

Breaking the Top- k Barrier of Hidden Web Databases*

Saravanan Thirumuruganathan[‡], Nan Zhang^{††}, Gautam Das^{‡†}

[‡]University of Texas at Arlington ^{††}George Washington University [†]Qatar Computing Research Institute

[‡]{saravanan.thirumuruganathan@mavs, gdas@cse}.uta.edu, ^{††}nzhang10@gwu.edu, [†]gdas@qf.org.qa

Abstract—A large number of web databases are only accessible through proprietary form-like interfaces which require users to query the system by entering desired values for a few attributes. A key restriction enforced by such an interface is the *top- k output constraint* - i.e., when there are a large number of matching tuples, only a few (*top- k*) of them are preferentially selected and returned by the website, often according to a proprietary ranking function. Since most web database owners set k to be a small value, the *top- k output constraint* prevents many interesting third-party (e.g., mashup) services from being developed over real-world web databases. In this paper we consider the novel problem of “digging deeper” into such web databases. Our main contribution is the meta-algorithm GetNext that can retrieve the next ranked tuple from the hidden web database using only the restrictive interface of a web database without any prior knowledge of its ranking function. This algorithm can then be called iteratively to retrieve as many top ranked tuples as necessary. We develop principled and efficient algorithms that are based on generating and executing multiple reformulated queries and inferring the next ranked tuple from their returned results. We provide theoretical analysis of our algorithms, as well as extensive experimental results over synthetic and real-world databases that illustrate the effectiveness of our techniques.

I. INTRODUCTION

A. Problem Motivation

Many web databases are “hidden” behind (i.e., only accessible via) a restrictive form-like interface which allows a user to form a search query by specifying the desired values for a few attributes; and the system responds by returning a small number of tuples matching the search query. Almost all such interfaces enforce the *top- k constraint* - i.e., when more than k tuples (where k is typically a predetermined small constant) match the user-specified query, only k of them are preferentially selected according to a (often proprietary) ranking function and returned to the user. For example, American Airline’s (AA) flight search-by-schedule¹ has a default value of $k = 10$. Similarly, Amazon’s best sellers list² for any category only displays the top-100 products.

How to properly set the value of k is an interesting design challenge for a web database owner. On one hand, the owner

may prefer a small k to (1) speed up query processing and shorten the returned webpage, and/or (2) thwart web/tuple scraping. However, in order to accommodate the needs of website users, the value of k should not be too small. Given these two conflicting goals, in practice k is often set to the minimum necessary value, according to the database owner’s belief, which provides the user with “enough” choices within the returned tuples. While such a strategy might suffice the simplest use-cases, it often cannot satisfy users with specific needs and also prevents many interesting third-party services from being developed over web databases - e.g.,

- Consider a third-party service which enables a user to filter query results according to attributes that cannot be specified in the original form-like interface. For example, American Airline’s (AA) flight search-by-schedule¹, a top-10 interface, does not allow a user to specify filtering conditions such as finding the top-10 flights with in-flight wifi. If a third-party service wants to provide such a feature, it must somehow “bypass” the *top- k constraint* because otherwise one might not be able to find enough (or any) wifi-equipped flights from the top-10 results.
- Consider a web aggregator or a web mashup which joins tuples from multiple hidden web databases and returns the joined results - e.g., a mashup joining Orbitz.com (a hotel booking website) with Tripadvisor.com (a hotel review website) to return the top- k cheapest hotels that have an average review of at least 4 stars. Once again, such a mashup must somehow break the *top- k constraint* because not enough matching tuples may be discovered from the mere k tuples returned by each web database.

To enable these third-party services and many other interesting applications (e.g., data analytics) that are currently disabled/handicapped by the *top- k constraint*, a trivial solution is for the third-party service provider to negotiate a private agreement with each web database owner in order to establish data-access channels beyond the *top- k web interface*. Nonetheless, such negotiations are difficult even between large organizations³ due to revenue sharing, security and myriad of other thorny issues - thus making the solution not scalable to a large number of web databases. As such, our focus in this paper is to develop automated third-party algorithms that only use the public interfaces of web databases without requiring any additional cooperation from the database owners.

*The work of Saravanan Thirumuruganathan and Gautam Das is partially supported by NSF grants 0812601, 0915834, 1018865, a NHARP grant from the Texas Higher Education Coordinating Board and grants from Microsoft Research and Nokia Research. The work of Nan Zhang is partially supported by NSF grants 0852674, 0915834, 1117297 and a GWU Research Enhancement Fund.

¹<http://www.aa.com/reservation/searchFlightsSubmit.do> By default $k = 10$. A user may configure k to be as large as 50. No page down is allowed.

²<http://www.amazon.com/Best-Sellers/zgbs>

³<http://online.wsj.com/article/SB121755825030403467.html>

Another seemingly straightforward solution to the above problems is *crawling* - i.e., the retrieval of *all* tuples in a hidden web database by issuing multiple queries through its web interface [1], [2]. Once all tuples are downloaded, they can be treated as a local database to support all of the above applications. Nonetheless, a key pitfall of this solution is its prohibitively high query cost (i.e., the numerous search queries one needs to crawl all tuples from a web database) - which can be simply infeasible for real-world web databases which often impose a per user/IP limit on number of queries one can issue over a given time frame (e.g., Google Search API allows only 100 free queries per user per day).

B. A Novel Problem: Breaking the Top- k Barrier

Given the pitfalls of crawling, we propose to study in this paper a novel problem of digging deeper into a web database to retrieve (more than k) top-ranked tuples which satisfy a user-specified search query - and thereby “breaking” the top- k barrier. Specifically, we consider the following fundamental operator:

GETNEXT: Given the top- h tuples ($h \geq k$) satisfying a user-specified query, retrieve the next-highest-ranked (i.e., No. $(h + 1)$) tuple from the hidden web database by issuing search queries through its public interface, without any knowledge of its ranking function.

One can see that, by calling GETNEXT iteratively, it is possible to retrieve as many top-ranked tuples as necessary for a user-specified query - thereby enabling both sample applications discussed above without the need of crawling all tuples from the database. Because of the query-number limitations enforced by web databases, an important objective in the design of GETNEXT is to maintain a small *query cost* - a goal shared by most existing studies on exploring hidden web databases (e.g., [3]–[5]).

It is important to understand that the most critical efficiency measure here is the *number of queries* the algorithm requires to issue through the web interface of the hidden database, *not* the actual processing time of issued queries and/or the local processing overhead. To understand why, note that our objective is to build a scalable third-party service that breaks the top- k barrier for a large number of users. A key obstacle for building such a service, as mentioned earlier, is the query allowance enforced by the hidden database owner - which completely shuts down our access to the database if we issue more queries than a pre-determined threshold. Compared with this hard cutoff, the variance of local processing overhead and/or query processing time is relatively negligible - especially given the fact that most well-designed hidden databases return query answers within a short period of time.

C. Outline of Technical Results

To design GETNEXT, the technical challenge may have subtle differences across various web databases, mainly because of the different ranking functions being used. At one extreme, some websites allow users to choose their own ranking function (from a predetermined set) - e.g., airlines websites allow

users to sort by attributes such as by price, departure time, etc. At the other extreme, a website might feature a complex and proprietary *query-specific* ranking function (e.g., “relevance” of a tuple to a query) that may never be deterministically inferred from other query answers. Other possible ranking functions include a global order that is nevertheless hidden from the input interface - e.g., Amazon uses popularity as the default ranking function but does not allow it to be specified in a search query. For most of the paper, we focus on the case where the ranking function is a query-independent global order of all tuples. The implications of other ranking-function variations on our solutions are discussed separately.

There are two key components of our proposed solution to GETNEXT: *candidate generation* and *candidate testing*.

