

# Exploiting Correlated Attributes in Acquisitional Query Processing

Amol Deshpande  
amol@cs.umd.edu  
University of Maryland

Carlos Guestrin  
guestrin@cs.cmu.edu  
CMU

Wei Hong  
wei.hong@intel.com  
Intel Research Berkeley

Samuel Madden  
madden@csail.mit.edu  
MIT

## Abstract

Sensor networks and other distributed information systems (such as the Web) must frequently access data that has a high per-attribute *acquisition* cost, in terms of energy, latency, or computational resources. When executing queries that contain several predicates over such expensive attributes, we observe that it can be beneficial to use correlations to automatically introduce low-cost attributes whose observation will allow the query processor to better estimate the selectivity of these expensive predicates. In particular, we show how to build *conditional plans* that branch into one or more sub-plans, each with a different ordering for the expensive query predicates, based on the runtime observation of low-cost attributes. We frame the problem of constructing the optimal conditional plan for a given user query and set of candidate low-cost attributes as an optimization problem. We describe an exponential time algorithm for finding such optimal plans, and describe a polynomial-time heuristic for identifying conditional plans that perform well in practice. We also show how to compactly model conditional probability distributions needed to identify correlations and build these plans. We evaluate our algorithms against several real-world sensor-network data sets, showing several-times performance increases for a variety of queries versus traditional optimization techniques.

## 1. Introduction

The past decade has seen the emergence of a number of massive-scale distributed systems, including the Web, the Grid, sensor networks [22], and PlanetLab [1]. Unlike many earlier distributed systems that consisted of a few tens of locally connected nodes, these networks are characterized by their large size, broad geographic distribution, high-latency, and ability to interface users to huge amounts of remote data.

One of the novel challenges associated with query processing in such environments is managing the high cost associated with obtaining the data to be processed. While traditional disk-based DBMSes work very hard to reduce the number of disk I/Os, the high costs of disk access are usually amortized by reading or writing pages that contain many records each. In contrast, in this new class of distributed systems, the cost of fetching a single field of a single tuple is frequently measured in seconds and this effort usually cannot be spread across several fields. For example, in a sensor network, the time to acquire a single reading from a calibrated digital sensor can be as long as 1.3 seconds [21]. Following the nomenclature of Madden et al. [19], we refer to such systems as *acquisitional*, to reflect the high costs of accessing data and the fact that it must be actively captured when a query is issued rather than being available on disk at the time of query execution.

In this paper, we explore a class of techniques aimed at mitigating the costs of capturing data in acquisitional systems. These techniques are based on the observation that, in such systems, locally available information with a low cost of acquisition is often highly correlated with much more expensive information. If these correlations are perfect (i.e., there exists a one-to-

one mapping from cheap attribute values to expensive attribute values), then a cheap attribute can simply be used in place of an expensive one. However, even in the absence of perfect correlation, a cheap attribute can provide information that can aid in selectivity estimation in query optimization.

For example, Figure 1 shows a scatter-plot of light values<sup>1</sup>, versus time of day from a single sensor in sensor network deployed in our lab. In this network, light is an expensive attribute, requiring the processor to be on for almost a second to acquire each sample.

However, time of day is readily available in just a few microseconds. Furthermore, looking at this figure, it is clear that light and time of day are correlated: given a time of day, light values can be bound to within a fairly narrow band, especially at night (e.g., during hours 0-5 and 16-24).

We focus on using correlations of this sort to assist in cost-estimation and query processing for multi-predicate range queries, of the form:

$$\begin{aligned} & \text{SELECT } a_1, a_2, \dots, a_n \\ & \text{WHERE } l_1 \leq a_1 \leq r_1 \ (P_1) \\ & \text{AND} \\ & \text{AND } l_k \leq a_k \leq r_k \ (P_k) \end{aligned} \quad (1)$$

Where  $k \leq n$  and  $P_n$  is shorthand for the range predicate in the associated WHERE clause. We are particularly interested in such queries in acquisitional systems where the cost of acquiring some of the attributes is non-negligible and where correlations exist between one or more attributes. In this paper, we illustrate many examples of such correlations in a variety of real world data sets. A simple way to account for such correlations is to generate a multi-dimensional probability distribution over attribute values and exhaustively search for a single predicate order that will minimize the expected cost. Unfortunately, due to the presence of correlations, expected case performance is far from optimal. For example, given Figure 1, the selectivity of a predicate like `light < 100 Lux` is substantially lower during the day than at night. Thus, if the user issues a multi-predicate query with one such predicate over light, the optimal order of the predicates may vary depending on the time of day. To take advantage of this insight, a query optimizer can *condition* on the time and choose a different plan depending on whether it is night or day.

This observation may seem counter-intuitive: query evaluation can become cheaper by observing additional attributes. If, however, such additional observations are low-cost and allow the query processor to determine with high confidence that a query predicate,  $P$ , will reject a particular tuple, this can offer substantial performance gains. The query processor can imme-

<sup>1</sup>Lux is a measure of light received per  $m^2$  of surface area. 100,000 Lux corresponds to full sunlight; 500 - 1000 Lux are typical of a well-lit room; a rural moonlight night is 1-2 Lux.

