
- [17] Ioannidis Y. (2003) *The History of Histograms (abridged)*. Proc. of 29th Int. Conf. on Very Large Data Bases, pp. 19-30, Berlin, Germany, ACM.
- [18] King I., Ng C. H. and Sia K. C. (2004) *Distributed content-based visual information retrieval system on peer-to-peer networks*. ACM Trans. Inf. Syst. 22(3). pp 477-501.
- [19] Lowe D. G. (2004) *Distinctive image features from scale-invariant keypoints*. Int. Journal of Computer Vision, vol. 60, pp. 91-110.
- [20] Lu J. and Callan J. (2004) *Federated search of text-based digital libraries in hierarchical peer-to-peer networks*. Proc. of SIGIR Workshop on Peer-to-Peer Information Retrieval, Sheffield, UK.
- [21] Müller W., Eisenhardt M. and Henrich A. (2005) *Scalable summary-based Retrieval in Peer-to-Peer Networks*. Proc. of the ACM Int. Conf. on Information and Knowledge Management, pp. 586-593, Bremen, Germany.
- [22] Nottelmann H. and Fuhr N. (2003), *Evaluating different methods of estimating retrieval quality for resource selection*. Proc. of 26th Int. ACM SIGIR Conf. on Research and Development in Information Retrieval, pp. 290-297, Toronto, Canada.
- [23] Novak D., Batko M., Zezula P. (2008) *Web-scale system for image similarity search: When the dreams are coming true*. Workshop on Content-Based Multimedia Indexing, pp. 446-453, London, UK, IEEE.
- [24] Shrivastava N., Buragohain C., Agrawal, D. and Suri, S. (2004) *Medians and beyond: new aggregation techniques for sensor networks*. Proc. of 2nd ACM Int. Conf. on Embedded networked sensor systems. pp. 239-249, Baltimore, MD, USA.
- [25] Thomas P. and Hawking D. (2009) *Server selection methods in personal metasearch: a comparative empirical study*. Information Retrieval, Springer, <http://www.springerlink.com/content/y3704552037635jj/>