Candidate Generation: Given the top- h tuples, the candidate generation step aims to identify a complete yet small set of tuples that can potentially have the rank $h + 1$. A key observation here is that the problem is equivalent to finding a small set of queries, each of which matches fewer than k tuples in the top- h , while together cover the rest of the database. One can see that, since each query in the set returns at least one non-top- h tuples, the No. $(h + 1)$ tuple must be returned by at least one query in the set. Based on this key observation, we propose a tuple-chain-construction based technique which further reduces the query cost required for candidate generation significantly.

Candidate Testing: Since the task is now reduced to testing which candidate is the No. $(h + 1)$ tuple, the key enabling question becomes how to perform pairwise rank-comparison between two tuples. Interestingly, for certain pairs of tuples, the comparison may be done with a single query to the hidden database. Specifically, consider issuing the most specific query that matches both tuples. If both are returned, then the result reveals their order. If only one is returned, then it must have a higher rank. The challenge, however, is in the worst-case scenario where neither is returned. In the paper, we start by resolving this scenario with a baseline approach that requires 2^m queries, where m is the number of attributes. Then, we propose two ideas - one connects with the well-studied problem of *minimal infrequent itemsets* mining [6], and the other is a heuristic of query-result inference - which significantly reduce the query cost for candidate testing.

D. Summary of Contributions

In summary, the main contributions of this paper are:

- We introduce the novel problem of breaking the top- k barrier of a hidden web database to retrieve top ranked tuples that match a user query. We consider several variants of the problem, and study necessary and sufficient conditions under which this problem can be solved.
- We propose BEYOND- h -GETNEXT and ORDERED-GETNEXT, two algorithms that iteratively uses the two fundamental operations, candidate generation and candidate testing, to retrieve the next-highest-ranked tuple. While BEYOND- h -GETNEXT guarantees the correct

retrieval of next ranked tuple⁴, ORDERED-GETNEXT further uses an effective heuristic of query-result inference to significantly reduce the query cost in practice without sacrificing correctness.

- Our contributions also include a careful theoretical analysis of BEYOND- h -GETNEXT and ORDERED-GETNEXT, as well as a thorough experimental evaluation over both synthetic datasets and real-world websites.

The rest of the paper is organized as follows. In §2, we discuss preliminaries - e.g., the models of hidden web databases and their ranking functions. §3 defines the problem of breaking the top- k barrier and outlines our proposed solution that uses GETNEXT. §4 and §5 detail the two main parts of our algorithm, candidate generation and candidate testing, respectively. In §6, we discuss extensions to the algorithms to handle special cases. §7 describes a detailed set of experiments over real-world datasets. §8 discusses related work, followed by the conclusion in §9.

II. PRELIMINARIES

In this section, we introduce a model for hidden databases and describe the different types of ranking functions used commonly in hidden databases.

A. Model of Hidden Databases

Consider a hidden database D with n tuples and m input attributes A_1, A_2, \dots, A_m . Given a tuple t and attribute A_i , let $t[A_i]$ be the value of A_i in t . Let $Dom(A_i)$ be the domain of A_i . For the purpose of this paper, we restrict our attention to categorical attributes and assume the appropriate discretization of numeric ones. We also consider all tuples distinct and without null values. Let $f(\cdot)$ be the ranking function which takes a tuple and a query as input and outputs an integer between 1 and n . Without loss of generality, we assume the output of $f(\cdot)$ to be unique for each tuple.

A user can query the system by specifying the desired values for a subset of A_1, \dots, A_m . Thus, a user query q is of the form `SELECT * FROM D WHERE $A_{i_1} = v_{i_1} \& \dots \& A_{i_s} = v_{i_s}$` , where $\{i_1, \dots, i_s\} \subseteq [1, m]$, and $v_{i_j} \in Dom(A_{i_j})$. The set of tuples matching query q is denoted as $Sel(q)$. If $|Sel(q)| > k$, an *overflow* occurs and only the top- k results are returned, along an overflow flag indicating that more tuples matching the query cannot be returned. If $|Sel(q)| = 0$, then an *underflow* occurs as no tuples match the query. Otherwise, i.e., when $|Sel(q)| \in [1, k]$, we say that q is valid. For the purpose of this paper, we make the realistic assumption that $k > 1$.

For the purpose of our paper, we assume that the interface only displays the top- k results and does not allow users to extract additional results by scrolling through the results. The only way to get additional results is to reformulate the input query. This is a reasonable assumption as many real world hidden web databases such as Yahoo! Autos limit the maximum number of page turns a user can perform.

TABLE I
INPUT TABLE, D

	A_1	A_2	A_3	A_4	A_5
t_1	0	0	0	0	1
t_2	0	0	0	1	1
t_3	0	0	1	0	1
t_4	0	1	1	1	1
t_5	1	1	1	0	1
t_6	1	1	1	1	1
t_7	0	0	0	0	0

B. Model of Ranking Function

There are two broad categories of ranking functions: *static* and *query-dependent*.

- A ranking function $f(\cdot)$ is *static* if for a given tuple t , $f(q, t)$ is constant for all queries q - i.e., the rank of a tuple is independent of the query being issued. An example in practice is the “sort by price” used by Yahoo! Autos. Note that the input tuple may feature not only A_1, \dots, A_m but also the non-input-specifiable attributes (e.g., “popularity” as discussed in §1).
- A ranking function is *query-dependent* if, for a given t , $f(q, t)$ varies for different queries q . An example of such a ranking function occurs in a fuzzy-matching scenario where all tuples are ordered according to the number of attribute matches between the query and each tuple.

As discussed in §1, we focus on static ranking functions in this paper. The reason for doing so is simple - if the ranking function is query dependent, no mechanism can be used to fetch the next ranked tuple. To understand why, note that in order to get tuples beyond top- k , it is necessary to reformulate the query. But this has the side effect of arbitrarily changing the ranking of tuples. Hence, with a query-dependent ranking function, no mechanism can guarantee the discovery of tuples with rank greater than k for a given query.

For the purpose of this paper, we conservatively assume that the ranking function is unbeknown to our algorithm. If the ranking function is known and is based on the attributes returned by the hidden web interface (such as sort by price), it is possible to leverage this information to design algorithms with significantly less query cost. We further discuss this variant in §5. In addition, we assume that it is possible to infer a unique global order of the top-ranked tuples to be extracted from the web interface. If such an order cannot be inferred from the interface, one of the possible partial orders would be returned, as we shall explain in §6.

Running Example: Table I shows a simple table which we shall use as running example throughout this paper. There are $m = 5$ Boolean attributes and $n = 7$ tuples which are ranked in the order given in the table. i.e., t_1 is the highest ranked tuple.

III. OVERVIEW OF GETNEXT

In this section, we first discuss the technical challenges of GETNEXT, and then outline the structure of our proposed two-

⁴if such an order can be uniquely determined from the top- k interface.

step solution - the details of each step shall then be developed in the next two sections, respectively.

A. Technical Challenges

To illustrate the main technical challenges, we consider a fundamental question: Given two tuples t and t' , how can we determine which one ranks higher? We start with a straightforward comparison - i.e., when t and t' match the same query q which returns at least one of the two tuples:

- if q returns t but not t' , then t is ranked higher,
- if q returns t' but not t , then t' is ranked higher, or
- if q returns both, then we can make the comparison based on the returned order.

In this case, we call two tuples *directly comparable*, with the higher-ranked tuple *directly dominating* the other one - i.e.,

Definition 1: [Domination] A tuple t is said to **directly dominate** another tuple t' , i.e., $t \succ t'$, if and only if t and t' are directly comparable and t ranks higher than t' .