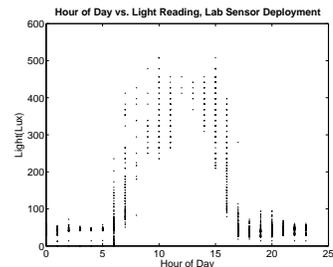


Figure 1: Hour of day vs. light values at a single sensor in authors' lab.

diately apply  $P$ , avoiding the cost of acquiring unnecessary and expensive attributes. In this paper, we show how to identify such correlations in data, and search for query plans that take advantage of those correlations. We demonstrate that the benefit of this technique can be substantial, offering a several-times performance speedup for a wide range of queries and query workloads. Our approach is *adaptive*, allowing a different query plan to be selected on a per-tuple basis, depending on the values of these conditioning attributes. To avoid the overhead of re-running the optimizer on every tuple, we pre-compute these conditional plans based on correlations observed before the query was initiated.

Throughout this paper, we use examples and data sets from sensor networks, though the techniques are general enough to apply to many different acquisitional environments; we return to a discussion of such environments in Section 7.

In summary, the contributions of this paper are:

- We show that correlations are prevalent in sensor network data, and that they can be exploited in multi-predicate range queries by introducing observations over inexpensive attributes correlated with query predicates.
- We introduce the notion of using *conditional plans* for query optimization with correlated attributes.
- We present an optimal, exponential-time algorithm for finding the best possible conditional plan given a multi-dimensional probability distribution over a set of attributes.
- We present a polynomial-time heuristic algorithm for reducing the running time of this algorithm and bounding the size of the generated plans.
- We present efficient ways of computing conditional probabilities over a set of predicates that allow us to avoid the exponential cost of representing an  $n$ -dimensional distribution over query variables.
- We evaluate the performance of our algorithms on several real-world data sets, showing substantial benefits over a large class of range queries.

## 2. Correlated attributes and conditional plans

In this section, we describe the problem we are seeking to address with our correlation-based techniques and the architecture of our query processing system. Though we use examples from the sensor network query processing system that we are building, the discussion in this section and the algorithms we develop in the later sections are equally applicable to other scenarios such as web querying, and even traditional query processing (Section 7).

### 2.1 Problem statement

We are concerned with the problem of generating an execution plan that will minimize the expected execution cost of a user-supplied range query of the form shown in query (1) above. Here, we assume that there are  $n$  attributes in the query table,  $X_1, \dots, X_n$ , and that the first  $m$  of them,  $X_1, \dots, X_m$  are referenced in the query (so  $m \leq n$ ). We assume that each attribute  $X_i$  takes on a value  $x_i \in \{1, \dots, K_i\}$ . Note that this definition requires real-valued attributes to be discretized appropriately (Cf. Section 4.3). In sensor networks, for example, this discretization is natural, as the resolution of the sensors is limited (e.g., to 10 bits when using the internal analog-to-digital (ADC) on the Berkeley Motes [7, 15].) We use the tuple  $\mathbf{x}$  to denote a particular assignment to all the attributes  $X_1 \dots X_n$ .

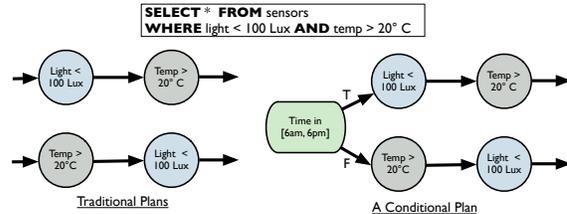


Figure 2: A conditional query plan that uses different ordering of query predicates depending on the time of day.

Each attribute is also associated with an *acquisition cost*, where  $C_i$  denotes the cost of acquiring the value of  $X_i$ . In a sensor network, this cost ranges from a few microjoules to tens of millijoules [21]; according to Madden et al. [19] the cost of acquiring a sensor reading once per second on a mote can be comparable to the cost of running the processor.

We denote the predicates in the query by  $\{\phi_1, \dots, \phi_p\}$ , and the logical *where* clause by  $\varphi$ . For simplicity, we use  $\varphi(\mathbf{x})$  to denote the truth value of  $\varphi$  for the tuple  $\mathbf{x}$ . Our goal is, thus, to find a minimum expected cost plan for determining whether  $\varphi(\mathbf{x})$  is true or false, where the expectation is taken over the possible assignments to the tuple  $\mathbf{x}$ .

Given such a query, a traditional query processor will choose an ordering of the query predicates,  $\phi_{i_1}, \dots, \phi_{i_p}$ , where  $i_1, \dots, i_p$  is a permutation of  $1, \dots, p$  (possibly based on the statistics that it maintains). This plan will then be evaluated, with attributed acquired as dictated by the plan.

We observe that, because of the natural correlations in the data, in many cases, such a rigid ordering of predicates might be highly inefficient compared to a more sophisticated execution plan that uses the observations acquired during query processing to change how it executes the rest of the plan. We illustrate this through an example.

Consider an example query containing two predicates  $\text{temp} > 20^\circ \text{C}$ , and  $\text{light} < 100 \text{ Lux}$ . Let the *a priori* selectivities of these two predicates be  $\frac{1}{2}$  and  $\frac{1}{2}$  respectively, and let the costs of acquiring the attributes be equal to 1 unit each. In that case, either of the two plans a traditional query processor might choose has expected cost equal to 1.5 units (Figure 2). However, we might observe that the selectivities of these two predicates vary considerably depending on whether the query is being evaluated during the day or at night. For instance, in Berkeley, during Summer, the predicate on  $\text{temp}$  is very likely to be false during the night, whereas the predicate on  $\text{light}$  is very likely to be false during the day. This observation suggests that using different plans depending upon the time of the day might be much more efficient. Figure 2 shows the *conditional plan* for this query that first checks the time of the day (*conditions* on the time of day), and evaluates the two query predicates in different order depending on the time. If we assume that the selectivity of the  $\text{temp}$  predicate drops down to  $\frac{1}{10}$  at night, and the selectivity of the  $\text{light}$  predicate is  $\frac{1}{10}$  during day, then the expected cost of this plan will be 1.1 units, a savings of almost 40%.

In this paper, we focus our conditional plans  $\mathcal{P}$  on simple binary decision trees<sup>2</sup>, where each interior node  $n_j$  specifies a binary *conditioning predicate*,  $T_j(X_i \geq x_i)$ , that splits the plan into two alternate conditional plans,  $\mathcal{P}_{T_j(\mathbf{x})=\text{T}}$  and  $\mathcal{P}_{T_j(\mathbf{x})=\text{F}}$ , where  $T_j(\mathbf{x})$  is the truth value of  $T_j$  on the tuple  $\mathbf{x}$ . At node  $n_j$ , the query processor will evaluate the predicate  $T_j(\mathbf{x})$  and choose to execute one of these two subplans depending on the value of

<sup>2</sup>Our approach can, of course, be extended to more general decision trees, as discussed in Section 7.

the predicate. Each conditioning predicate depends only on the value of a single attribute.

## 2.2 Plan evaluation cost

During execution of a plan  $\mathcal{P}$ , we simply traverse the binary tree defined by  $\mathcal{P}$ , acquiring any attributes we need to evaluate the conditioning predicates. For a particular tuple  $\mathbf{x}$ , *i.e.*, an assignment to all attributes, our plan will traverse a single path to a leaf of the binary tree  $\mathcal{P}$ , which is labeled with T or F indicating the truth value of  $\varphi(\mathbf{x})$ . The cost of this traversal is the sum of the cost of the attributes that are acquired by plan  $\mathcal{P}$  in this traversal. Specifically, at each node  $n_j$  in this traversal, if the attribute in the predicate  $T_j$  has already been acquired, then this node has zero *atomic cost*. However, if the attribute  $X_i$  in  $T_j$  has not yet been acquired, then the atomic cost of this node is  $C_i$ . For simplicity, we annotate each node  $n_j$  of our plan with this atomic cost  $C(n_j)$ . Note that the atomic cost of a leaf is 0, as no attributes are acquired at this point. We can now formally define the *traversal cost*  $C(\mathcal{P}, \mathbf{x})$  of applying plan  $\mathcal{P}$  to the tuple  $\mathbf{x}$  of the plan recursively by:

$$C(\mathcal{P}, \mathbf{x}) = \begin{cases} 0 & \text{if } |\mathcal{P}| = 1, \\ C(\text{Root}(\mathcal{P})) + C(\mathcal{P}_{T_j(\mathbf{x})}, \mathbf{x}), & \text{otherwise,} \end{cases} \quad (1)$$

where  $\text{Root}(\mathcal{P})$  is the root node of the tree for plan  $\mathcal{P}$ ,  $C(\text{Root}(\mathcal{P}))$  is the atomic cost of this root node as defined above, and the cost is 0 when we have reached a leaf, *i.e.*, when  $|\mathcal{P}| = 1$ .

Optimal planning is an optimization problem that involves searching the space of available conditional plans that satisfy the user's query for the plan  $\mathcal{P}^*$  with minimal expected cost:

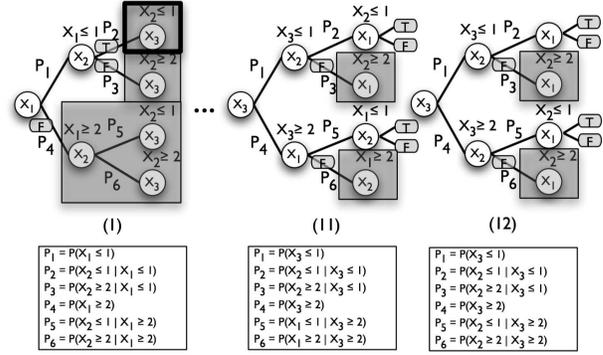
$$\begin{aligned} \mathcal{P}^* &= \arg \min_{\mathcal{P}} C(\mathcal{P}), \\ &= \arg \min_{\mathcal{P}} E_{\mathbf{x}} [C(\mathcal{P}, \mathbf{x})], \\ &= \arg \min_{\mathcal{P}} \sum_{\mathbf{x}} P(\mathbf{x}) C(\mathcal{P}, \mathbf{x}). \end{aligned} \quad (2)$$