A tuple can dominate another tuple directly or indirectly. Suppose tuple $t \succ u$ and $u \succ v$. Even if t and v are not directly comparable, we can infer that t indirectly dominates v . By default, we use the term domination to refer to direct domination.

For example, consider the running example with a top-2 interface. We can observe that t_1 and t_3 are directly comparable using the query q_1 : `SELECT * FROM D WHERE $A_1 = 0$ AND $A_2 = 0$ AND $A_4 = 0$ AND $A_5 = 1$` with t_1 ranked higher than t_3 . Similarly, tuples t_2 and t_3 are directly comparable using the query q_2 : `SELECT * FROM D WHERE $A_1 = 0$ AND $A_2 = 0$ AND $A_5 = 1$` . The result includes t_2 but not t_3 - i.e., t_2 ranks higher.

A key observation here is that if two tuples are directly comparable, then we need only one query to determine their domination relationship: the *most specific* query which matches both tuples - i.e., the query which contains one predicate for each attribute on which both tuples share the same value. To understand why, note that if this query cannot return at least one of the two tuples, then no other query can - i.e., the two tuples are not directly comparable. For the running example, both q_1 and q_2 shown above are the most specific queries matching the two corresponding tuples.

While the possibility of direct comparison shows promises for ranking tuples in the database, it also illustrates the key technical challenge for GETNEXT: not every pair of tuples are directly comparable with each other - e.g., neither t_6 nor t_7 in the running example can be returned by the most specific query that matches both of them (i.e., `SELECT *`).

In this case, the comparison of the two tuples requires one to identify a “bridge” of tuples between them - e.g., $t \succ t_x \succ t'$ for comparing t with t' . The problem, however, is it is unclear how one can find the bridging tuples without actually crawling all tuples from the database and incurring a prohibitively high query cost. In the next subsection, we outline the structure of our proposed solution to address this challenge.

B. Outline of Our Proposed Solution

Our proposed solution for GETNEXT is a two-step process:

- **Candidate Generation:** In this step, we identify a small set of *candidate tuples* which are guaranteed to contain the No. $h + 1$ tuple. If the output set has a size of 1, then we can directly output the No. $h + 1$ tuple. Otherwise, we call the following candidate testing step. §4 describes our design for candidate generation.
- **Candidate Testing:** In this step, we take the set of candidate tuples as input and compare between them to determine which tuple is indeed the No. $h + 1$. §5 describes our design for candidate testing.

IV. CANDIDATE GENERATION

Given the current set of top ranked tuples, the candidate generation step is supposed to produce a set of candidate tuples, one of which is guaranteed to be the next ranked tuple in the database. In subsection §6-3, we discuss finding the next ranked tuple matching a selectivity constraint. The determination of the exact next-ranked tuple from the candidate set is done using the candidate testing oracle described in §5. In this section, we first describe a baseline approach for candidate generation, and then introduce a more efficient algorithm using a notion of directed acyclic graphs (DAG) of tuples. The DAG based algorithm exploits the ordering information provided by query answers to potentially complete multiple rounds of candidate generation in a single iteration (i.e., it may answer multiple consecutive GETNEXT calls without additional query cost). Recall from §2 that we make the realistic assumption of $k > 1$.

A. Baseline Approach

The essence of candidate generation can be stated as follows. Given the top- h tuples, candidate generation needs to identify a set of queries that is guaranteed to “cover” (i.e., return) the next-ranked (i.e., No. $(h + 1)$) tuple. One can see that such a set of queries must together match all possible tuples in the database - in order to ensure that no other tuple has a higher rank than the next-ranked tuple being covered.

We start by considering a simple baseline approach as follows: First, find a set of attributes \mathcal{A} such that if we partition the top- h tuples based on their value combinations for attributes in \mathcal{A} , then each partition contains fewer than k elements. Since each tuple is unique, such an \mathcal{A} already exists. After finding $\mathcal{A} = \{A_{i_1}, \dots, A_{i_j}\}$, we construct queries of the form q_i : `SELECT * FROM D WHERE $A_{i_1} = v_{i_1}$ AND \dots AND $A_{i_j} = v_{i_j}$` for all possible value combinations of $v_{i_1} \in Dom(A_{i_1}), \dots, v_{i_j} \in Dom(A_{i_j})$, and execute all such queries. One can see that these queries completely cover the database domain and thus return a candidate set for the No. $(h + 1)$ tuple. To understand why, note that the No. $(h + 1)$ tuple must be returned by one of the queries issued, because otherwise the query which matches the No. $(h + 1)$ tuple must return a tuple that directly dominates the No. $(h + 1)$ tuple.

Example 1: Given the top-3 tuples in the running example, suppose we want to retrieve the next ranked tuple. We identify




Fig. 1. DAG used in Examples 1 and 2

an attribute, say A_3 (or A_4), such that the number of tuples having the values 0 and 1 are less than $k = 3$. We execute two queries by augmenting q - specifically, q_1 : `SELECT * FROM D WHERE $A_3 = 0$` returns new tuples $\{t_7\}$ and q_2 : `SELECT * FROM D WHERE $A_3 = 1$` returns new tuples $\{t_4, t_5\}$. The candidate set for 4-th ranked tuple is the set $\{t_4, t_5, t_7\}$. If we want to retrieve the 5-th ranked tuple, we can choose any of the attributes A_2, A_3 or A_4 to partition the top-4 tuples.

Analysis: The number of queries executed to identify the candidate set depends on the domain of the attribute(s) selected. Given an attribute set \mathcal{A} , the number of queries executed is $\prod_{A \in \mathcal{A}} |Dom(A)|$.

B. DAG based Approach

In this subsection, we develop a DAG-based algorithm which leverages the order information provided in the query results to further reduce the number of returned candidate tuples, and to identify the candidate sets for multiple next-ranked tuples at a single iteration. In other words, our DAG based approach retrieves the candidate sets for as many next ranked tuples as possible so that subsequent *GetNext* do not incur any additional query cost.

The data structure used in our approach is a directed acyclic graph (DAG) called the *dominance directed graph*. Each node in the DAG correspond to a tuple and a directed edge exist from node u to node v if u dominates v . Given the result of any query q , we can form an DAG from it results. If the query returned $|q|$ tuples, then the DAG would have at most $\binom{|q|}{2}$ edges and an linear chain of $|q|$ tuples as a subgraph. An example of the DAG formed from queries q_1 and q_2 from Example 1 is in Figure 1. Given a set of queries q_i , we can form a set of linear chains from their results. Let S_i denote the i -th linear chain and S be the set of all linear chains. The notation $head(S_i)$ returns the tuple with highest rank in S_i while $head(S)$ returns the set of highest ranked tuples in each chain.

The primary aim of this approach is to identify a linear of chain of consecutively ranked tuples, if any. If such a chain exists, then the tuples from the chain can be returned for the subsequent *GetNext* calls without additional query cost. We use two observations to extract this chain. First, the only tuples that can dominate the candidates for t_{h+1} are the ones in the top- h . Second, since the database has a fixed (but hidden) global order of all tuples, there always exists a dominance relationship (i.e., direct comparison) between the tuples with rank h and $h + 1$. If not, the ranks of these two tuples can be

flipped without violating any other relative rankings.

To see how these observations are useful, consider the augmented queries from the baseline approach. Each such query q_i results in a linear chain S_i . We can see that $head(S_i)$ dominates other tuples from S_i . Hence, $head(S_i)$ is the only tuple from S_i that needs to be added to candidate set. Since tuples t_h and t_{h+1} must be directly comparable, we need to consider only the head of each linear chain and compare it with tuple t_h .