Using the recursive definition of the cost  $C(\mathcal{P}, \mathbf{x})$  of evaluating a tuple  $\mathbf{x}$  in Equation (1), we can similarly specify a recursive definition of the expected cost  $C(\mathcal{P})$  of a plan. For this recursion, we must specify, at each node  $n_j$  of the plan, the conditional probability that the associated predicate  $T_j$  will be true or false, given the predicates evaluated thus far in the plan. We use  $\mathbf{t}$  to denote an assignment to this set of predicates. Using this notation, the expected plan cost is given by:

$$C(\mathcal{P}, \mathbf{t}) = \begin{cases} 0 & \text{if } |\mathcal{P}| = 1, \\ \begin{aligned} &C(\text{Root}(\mathcal{P})) + \\ &P(T_j | \mathbf{t})C(\mathcal{P}_{T_j}, \mathbf{t} \wedge T_j) + \\ &P(\neg T_j | \mathbf{t})C(\mathcal{P}_{\neg T_j}, \mathbf{t} \wedge \neg T_j), \end{aligned} & \text{otherwise,} \end{cases} \quad (3)$$

where  $C(\mathcal{P}, \mathbf{t})$  is the expected cost of the (sub)plan  $\mathcal{P}$  starting from its root node  $\text{Root}(\mathcal{P})$ , given that the predicates  $\mathbf{t}$  have been observed. At this point, the expected cost depends on the value of the new predicate  $T_j$ . With probability  $P(T_j | \mathbf{t})$ ,  $T_j$  will be true and we must solve the subproblem  $C(\mathcal{P}_{T_j}, \mathbf{t} \wedge T_j)$ , *i.e.*, the subplan  $\mathcal{P}_{T_j}$  after observing the original predicate values  $\mathbf{t}$  and the new predicate value  $T_j = \text{T}$ . Similarly, with probability  $P(\neg T_j | \mathbf{t}) = 1 - P(T_j | \mathbf{t})$ ,  $T_j$  will be false and we must solve  $C(\mathcal{P}_{\neg T_j}, \mathbf{t} \wedge \neg T_j)$ , *i.e.*, the subplan  $\mathcal{P}_{\neg T_j}$  after observing  $\mathbf{t}$  and the  $T_j = \text{F}$ . As before, when we reach a leaf ( $|\mathcal{P}| = 1$ ), the cost is 0. Now the expected cost of a plan  $\mathcal{P}$  is defined using Equation (3) by  $C(\mathcal{P}, \emptyset)$ .

We present several search algorithms for finding minimal cost plans in Section 3. In the remainder of this section, we illustrate the concepts presented thus far by considering the *naïve*



**Figure 3:** Some possible plans for the two predicate query  $X_1 = 1 \wedge X_2 = 1$ , with three attributes,  $X_1$ ,  $X_2$ , and  $X_3$  available for use in the query. Labels in the nodes indicate the attribute acquired at that point, and labels on the edges denote the probability of the outcome along that edge. Each  $P_i$  is given by the conditional probability expansions at the bottom of the figure. Terminal points in the tree are labeled with their outcomes: T is the tuple passes the query or F if it fails. Grayed out regions do not need to be explored, because either a predicate fails or all predicates are satisfied before they are reached.

generate-and-test algorithm that enumerates all possible conditional plans over a set of attributes.

### Example:

Consider the simple example of exhaustively enumerating the plans for three attributes,  $\{X_1, X_2, X_3\}$ , each with binary domain  $\{1, 2\}$ . Our query in this example is simply  $\varphi = (X_1 = 1 \wedge X_2 = 1)$ . Figure 3 shows three possible plans in this enumeration; there are 12 total possible plans in this case. Each node in this figure is labeled with the attribute acquired at that node, and the  $P_i$  values denote the conditional probability of the outcome along each edge occurring, given the outcomes already specified in the plan branch. Terminal outcomes are denoted as T or F, indicating that a tuple that reaches this node is output or rejected.

Note that, if at some branch of the plan we can prove the truth value of  $\varphi$  we do not need to further expand the tree. For example, in Figure 3, the grayed-out regions represent branches of the plan that do not need to be evaluated since a predicate has failed. The outlined box at the top of Plan (1) indicates that all query predicates have been evaluated on this branch, so  $X_3$  does not need to be acquired.

Given the plans in Figure 3, it is straightforward to read off the expected cost as defined in Equation (3). For example, for Plan (11) the cost is:

$$\begin{aligned} C(\text{Plan (11)}) &= C_3 + \\ &P(X_3 \leq 1)(C_2 + P(X_2 \leq 1 | X_3 \leq 1)C_1) + \\ &P(X_3 \geq 2)(C_1 + P(X_1 \leq 1 | X_3 \geq 2)C_2), \end{aligned}$$

where, for example, when branching on  $X_2$ , we do not need to consider the branch for the assignment  $X_2 = 2$  as this assignment makes  $\varphi$  false and we replace the grayed-out box by a leaf with value F.

At this point, the observation from the introduction bears repeating: the cheapest possible plan is not always the one that immediately acquire the query predicates. In our example, plan (12) could be cheaper than plan (1), if observing  $X_3$  has low cost and dramatically skews the probabilities of the attributes  $X_1$  and  $X_2$ . In particular, if  $X_3 = 1$  increases the probability of  $X_2 = 2$ , then observing  $X_3$  may allow us to select the particular attribute that is more likely to determine if  $\varphi = \text{F}$ . Thus, if  $X_3 = 1$ , the query processor may avoid acquiring  $X_1$ , which it does first in plan (1).

## 2.3 Estimating event probabilities

We now look at the question of determining the values of the conditional probabilities on the edges of a plan  $\mathcal{P}$ . The problem is to determine the probability that a predicate  $T_j$  of the form  $T_j(X_i \geq x_i)$  will be satisfied, given a set of conditioning predicates, or *a priori* conditions,  $\mathbf{t}$ . We can write this probability as  $P(T_j | \mathbf{t}) = P(T_j | t_0, \dots, t_{|\mathbf{t}|})$ , where  $t_i$  is the truth assignment to a predicate  $T_i$  that appears in  $\mathbf{t}$ . Bayes rule tell us that this quantity is equal to:

$$\frac{P(T_j, t_0, \dots, t_{|\mathbf{t}|})}{P(t_0, \dots, t_{|\mathbf{t}|})}.$$

Given this definition, a *naïve* method for computing such probabilities is to scan the historical data, computing the rows  $r$  that satisfy the predicates in  $\mathbf{t}$ , and the subset of  $r$ ,  $r_{sat}$ , that satisfies  $T_j = \mathbf{T}$ . The probability  $P(T_j = \mathbf{T} | \mathbf{t})$  is simply  $|r_{sat}|/|r|$ .

The disadvantage of this approach is that it requires a linear scan of the entire data set for each probability computation (although the space requirements are small, as it is trivial to compute the sizes of  $r$  and  $r_{sat}$  without actually materializing the rows.) As we will see, our algorithms can require tens of thousands of such computations, so linear scans can be costly, especially for large data sets. Furthermore, if the number of predicates in  $\mathbf{t}$  is large, then the number of records that satisfy these predicates may be very small. In such cases, our estimate of  $P(T_j = \mathbf{T} | \mathbf{t})$  will have very high variance. We describe more efficient (and potentially more accurate) techniques for estimating conditional probabilities in Section 5 below.

## 2.4 Revisiting the cost model

The cost model that we described in the earlier section focuses only on the acquisition costs of the attributes. There may be scenarios where other factors may be equally important. One of our targeted applications for the techniques developed in this paper is sensor networks query processing. In that setting, the communication cost incurred in transmitting the plan to the individual sensor nodes also needs to be taken into account. Larger conditional plans with more branches require more communication and the overheads of transmitting them could outweigh the benefits of using conditional plans. The plan size also important due to the limited storage available in sensor nodes.

A simple approach is to bound the plan size to be under some fixed size, where that size can be selected to easily fit into device RAM. Alternatively, we can modify our optimization problem to include both the cost of acquiring predicates and the cost of communicating the plan. In this case, the optimization becomes:

$$\arg \min_{\mathcal{P}} (C(\mathcal{P}) + \alpha \zeta(\mathcal{P})),$$

where  $\zeta(\mathcal{P})$  denotes size in bytes of  $\mathcal{P}$ , and  $\alpha$  is a scaling factor equal to  $\frac{\text{cost to transmit a byte}}{\text{\# tuples processed in query lifetime}}$  that allows us to capture the intuition that, as the running time of a continuous query gets large, the time spent in query execution will dominate the cost of sending the plan. We focus our presentation on limiting plan sizes, though this joint optimization problem could be addressed with an extension of our approach.

## 2.5 Architecture

Before detailing our algorithms for plan generation, we briefly discuss the architecture in which queries are executed. Current sensor network hardware [7], with 4K of RAM and a 4Mhz processor does not have sufficient power to run our optimization algorithms. However, once the plan is generated, the online plan execution step (*i.e.*, a simple traversal of a binary

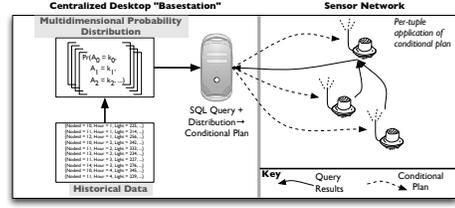


Figure 4: Query processing architecture. Conditional plans are generated from multi-dimensional probability distributions on the basestation and sent into the network. Sensors execute the plans locally, producing results which are routed back to the basestation for display.

tree) requires minimal computational power. As a result, we build the conditional plans offline, on a well-provisioned *basestation*, using historical readings collected over time. Once generated, these plans are distributed to the sensors (or other query processing nodes) for execution. Just as optimizer statistics are periodically collected and re-analyzed in a traditional DBMS, we envision that conditional plans may be re-generated at the user's request, or when the query processor detects substantial changes in the correlations.

Figure 4 illustrates the major components of our architecture. The basestation collects historical data, computes conditional probabilities over that data, using either a multidimensional probability distribution or by repeatedly scanning the data (as described above), and uses those probabilities plus the user's query to generate a conditional plan. This plan is sent into the network, executed by the sensors, and the results are transmitted back to the basestation.

## 3. Optimal solution

This section focuses on obtaining the minimal cost plan for determining the truth value of a logical clause  $\varphi$  as defined in Equation (2). We first determine the complexity class of this problem, and then present an exhaustive algorithm that can compute such optimal plan.