The overview of the algorithm is as follows. We have a list of linear chains (from augmented queries of prior *GetNext* invocations) and the linear chain, say S_i , from which tuple t_h was extracted. We perform pairwise comparison between tuples from different linear chains. An edge is added from node u to node v , if they are directly comparable and u ranks higher than v . Then we compare the tuple t_h with the head of each chain *except* S_i . If *none* of the heads are directly comparable with t_h , then we can assign $head(S_i)$ to be the next ranked tuple without even performing candidate testing. This is possible due to the fact that consecutively ranked tuples are always comparable. If some of them are comparable with t_h , only these form the candidate set for t_{h+1} . The candidate tuples are then compared pairwise with each other to identify non dominated tuples. The domination can be either direct or indirect. It is easy to see that tuple t_{h+1} is guaranteed to be among the non dominated tuples that are also comparable to tuple t_h .

If there are multiple candidate tuples for t_{h+1} , then the candidate testing oracle must be invoked. If not, we are guaranteed that the only candidate tuple must have rank $h + 1$. The candidate tuple is then removed from its linear chain and the process is continued till the number of candidates for the next ranked tuple is more than 1. This can potentially result in multiple consecutive next ranked tuples to be retrieved.

Example 2: Consider the same setting as Example 1. We wish to extract 4-th ranked tuple from a top-3 interface. Using attribute A_3 , we construct two augmented queries q_1 and q_2 resulting in two linear chains S_1 and S_2 . The last tuple t_3 belonged to linear chain S_2 . The resulting DAG can be seen from Figure 1. Both the tuples t_7 and t_4 are comparable with t_3 and do not dominate each other. However, t_7 is indirectly dominated by t_4 through t_5 . Hence we can immediately declare t_4 as the 4-th ranked tuple. Since t_5 also dominates t_7 , it is identified as the 5-th ranked tuple. Note that in both the cases, no calls were made to the candidate testing section. Additionally, we identified two consecutively ranked tuples in a single invocation of *GetNext*.

Analysis : At each iteration, let the number of linear chains be l . The query cost for pairwise comparison of tuples between chains is $\prod_{i=1}^l |S_i|$. We also require an addition l queries to compare tuple t_h with the heads of each chain. Thus, the algorithm requires at most $l + \prod_{i=1}^l |S_i|$ in any iteration. Note that subsequent iterations do need any additional queries till one of the chains is completely consumed as the comparison information between tuples has already been identified.

V. CANDIDATE TESTING

In this section, we consider the candidate testing problem - i.e., based on prior knowledge of the top- h ranked tuples t_1, \dots, t_h , what queries does one need to issue to the hidden database in order to test whether a given tuple t has rank $h+1$? We start with two baseline approaches which can require prohibitively high query costs in practice, and then present our two ideas for improving their efficiency: (1) a reduction to beyond- h minimal queries - which significantly reduces both worst- and average-case query costs, and (2) a heuristic query ordering - which further reduces the query cost in practice. It must be noted that if the ranking function is known and based on the attributes returned by the hidden database (e.g. sort by price), then the next ranked tuple can be directly identified from the candidate tuples without an explicit candidate testing phase or querying the hidden database for comparison.

A. Baseline Approaches

To prove that t indeed has rank $h+1$, we have to ensure that no tuple in the database, other than the top- h ones, dominates t . A seemingly straightforward baseline approach is then to first crawl all other tuples from the database, and then compare each of them with t to identify any dominance relationship. The problem with this approach, however, is that the crawling step requires at least n/k queries - where n is the number of tuples in the database and k is as in the top- k interface - because each query returns at most k tuples. Most common hidden web databases routinely have hundreds of thousands of tuples with a relatively small value of k , resulting in a prohibitive query cost to test a single tuple.

We now consider another baseline which is enabled by the following observation: according to the definition of dominance relationship shown in §3, the only queries which may “reveal” a tuple dominating t are those that actually match t - i.e., queries of the form

$$q : \text{SELECT } * \text{ FROM } D \text{ WHERE } A_{i_1} = t[A_{i_1}] \text{ AND } \dots \\ \text{AND } A_{i_r} = t[A_{i_r}] \quad (1)$$

where $\{i_1, \dots, i_r\} \subseteq \{1, \dots, m\}$ (recall that m is the number of attributes). Specifically, t has rank $h+1$ if and only if every query of the form (1) either returns t as the highest-ranked non-top- h tuple, or returns only tuples in the top- h .

Thus, our second baseline is to issue all queries matching t . One can see that the query cost for the second baseline is $\binom{m}{0} + \dots + \binom{m}{m} = 2^m$. While this number is often much smaller than n/k for a practical hidden database (because there are usually only a few, e.g., 5 or 10, attributes that can be specified on the input web interface), issuing 2^m queries for each candidate tuple may still lead to an extremely high query cost. In the following two subsections, we develop our two ideas for reducing query cost respectively.

B. Beyond- h Minimal Queries

Our first idea is to reduce the space of queries required for rank testing from all queries which match t (i.e., of the form

in (1)) to a much smaller subset which we refer to as the *beyond- h minimal queries*. In the following, we first define beyond- h minimal queries and show the completeness of such queries - i.e., issuing them suffices for rank testing. Then, we describe a (somewhat surprising) mapping of beyond- h minimal queries to finding minimal infrequent itemsets - a problem that has been extensively studied in the database and data mining communities (e.g., see survey in [7]). Finally, we leverage the existing results on minimal infrequent itemsets to derive an upper bound on the number of beyond- h queries.

Definition and Completeness: For any query q which matches t , we use $S(q)$ to represent the *companion attribute set* of the query - i.e., the set of attributes involved in the query. For example, $S(q) = \{A_{i_1}, \dots, A_{i_r}\}$ for q in (1). Then, we call q a *beyond- h minimal query* if and only if it satisfies both of the following two conditions:

- q must return at least one non-top- h tuples - i.e., q must match fewer than k tuples in t_1, \dots, t_h
- any query q' which matches t and has $S(q') \subset S(q)$ must only return top- h tuples - i.e., q' must match at least k tuples in t_1, \dots, t_h .

One can see from the definition that, as the name suggests, q is a “minimal” query which returns any tuple beyond the top- h . We now explain why issuing only beyond- h minimal queries suffices for rank testing. Consider the testing of whether t is the tuple with rank $h+1$. A key observation here is that any query q_0 which matches t but is not a beyond- h minimal query must satisfy one of the following two conditions:

- If q_0 matches at least k tuples in t_1, \dots, t_h , then one can already infer the answer to q_0 from the knowledge of t_1, \dots, t_h - i.e., q_0 is useless for rank testing.
- If q_0 matches fewer than k tuples in t_1, \dots, t_h but is not a beyond- h minimal query, then there must exist a beyond- h minimal query q'_0 such that $S(q'_0) \subset S(q_0)$. If q'_0 returns t as the top-ranked tuple besides top- h , then we are already certain that no non-top- h tuple matching q_0 can outrank t . Otherwise, we are already certain that t cannot have rank $h+1$ - i.e., in either case, we do not need to issue q_0 .

Example : Considering the running example from Table I, we can see that $A_3 = 1$ and $A_4 = 1$ are two examples of beyond- h queries for t_4 .

Mapping: We now show that the problem of finding all beyond- h minimal queries is equivalent to finding all minimal infrequent itemsets over a transactional database. To understand why, consider the following procedure which maps the top- h tuples to h transactions. We first map each attribute A_j ($j \in [1, m]$) to an item s_j . Then, for each tuple t_i ($i \in [1, h]$), we map it to a transaction r_i by including in r_i all items corresponding to the attributes on which t_i and the testing tuple t share the same value - i.e.,

$$r_i = \{s_j | t_i[A_j] = t[A_j]\}. \quad (2)$$

We can see that, with this mapping, the companion attribute set of each beyond- h minimal query q , i.e., $S(q)$, becomes a

minimal infrequent itemset over the h transactions, with the frequency threshold being k/h . This observation can be readily made from the definition of beyond- h minimal queries: Since such a query must match fewer than k tuples in t_1, \dots, t_h , $S(q)$ is infrequent given the threshold of k/h . Since no subset of $S(q)$ can match fewer than k tuples in top- h , $S(q)$ must be minimally infrequent. One can see that the inverse also holds - i.e., there is a one-one mapping between $S(q)$ and a minimal infrequent itemset.