### 3.1 Problem complexity

The optimization problem outlined in the previous section is an instance of the general *minimum cost resolution strategy problem (MRSP)* [12], which seeks to find a minimum cost strategy for evaluating a boolean formula  $\varphi$ . The complexity of MRSP has been analyzed for various special cases. In this section, we analyze a case that is particularly relevant for our query optimization task. We focus on simple boolean formulas  $\varphi$  that are formed by the conjunction of unary predicates over a subset  $X_1, \dots, X_m$  of our  $n$  correlated attributes. This simple case includes our example Query (1), and, of course, if we were to include disjunctions the complexity will usually not decrease. For this problem class, we prove that:

**THEOREM 3.1.** *Let  $X_1, \dots, X_n$  be a set of attributes, and let  $\varphi$  be a conjunction of unary predicates on  $X_1, \dots, X_m$ , where  $m \leq n$ , then:*

1. *even if we have an oracle that instantaneously computes the conditional probabilities  $P(T_j | \mathbf{t})$ , for any predicate  $T_j$  and any assignment to any set of predicates  $\mathbf{t}$ , then, for some cost  $K$ , the problem of deciding whether a plan for determining the truth value of  $\varphi$  with expected cost lower than  $K$  exists is #P-hard<sup>3</sup>.*

<sup>3</sup>#P is a complexity class containing counting problems associated with decision problems in NP, e.g., counting the number of satisfiable assignments of a SAT formula is #P-complete.

2. if  $\mathcal{D}$  is a set of  $d$  tuples, and the expected cost of a plan  $\mathcal{P}$  is approximated by:

$$C(\mathcal{P}) = E_{\mathbf{x}} [C(\mathcal{P}, \mathbf{x})] \approx \frac{1}{d} \sum_{\mathbf{x} \in \mathcal{D}} C(\mathcal{P}, \mathbf{x}), \quad (4)$$

then, for some cost  $K$ , the problem of deciding whether a plan for determining the truth value of  $\varphi$  with approximate expected cost (as defined in Equation (4)) lower than  $K$  exists is NP-complete.  $\square$

Item 1 addresses a very general exact optimization case. Here, we decouple the problem of optimizing the plan from that of computing the conditional probabilities required to evaluate a plan, by adding an instantaneous oracle that can provide these probabilities. Even with this oracle, the complexity of the problem we would like to solve is still #P-hard (reduction from #3-SAT), indicating that efficient solution algorithms are unlikely to exist. Alternatively, as addressed in Item 2, we may wish to optimize our plan with respect to a particular dataset  $\mathcal{D}$ . This is a simpler case, as we only need to consider the tuples in  $\mathcal{D}$ , allowing us to ignore exponentially-many possible tuples in the optimization that are not in  $\mathcal{D}$ . Unfortunately, this simpler problem is still NP-complete (reduction from the complexity of binary decision trees [16]).

### 3.2 Exhaustive algorithm

The hardness results in the previous section indicate that a polynomial time algorithm for obtaining an optimal plan  $\mathcal{P}^*$  is unlikely to exist. In this section, we describe an optimal depth-first search algorithm that includes caching and pruning techniques to attempt to find  $\mathcal{P}^*$ . Even with caching and pruning, the worst case complexity of this algorithm is still exponential in the number of attributes. However, as we are computing our plans off-line, and then downloading these plans onto the sensor nodes, we expect that, for query tables over a small number of attributes, this exhaustive algorithm should be sufficient. These methods are then compared in Section 6.

Our dynamic programming algorithm (Figure 5) is based on the observation that once a predicate is used to split the conditional plan, it sub-divides the original problem into two independent subproblems. Each subproblem covers a disjoint subspace of the attribute-domain space covered by the original problem, and as such, can be solved independently of the other subproblem. A subproblem is thus defined by the attribute-domain space covered by it, or, specifically, by the ranges of the values that the attributes can take given the conditioning predicates used before this subproblem was generated.

We will denote a subproblem where each attribute  $X_i$  can take values in the range  $R_i = [a_i, b_i]$  by  $\text{Subproblem}(\varphi, R_1 = [a_1, b_1], \dots, R_n = [a_n, b_n])$ . Our initial problem is denoted by  $\text{Subproblem}(\varphi, R_1 = [1, K_1], \dots, R_n = [1, K_n])$ . We will specify the dynamic programming algorithm by defining the expected completion cost of a problem in terms of the costs of its subproblems. Here, we denote the cost of the optimal plan for  $\text{Subproblem}(\varphi, R_1, \dots, R_n)$  by  $J(\varphi, R_1, \dots, R_n)$ .

The subproblems that are formed from a particular problem are defined by the predicate used to split this problem. The potential predicates that can be used to further divide  $\text{Subproblem}(\varphi, R_1, \dots, R_n)$  are of the form  $T(a_i \leq X_i < x_i)$  and  $T(x_i \leq X_i \leq b_i)$ . That is, the range  $R_i = [a_i, b_i]$  of  $X_i$  is divided into  $[a_i, x_i - 1]$  and  $[x_i, b_i]$ . Our goal is to choose the optimal splitting attribute  $X_i$  and the optimal assignment from the range  $[a_i, b_i]$ .

The expected cost of a problem can now be defined recursively by the sum of three terms: the cost of acquiring the splitting attribute, plus the cost of each subproblem weighted by the

```

EXHAUSTIVEPLAN( $\varphi, R_1, \dots, R_n, \bar{C}$ )
// If ranges can determine truth value, or this subproblem has been cached, this
// branch is complete.
IF THE RANGES  $R_1, \dots, R_n$  ARE SUFFICIENT TO DETERMINE TRUTH OF
 $\varphi$ , OR ALL QUERY ATTRIBUTES HAVE BEEN OBSERVED:
  RETURN  $[0, \emptyset]$ .
IF SUBPROBLEM  $R_1, \dots, R_n$  HAS BEEN CACHED:
  RETURN CACHED RESULT.
// Find the optimal attribute to observe here, and, recursively, the optimal plan
// for the current ranges.  $\bar{C}$  is a bound that lets us avoid unnecessary search in
// high cost branches.
LET  $C_{min} \leftarrow \bar{C}$ , AND  $\mathcal{P} \leftarrow \emptyset$ .
FOR  $i \leftarrow 1$  TO  $n$ :
  // If the  $i$ th attribute has not been observed yet, pay its observation cost.
  IF  $R_i$  IS  $[1, K_i]$ , THEN:
    LET  $C' \leftarrow C_i$ .
  ELSE:
    LET  $C' \leftarrow 0$ .
  IF  $C' < C_{min}$ , THEN:
    // Iterate through each value of the  $i$ th attribute, computing the expected
    // cost recursively according to Equation 5.
    FOR  $x_i \leftarrow a + 1$  TO  $b$ , WHERE  $R_i = [a, b]$ :
      LET  $[C_{<x_i}, \mathcal{P}_{<x_i}] \leftarrow \text{EXHAUSTIVEPLAN}(\varphi,$ 
         $R_1, \dots, [a, x_i - 1], \dots, R_n, C_{min} - C')$ .
      LET  $P_{<x_i} \leftarrow P(X_i \in [a, x_i - 1] \mid R_1, \dots, R_n)$ .
      LET  $C' \leftarrow C' + P_{<x_i} C_{<x_i}$ .
    IF  $C' < C_{min}$ , THEN:
      LET  $[C_{\geq x_i}, \mathcal{P}_{\geq x_i}] \leftarrow \text{EXHAUSTIVEPLAN}(\varphi,$ 
         $R_1, \dots, [x_i, b], \dots, R_n, C_{min} - C')$ .
      LET  $P_{\geq x_i} \leftarrow (1 - P_{<x_i})$ .
      LET  $C' \leftarrow C' + P_{\geq x_i} C_{\geq x_i}$ .
    // If a lower cost split is found, the plan is modified to include new
    // split.
    IF  $C' < C_{min}$ , THEN:
      LET  $C_{min} \leftarrow C'$ .
      LET  $\mathcal{P} \leftarrow \left\{ \begin{array}{l} T(X_i < x_i) \rightarrow \mathcal{P}_{<x_i}, \\ T(X_i \geq x_i) \rightarrow \mathcal{P}_{\geq x_i} \end{array} \right\}$ .
  // Only cache results if an optimal plan is obtained, rather than early stopping
  // by pruning.
  IF  $C_{min} < \bar{C}$ , THEN:
    CACHE  $[C_{min}, \mathcal{P}]$  AS THE OPTIMAL PLAN FOR  $R_1, \dots, R_n$ .
  RETURN  $[C_{min}, \mathcal{P}]$ .

```

Figure 5: Exhaustive planning algorithm.

probability that the observed attribute value will lead to this particular subproblem. Specifically, if we choose to split on attribute  $X_i$ , we must first pay the cost of acquiring  $X_i$ , which we now denote by  $C'_i$ . If the original problem has not yet acquired this attribute, *i.e.*, if the range  $R_i$  of  $X_i$  still spans all possible values ( $R_i = [1, K_i]$ ), then we must pay the acquisition cost  $C'_i = C_i$ . However, if  $X_i$  has already been acquired (in which case,  $[a_i, b_i]$  will be a strict subset of  $[1, K_i]$ ), then  $C'_i = 0$ .

Now, we must consider the recursive cost of the subproblems. The particular choice of subproblem depends on the actual observed value of  $X_i$ . Thus, the probability we will need to solve  $\text{Subproblem}(\varphi, R_1, \dots, [a_i, x_i - 1], \dots, R_n)$  depends on the probability that  $X_i$  will take on a value in  $[a_i, x_i - 1]$  given the ranges  $R_1, \dots, R_n$  observed thus far, *i.e.*,  $P(X_i \in [a_i, x_i - 1] \mid R_1, \dots, R_n)$ . The optimal choice of splitting attribute and value is now the one that minimizes the expected cost of the resulting subproblems.

$$\begin{aligned}
J(\varphi, R_1, \dots, R_n) &= \min_i \min_{x_i \in [a_i+1, b_i]} C'_i + \\
&P(X_i \in [a_i, x_i - 1] \mid R_1, \dots, R_n) \times \\
&J(\varphi, R_1, \dots, [a_i, x_i - 1], \dots, R_n) + \\
&P(X_i \in [x_i, b_i] \mid R_1, \dots, R_n) \times \\
&J(\varphi, R_1, \dots, [x_i, b_i], \dots, R_n). \quad (5)
\end{aligned}$$