Example : Suppose we have extracted the top three tuples and want to determine if tuple t_4 is indeed the 4-th ranked tuple. We first map tuples t_1, t_2, t_3 to transactions as $r_1 = \{A_1 = 0, A_5 = 1\}$, $r_2 = \{A_1 = 0, A_4 = 1, A_5 = 1\}$ and $r_3 = \{A_1 = 0, A_3 = 1, A_5 = 1\}$. The threshold is $\frac{3}{3} = 1$. The infrequent itemsets are $A_3 = 1$ and $A_4 = 1$ which correspond to beyond- h queries for t_4 . Also, the number of beyond- h queries is dramatically smaller than the 2^5 queries needed in the previous approach.

While (as we shall show below) the mapping enables us to derive an upper bound on the number of beyond- h minimal queries, we would like to remark here two major differences between our problem and the traditional problem of finding minimal infrequent itemsets.

First, even though finding all minimal infrequent itemsets is known to be #P-complete, the time complexity is *not* really a concern for our problem because our input size m - i.e., the number of attributes - is usually much smaller than the number of items in a transactional database. As such, we could simply enumerate all 2^m possible itemsets (and find the minimal infrequent ones) without causing significant overhead. What is a major concern for us, however, is the number of minimal infrequent itemsets because it translates to the number of queries we have to issue through the web interface - a costly and time-consuming process.

Second, our frequency threshold, i.e., k/h , is generally much larger than the threshold traditionally considered for minimal infrequent itemsets. As we mentioned in §1, even an $h = 2k$ may bear significant interest as third-party analyzers are most likely interested in those highly ranked, albeit outside top- k , tuples. As we shall show below, this unusually high threshold enables us to improve the upper bound on the number of beyond- h minimal queries when h is small.

Upper Bound: First, according to the existing results on the number of minimal infrequent itemsets, that the number of beyond- h minimal queries can be bounded by $\binom{m}{m/2}$. We now show that when h is small, specifically $h \leq m/2 + k - 1$, the number of beyond- h minimal query q has another upper bound of $\binom{m}{h-k+1}$.

An important observation here is that the number of predicates in a beyond- h minimal query, say q , is at most $h - k + 1$. To understand why, consider a query-construction process in which we start with the SELECT * query, and then gradually add into it one conjunctive predicate in q (i.e., one attribute in $S(q)$) at a time, until the query matches fewer than k tuples in the top- h . One can see that each predicate being added, say

$A_i = t[A_i]$, must remove at least one top- h tuple from the set of tuples matching the previous query, because otherwise one can always remove $A_i = t[A_i]$ from q without changing the answer to q - contradicting the fact that q is beyond- h minimal. As such, once $h - k + 1$ predicates are added to the query, the number of top- h tuples matching the query must drop to below k - i.e., $S(q)$ contains at most $h - k + 1$ attributes. Again, since all beyond- h minimal queries forms an anti-chain, the number of them is at most $\binom{m}{h-k+1}$ when each beyond- h minimal query contains at most $h - k + 1$ predicates and $h - k + 1 \leq m/2$.

In summary, we have the following theorem:

Theorem 1: Given the top- h tuples, the maximum number of queries one needs to issue for testing whether a tuple has rank $h + 1$ over a database of m attributes and n tuples, $c(n, m, h + 1)$, satisfies

$$c(n, m, h + 1) \leq \binom{m}{\min(h - k + 1, m/2)}. \quad (3)$$

C. Query Ordering

Our next idea to reduce query cost that works very well in practical hidden databases is a heuristic - query ordering. Recall that *beyond- h query* is a minimal query that returns at least one non-top- h tuple. Given a candidate tuple t , if all its corresponding beyond- h minimal queries returns t as the highest ranked non-top- h tuple, then we can conclude that no other tuple dominates t and hence t has rank $h + 1$. Note that to make this conclusion, it is mandatory to execute all the beyond- h queries.

The key idea in query ordering is that of *elimination*. If we can eliminate all but one tuple from the candidate set, then the remaining tuple has to be the next ranked tuple and we can make that conclusion even without executing any of the beyond- h queries for it. This is due to the fact that the candidate generation step produces a set of tuples one of which is guaranteed to be in the next ranked tuple. The query ordering heuristic takes the idea a little further.

Given a candidate tuple t and one of its beyond- h queries q , there are two possible results : (1) t is the top ranked non-top- h tuple (2) t is *not* the top ranked non-top- h tuple. In the first case, the query q did not give any contradicting evidence for t and the next beyond- h query needs to be executed. On the other hand, the second outcome provides an evidence that disqualifies t from being the next ranked tuple. i.e. the procedure for testing t can be terminated early. The heuristic tries to reorder the execution of beyond- h queries so that if t is not the No. $h + 1$ ranked tuple, it is detected earlier.

While reordering the queries of a single candidate tuple is useful by itself, the maximum advantage is obtained when the set of beyond- h queries of *all* the tuples in candidate set are reordered. By ordering queries based on the chance that it eliminates atleast one candidate tuple and executing them in that order, we eliminate as many candidates as possible in the least number of queries. Furthermore, while executing the queries, any candidate tuple dominate by others can be immediately rejected.

The heuristic relies on two factors that make a beyond- h query q useful. Note that both the factors implicitly favor shorter queries over longer ones.

- The number of tuples in *candidate set* matched by q . If q matches l tuples in candidate set, we can immediately eliminate the $l - 1$ dominated candidates after executing q as they cannot have rank $h + 1$.
- The *expected* number of tuples in the *database* that is matched by q . If q matches a large fraction of database, then there is a high likelihood that one of such tuples will be ranked higher than candidate tuple t . Of course, since the entire database is not available to us, we estimate the fraction by assuming a random database where the attribute values are uniformly distributed. While this assumption does not always hold, it serves as a useful approximation and heuristic. Given a boolean database with 5 attributes and any query with two attributes can be expected to match 25% of the tuples.

In summary, the query ordering heuristic pools the beyond- h queries of all candidate tuples and reorders them based on a weighted combination of the two factors described above. The weights can be determined using domain knowledge of the hidden database. The queries are executed in the order so as to eliminate the candidate tuples as early as possible. Any candidate tuple dominated by a non top- h tuple or other candidate tuple are eliminated. The process is continued till only one candidate remains.

Example : Suppose we wanted to determine if t_3 or t_4 is the third ranked tuple. $A_3 = 1$ and $A_4 = 1$ are two of the beyond- h queries for t_4 while the corresponding ones for t_3 are $A_3 = 1$ and $A_4 = 0$. Since the query $A_3 = 1$ matches both t_3 and t_4 , it is executed before either of $A_4 = 1$ or $A_4 = 0$. After executing $A_3 = 1$, we note that t_3 is ranked higher than t_4 in the result and hence declare it as the 3-rd ranked tuple.

Analysis : The query cost of heuristic is bounded by the upper bound for the number of beyond- h queries for the tuples in candidate set. In the worst case, this procedure degenerates to executing all the beyond- h queries for *all but one* of the candidate tuples.

VI. ALGORITHM DESIGN AND EXTENSIONS

In this section, we integrate the candidate generation and testing techniques discussed in previous two sections to develop our final algorithms for GETNEXT. In addition, we shall describe different extensions of our algorithms such as retrieving the top ranked tuples when no unique total order exists among them or retrieving top ranked tuples that satisfy additional user specified filters.

A. Algorithms BEYOND- h -GETNEXT and ORDERED-GETNEXT

We start by integrating our DAG-based candidate generation algorithm with the beyond- h queries based candidate testing algorithm to develop the BEYOND- h -GETNEXT algorithm. To be the next ranked tuple, any candidate tuple must be the top ranked non top- h tuple for each of its beyond- h

queries. Algorithm 1 depicts the pseudocode of BEYOND- h -GETNEXT.

Algorithm 1 BEYOND- h -GETNEXT

- 1: **Input parameters :** $topH$, the set of top ranked tuples
 - 2: Get candidates for t_{h+1} using candidate generation
 - 3: **for** each candidate tuple t **do**
 - 4: Generate and execute beyond- h -queries for t
 - 5: If any tuple other than top- h tuples dominate t , reject t
 - 6: **end for**
 - 7: **return** unrejected tuple as t_{h+1}
-

We also integrate our candidate generation algorithm with the heuristic candidate testing algorithm to develop the ORDERED-GETNEXT algorithm. The only difference between ORDERED-GETNEXT and BEYOND- h -GETNEXT is in the rank testing phase. In ORDERED-GETNEXT, we first identify the beyond- h queries for *all* candidate tuples and order them based on their likelihood of rejecting a candidate tuple. The queries are executed until all but one candidate tuples have been rejected. The remaining tuple is declared as No. $h + 1$ tuple. Algorithm 2 depicts the pseudocode of ORDERED-GETNEXT.

Algorithm 2 ORDERED-GETNEXT

- 1: **Input parameters :** $topH$, the set of top ranked tuples
 - 2: Get candidates for t_{h+1} using candidate generation
 - 3: Collect the beyond- h queries of all candidates and order them based on likelihood to reject candidates
 - 4: **for** each query **do**
 - 5: Execute query
 - 6: Reject any candidate tuple dominated by other candidates or a non top- h tuple
 - 7: If only one candidate left, break
 - 8: **end for**
 - 9: **return** remaining tuple as t_{h+1}
-

B. Absence of Total Order within Top Ranked Tuples

One of the assumptions that was made by the algorithms was that the set of top ranked tuples that we wish to retrieve are totally ordered and the order is inferable from the hidden database interface. Specifically, we assumed that tuples t_h and t_{h+1} was directly comparable. In this subsection, we discuss how to handle the different scenarios when the assumption does not hold.

Two tuples can be compared with each other either directly or indirectly and similarly the dominance relationship can be established directly or indirectly through other intermediate tuples. For example, we might have two tuples t and v that are not directly comparable. However, if $t \succ u$ and $u \succ v$, then we can indirectly infer their dominance relationship. If two tuples are not comparable at all, even indirectly, then their dominance relationship cannot be established. Choosing either of the tuples to be the next ranked tuple results in a potentially

valid total ordering from the limited information available. The possibility of two tuples not comparable affects both the candidate generation and testing steps.

Candidate Generation: In candidate generation, if the head tuple of *every* linear chain was not comparable to t_h , then we cannot assign the head tuple from the linear chain from which t_h was extracted to be the next ranked tuple (as it is not comparable to t_h). All the non dominated candidate tuples are sent to candidate testing for identifying the next ranked tuple.

Candidate Testing: If multiple tuples from candidate set are not dominated by any other tuple other than the ones in top- h (including other tuples in candidate set), then each of them can potentially be considered as the next ranked tuple. Hence, one of the non dominated candidate tuples is selected uniformly at random as the next ranked tuple and the process is continued. This random selection creates one of the valid partial order of the top ranked tuples. Since the output total order is no longer accurate, a metric must be chosen to measure the distance between the actual total order and the partial order. The accuracy measure used is the expected distance between a randomly generated total order and the actual total order. The distance between two ranked list can be computed using Kendall τ or the Spearman's footrule.

C. Top Ranked Tuples with Selectivity Constraints

The discussions in the previous sections described techniques to retrieve the top ranked tuples from the entire database. An equally important and practical scenario is one where the user is interested in the top ranked tuples over a subset of the database. For example, the user might be interested in the cheapest flights with in-flight wifi. An alternate perspective is to view the problem as retrieving top ranked tuples where some of the attribute values are already preset by the user, for e.g. wifi. The specified attributes then partition the entire hidden database into two partitions - one which matches the specified attributes and another which does not match the specified attributes. In this subsection, we discuss how to extend the techniques discussed so far to solve this problem.

An initial approach one might come up with is to keep retrieving top tuples from the *entire* database incrementally till we have adequate number of tuples satisfying the user selectivity constraints. This might be the only possible approach if the user selectivity constraint cannot be filtered through the interface of hidden database. For e.g. if the user is interested in top-10 flights with in-flight wifi. However if the constraint cannot be entered via the airline interface, we can keep retrieving top ranked tuples till we have accumulated 10 flights with in-flight wifi. If the filters are too selective, then the number of tuples to be fetched before we return the user results could be very high.

However, if the user's constraints can be entered via the hidden database interface (but user still needs more than k results), then an alternate approach is possible. As an example, the user might be interested in top-20 flights with wifi on a top-10 interface where the wifi availability is an input attribute.

We can directly apply the techniques for extracting top ranked tuples over the subset of database that satisfies the selectivity constraints instead of applying it on the original database. This corresponds to prefixing the selectivity constraints to each of the queries executed by the algorithms. The candidate generation phase produces only tuples that satisfy the constraints.

The algorithms that work only on the database subset might seem to be a more efficient approach to solve the problem and in most scenarios it is. However, there are few factors that influence the output. First, if the selectivity constraints are coarse or not too selective, then a large section of database would be covered. This in turn, increases the chances of finding a correct set of top ranked tuples satisfying the constraints. If the number of tuples that match are small, then there is a high likelihood that the tuples are not comparable. In this case, we are left with a partial order of tuples instead of a total order.

Secondly, even if a total order exist among the top ranked tuples in the subset, it might not be possible to order them by only looking at the candidate tuples matching the constraints. This is because, the tuple(s) that helped to indirectly compare and order the candidate tuples, say t_x and t_y could itself not satisfy the selectivity constraint. In this case, the two tuples are incomparable, even though a global order exist between them. In both the scenarios, we are potentially left with a partial order. The techniques used in linearizing the partial order from §6-B can be used to solve this issue.

VII. EXPERIMENTAL RESULTS

In this section we describe our experimental setup, compare the performance of algorithms for candidate generation and candidate testing and show the efficiency and accuracy of our methods.

A. Experimental Setup

Hardware and Platform: All our experiments were performed on a quad-core 2 GHz AMD Phenom machine with 8 GB of RAM. The algorithms were implemented in Python.

Datasets: We used both synthetic and real-world data sets in the experiments. The synthetic dataset we used is a boolean one with 200,000 tuples and 80 attributes. The tuples are generated as i.i.d. data with each attribute having probability of $p = 0.5$ to be 1 (except for one experiment where we created different datasets with different values of p). We refer to this dataset as the BOOL-IID dataset. The real-world dataset we used consists of data crawled from the Yahoo! Autos website⁵, a real-world hidden database. It contains 200,000 used cars for sale in the Dallas-Fort Worth metropolitan area. There are 32 Boolean attributes such as A/C, Power Locks, etc, and 6 categorical attributes, such as Make, Model, Color, etc. The domain size of categorical attributes ranges from 5 to 16.

Real-World Online Experiment: In addition to the offline experiments described above, we also directly applied our techniques *online* over Amazon.com (specifically Amazon's Product Advertising API⁶) to discover the top-250 (according

⁵<http://autos.yahoo.com/>

⁶<https://affiliate-program.amazon.com/gp/advertising/api/detail/main.html>




Fig. 2. Query cost vs h on Boolean dataset




Fig. 3. Effect of different ranking functions on Autos dataset




Fig. 4. Query cost versus k




Fig. 5. Query cost versus database size




Fig. 6. Query cost versus skew

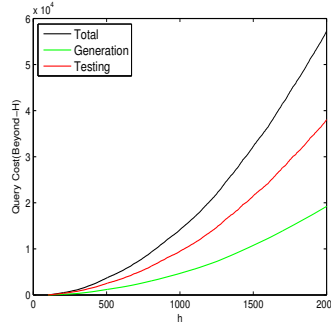


Fig. 7. Query cost versus large h

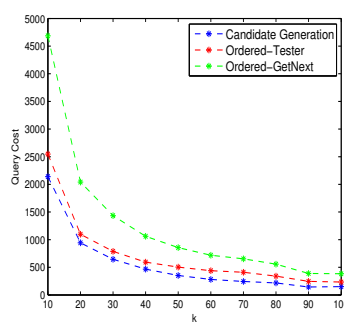


Fig. 8. Comparing candidate generation versus testing

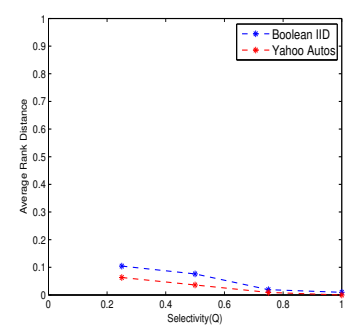


Fig. 9. Rank distance versus query selectivity

to sales rank) Amazon DVD titles from a top-100 interface⁷ provided by the API. Since the individual item description provided by Amazon.com reveals the sales rank of the item, we were able to verify the correctness of all results discovered by our algorithm. For this online experiment, (top- k) search query can be constructed using 15 categorical attributes such as Actor, Artist, Publisher, etc., with their domain sizes ranging from 5 to over 1,000. Amazon.com has a limit of 2,000 queries per IP address per hour.

Algorithms: We tested two algorithms BEYOND- h -GETNEXT and ORDERED-GETNEXT. However, since both these algorithms use the same candidate generation technique, we highlight the behavior of the candidate generation and testing phase separately. In other words, we plot the performance of algorithm GETNEXT for different parameters and then compare the performance of different candidate testing algorithms. This choice of presentation accentuates the improvements provided by the beyond- h -queries and the heuristic query ordering that gets masked when directly comparing BEYOND- h -GETNEXT and ORDERED-GETNEXT.

Performance Measures: We use query cost, the number of queries executed on the hidden database as the performance

⁷By default Amazon’s Product Advertising API provides a top-10 interface, while allowing a user to “Page Down” for up to 9 times, essentially leading to a top-100 interface.

measure. This includes the queries used to retrieve candidate tuples, queries to compare candidates and the beyond- h queries for each candidate. When the total order cannot be inferred, we use expected distance between randomly generated total order and the actual total order. The distance between two ranked lists is computed using Kendall- τ metric.

B. Experimental Results

In the following discussion we denote the number of top ranked tuples retrieved from the hidden database as h . In other words, it denotes the maximum number of invocations of GETNEXT by the third party service.

Query cost versus h : In our first experiment, we evaluated the performance of our algorithms BEYOND- h -GETNEXT and ORDERED-GETNEXT on the boolean dataset by investigating the query cost as a function of h for various different values of k . As Figure 2 shows, the query cost increases with increasing h , as is expected. Moreover, significant savings are achieved by using the ordering heuristic in ORDERED-GETNEXT. We also notice that k plays an important role in the efficiency of the algorithms: larger k results in more efficient performance. To consider a specific performance point, when $k = 75$ and $h = 200$, ORDERED-GETNEXT requires less than 300 additional queries to retrieve the extra 125 tuples.

We also performed similar experiments on the Autos dataset and observed similar trends, with ORDERED-GETNEXT out-




Fig. 10. Query Cost versus Query Selectivity

performing BEYOND- h -GETNEXT (Figure 3). Additionally, we also investigated the effect the specific ranking function used has on the performance of our algorithms. As Figure 3 shows, we used three different ranking attributes: TxnID (a unique ID for each tuple), as well as attributes such as Price and Miles. We note that the performance of our algorithms vary for different ranking functions, but nevertheless are still very efficient in all cases (and as noted earlier, our algorithms do not try to take advantage of any knowledge of these ranking functions).

Query cost versus k : In our next experiment, we investigated the effect of k on the query cost for fixed values of h , for both the boolean dataset as well as the Autos dataset. As Figure 4 shows, the positive effect of larger values of k on the query cost is dramatic, with larger values of k being very effective in reducing the query cost of our algorithms. This is to be expected, as our earlier arguments in the paper have shown that large k significantly reduces the number of queries needed in the candidate generation and testing procedures (since the number of minimal infrequent itemsets in a database rapidly reduces with increasing support threshold).

Query cost versus database size: Since our algorithms are designed to retrieve only the top- h tuples from the database, the actual size of the database should not have a significant impact on the performance of our algorithms. This is verified in Figure 5, which shows that the query cost remained practically unchanged for ORDERED-GETNEXT, even though we try our experiments on various fractional sizes of the original databases (the slight dip in query cost is attributable to the uncertainty of the sampling process). In this experiment, $k = 100$ and $h = 200$.

Query cost versus skew: We experimented with ORDERED-GETNEXT ($k = 100$, $h = 200$) on several boolean databases created with different values of skew parameter p . As Figure 5 shows, the algorithm is most efficient when the database has equiprobable 1s and 0s, but the cost increases when the proportion becomes unbalanced. This is attributable to the fact that when the database contains more 1s (or more 0s), the algorithm has to “dig deeper” - i.e., issue a larger number of (and more specific) queries in order to generate all candidates.

Effect of large h : Our earlier experiments were focused on values of h that were at most a small factor larger than k .




Fig. 11. Query Cost versus h on Amazon DVD website

Such values are meaningful in actual applications where a user is interested in seeing a few more tuples than what has been returned to her by the original query. But we were also interested in stress-testing our algorithms on much large values of h to see how they performed. Figure 7 shows the results of such an experiment using ORDERED-GETNEXT on the Autos dataset, where k was set at 100. As can be seen, the query cost increases quite significantly for much larger values of h , which leads to the conclusion that beyond a certain point, it is actually preferable to *crawl* the database and extract the top- h queries rather than use ORDERED-GETNEXT. The figure also profiles the separate query costs of the candidate generation and testing procedures.

Comparing generation versus testing procedures: In Figure 8, we compare the query costs of the two main procedures: candidate generation and candidate testing. We ran ORDERED-GETNEXT over the Autos dataset for $h = 200$ and varied k . As can be seen, the query cost is almost equally divided between the generation and test procedures for almost all points of the curve, with testing being slightly more expensive.

Effect of query selectivity: In Figures 9 and 10, we investigate the impact of selectivity. If a query is extremely selective, then it is clear that no algorithm can extract a total order of the top- h tuples. In such situations, our algorithms return a partial order of the top- h tuples. As discussed in VI-B, we compare a random total order that conforms to the returned partial order against the true top- h tuples for that query using Kendall- τ measure. As the query becomes less selective, the rank distance increases and its query cost becomes less, which is to be expected as the candidate testing procedure gets opportunities to terminate early as one needs a smaller number of queries to exclude a tuple from consideration. Similarly, as the query selectivity drops, our algorithm can retrieve the actual total order. Our experiments uses ORDERED-GETNEXT for both datasets, with $k = 100$ and $h = 200$.

Experiment against Amazon DVD Titles : To show the practicality of our algorithms, we retrieved the top-250 Amazon DVD titles in terms of their sales rank. Note that by default, Amazon only displays the top-100 items in any category. The correctness of our algorithm is verified by the checking the individual item description pages of the items discovered by

GETNEXT (which reveals the actual sales ranking of the items). The queries were made using the Amazon Product Advertising API and the maximum value of k is 100. A sample query to get the top-10 PG rated DVDs ordered by their salesrank is shown in footnote⁸. Figure 11 shows that when $k = 100$, the top-250 titles can be retrieved using fewer than 500 queries, well below the 2000 queries-per-hour-per-IP-address limit imposed by Amazon.com. The figure also shows the behavior of both BEYOND- h -GETNEXT and ORDERED-GETNEXT for different values of k and h .

VIII. RELATED WORK

Information Integration and Extraction for Hidden databases: A significant body of research has been done on information integration and extraction over hidden databases - see tutorials [8], [9]. Due to space limit, we only list a few closely-related work: [10] proposes a crawling solution. Parsing and understanding web query interfaces has been extensively studied (e.g., [11], [12]). The mapping of attributes across different web interfaces has also been addressed (e.g., [13]). Also related is the work on integrating query interfaces for multiple web databases in the same topic-area (e.g., [14], [15]). Our paper provides results orthogonal to these existing techniques as it represents the first formal study on retrieving top- h ($h > k$) tuples matching a user-specified query by reformulating the query through a top- k interface.

Data Analytics over Hidden Databases: There has been prior work on crawling, sampling, and aggregate estimation over the hidden web, specifically over text [16], [17] and structured [10] hidden databases and search engines [18]–[20]. Specifically, sampling-based methods were used for generating content summaries [21]–[23], processing top- k queries [24], etc. Prior work (see [3] and references therein) considered sampling and aggregate estimation over structured hidden databases.

Top- k Query Processing: There have been extensive studies on retrieving the top- k tuples over a traditional database - see [25] for a survey. Our approach differs by allowing the retrieval of top- h tuples through a restricted top- k web interface.

Frequent Itemset Mining: In §5, we map the discovery of beyond- h queries to the problem of infrequent-minimal-itemset mining - a problem well studied in data mining [7]. [6] provides additional details about algorithms and properties for infrequent itemset mining.

IX. CONCLUSION

In this paper we have initiated study on the problem of retrieving the top- h ($h > k$) tuples from a hidden web database that only provides a top- k search interface. To address the fundamental operator GETNEXT, we proposed a two-step process, candidate generation and candidate testing, and developed efficient algorithms for both steps. We conducted comprehensive set of experiments over synthetic datasets and

real-world hidden databases which demonstrate the effectiveness of our proposed techniques. There are multiple exciting directions for future research. We intend to investigate the possibility of retrieving the top ranked tuples approximately - for e.g., retrieve as many top ranked tuples under budget cost or in a rank agnostic fashion. Further, we plan to build attractive demonstrations of mashup applications against real-world hidden web databases.

REFERENCES

- [1] J. Madhavan, D. Ko, L. Kot, V. Ganapathy, A. Rasmussen, and A. Y. Halevy, "Google's Deep Web crawl," *Proceedings of The Vldb Endowment*, vol. 1, pp. 1241–1252, 2008.
- [2] M. Álvarez, J. Raposo, A. Pan, F. CACHED, F. Bellas, and V. Carneiro, "Crawling the content hidden behind web forms," in *Proceedings of the 2007 international conference on Computational science and Its applications - Volume Part II*, ser. ICCSA'07. Springer-Verlag, 2007, pp. 322–333.
- [3] A. Dasgupta, X. Jin, B. Jewell, N. Zhang, and G. Das, "Unbiased estimation of size and other aggregates over hidden web databases," in *SIGMOD*, 2010.
- [4] A. Dasgupta, G. Das, and H. Mannila, "A random walk approach to sampling hidden databases," in *SIGMOD*, 2007.
- [5] X. Jin, N. Zhang, and G. Das, "Attribute domain discovery for hidden web databases," in *SIGMOD*, 2011.
- [6] D. J. Haglin and A. M. Manning, "On minimal infrequent itemset mining," in *International Conference on Data Mining*, 2007.
- [7] J. Han, H. Cheng, D. Xin, and X. Yan, "Frequent pattern mining: current status and future directions," *DMKD*, 2007.
- [8] K. Chang and J. Cho, "Accessing the web: From search to integration," in *Tutorial, SIGMOD*, 2006.
- [9] A. Doan, R. Ramakrishnan, and S. Vaithyanathan, "Managing information extraction," in *Tutorial, SIGMOD*, 2006.
- [10] S. Raghavan and H. Garcia-Molina, "Crawling the hidden web," in *VLDB*, 2001.
- [11] E. Dragut, T. Kabisch, C. Yu, and U. Leser, "A hierarchical approach to model web query interfaces for web source integration," in *VLDB*, 2009.
- [12] Z. Zhang, B. He, and K. Chang, "Understanding web query interfaces: best-effort parsing with hidden syntax," in *SIGMOD*, 2004.
- [13] B. He, K. Chang, and J. Han, "Discovering complex matchings across web query interfaces: A correlation mining approach," in *KDD*, 2004.
- [14] E. Dragut, C. Yu, and W. Meng, "Meaningful labeling of integrated query interfaces," in *VLDB*, 2006.
- [15] B. He and K. Chang, "Statistical schema matching across web query interfaces," in *SIGMOD*, 2003.
- [16] Z. Bar-Yossef and M. Gurevich, "Mining search engine query logs via suggestion sampling," in *VLDB*, 2008.
- [17] K. Bharat and A. Broder, "A technique for measuring the relative size and overlap of public web search engines," in *WWW*, 1998.
- [18] K. Liu, C. Yu, and W. Meng, "Discovering the representative of a search engine," in *CIKM*, 2002.
- [19] M. Shokouhi, J. Zobel, F. Scholer, and S. Tahaghoghi, "Capturing collection size for distributed non-cooperative retrieval," in *SIGIR*, 2006.
- [20] Z. Bar-Yossef and M. Gurevich, "Efficient search engine measurements," in *WWW*, 2007.
- [21] J. Callan and M. Connell, "Query-based sampling of text databases," *ACM TOIS*, vol. 19, no. 2, pp. 97–130, 2001.
- [22] P. Ipeirotis and L. Gravano, "Distributed search over the hidden web: Hierarchical database sampling and selection," in *VLDB*, 2002.
- [23] Y.-L. Hedley, M. Younas, A. E. James, and M. Sanderson, "Sampling, information extraction and summarisation of hidden web databases," *Data and Knowledge Engineering*, vol. 59, no. 2, pp. 213–230, 2006.
- [24] N. Bruno, L. Gravano, and A. Marian, "Evaluating top-k queries over web-accessible databases," in *ICDE*, 2002.
- [25] I. Ilyas, G. Beskales, and M. Soliman, "A survey of top-k query processing techniques in relational database systems," *ACM Computing Surveys*, vol. 40, 2008.

⁸[http://ecs.amazonaws.com/onca/xml?Service=AWSECommerceService&AWSAccessKeyId=\[fill\]&Operation=ItemSearch&SearchIndex=DVD&ResponseGroup=Large,SalesRank&Sort=salesrank&AudienceRating=PG&Timestamp=\[fill\]&Signature=\[fill\]](http://ecs.amazonaws.com/onca/xml?Service=AWSECommerceService&AWSAccessKeyId=[fill]&Operation=ItemSearch&SearchIndex=DVD&ResponseGroup=Large,SalesRank&Sort=salesrank&AudienceRating=PG&Timestamp=[fill]&Signature=[fill])