To complete this recursion, we must define the base cases. First, if all of the attributes in  $\varphi$  have already been observed, then the cost is 0. Specifically,  $J(\varphi, R_1, \dots, R_n) = 0$ , if each  $R_i = [a_i, b_i]$  is a strict subset of  $[1, K_i]$ , for each query attribute  $i = 1, \dots, m$ . Similarly, we may reach a point where the ranges  $R_1, \dots, R_n$  are sufficient to evaluate the truth value

of  $\varphi$ . For example, if  $\varphi$  is a conjunctive query and the  $j$ th predicate in the query is  $\phi_j(X_i \geq 10)$ , and we reach a subproblem  $J(\varphi, R_1, \dots, R_n)$  where  $R_i = [1, 9]$ , then the cost of this subproblem is 0. These base cases now complete the recursion. The optimal plan is obtained by caching, for each subproblem  $J(\varphi, R_1, \dots, R_n)$ , the attribute  $X_i$  and the assignment  $x_i$  that minimizes Equation (5).

Figure 5 shows the pseudo-code of our dynamic programming algorithm based on this recursion, that searches the plan space in a depth-first manner. Other than the caching optimization inherent in dynamic programming, we also use pruning to cut down on the search space explored. Our pruning strategy is simple: when exploring possible splitting attribute values, we store the cost of the best plan explored thus far ( $\bar{C}$ ), if the current branch exceeds this cost, we no longer need to explore. More elaborate pruning techniques, such as branch-and-bound, could potentially be effective in this problem.

**Complexity:** The complexity of EXHAUSTIVEPLAN depends on the total number of subproblems that may be generated during the execution. This number is bounded by  $\prod_{i=1}^n \frac{K_i \times (K_i - 1)}{2}$ , i.e., the number of possible ranges  $R_1, \dots, R_n$ . For each subproblem, we have to consider (1) each possible splitting attribute and assignment, in the worst case,  $\sum_{i=1}^n (K_i - 1)$ , and (2) compute the conditional probabilities required by the algorithm. Hence the running complexity of this algorithm in the worst case (when pruning does not help in reducing the search space) is  $O(nKK^{2n} + K_p K^{2n})$ , where  $K = \max_i K_i$  is the maximum number of possible assignments for an attribute, and  $K_p$  denotes the complexity of computing the conditional probabilities — we will discuss this in detail in Section 5; thus, this optimal algorithm is only feasible for a small number of attributes, each with a small number of assignments.

## 4. Heuristic solutions

The complexity of the exhaustive algorithm described in the last section is prohibitive except for the simplest of problems. In this section, we present several heuristics for finding good conditional plans; we compare these algorithms in Section 6.

### 4.1 Sequential Plans

A *sequential plan* is defined to be a plan that does not use any conditioning predicates to *split* the plan, choosing instead a sequential order on predicates that is followed regardless of the observed values, until a stopping condition is reached. As we will see later, sequential plans form the building blocks of our heuristic conditional planning algorithm. We consider three algorithms for choosing sequential plans.

#### 4.1.1 Naïve Algorithm

For conjunctive queries, *Naïve* algorithm, used by most traditional query optimizers [18, 13, 3], simply orders predicates by  $\frac{\text{cost}}{1 - \text{selectivity}}$ , where selectivity is the marginal probability that the predicate does not output a tuple as computed from the historical data. Because this algorithm does not take into account data correlations, we expect it to perform poorly over correlated data sets. We are not aware of any extensions of this algorithm for arbitrary queries.

#### 4.1.2 Optimal Sequential Plan (OptSeq)

Interestingly, the exhaustive algorithm described in Section 3 can be used to compute the optimal sequential plan for a conjunctive query. We will use  $\hat{P} = \text{OPTSEQUENTIAL}(\varphi, R_1, \dots, R_n)$  to define the optimal sequential plan  $\hat{P}$  given that we have already observed the ranges

$R_1, \dots, R_n$  to the attributes  $X_1, \dots, X_n$ . Note that any conjunctive query  $\varphi$  over  $X_1, \dots, X_m$  can be written as  $\varphi = \bigwedge_{i=1}^m \phi_i(l_i \leq X_i \leq r_i)$ . If we observe the value of some query attribute  $X_i$ , then either  $\varphi$  is proven to be false, or we need to observe other query attributes. In general, the choice of next query attribute will depend on the particular observed value of  $X_i$ , yielding a complex conditional plan. However, we can redefine our optimization problem to restrict potential conditioning predicates to be only the query predicates  $\phi_i(l_i \leq X_i \leq r_i)$ . If  $\phi_i$  is observed to be false, we can terminate the search as above. Otherwise, we continue by choosing another predicate  $\phi_j$ , and thus observing attribute  $X_j$ . If  $\phi_j$  is false, we have proven that  $\varphi$  is false, otherwise we simply recurse. Clearly, this procedure will lead to a sequential order over attributes. The same optimal dynamic programming algorithm presented in the previous section can be applied in this redefined version of the problem to obtain this optimal sequential plan. This redefined problem is specified as follows:

- We redefine each query attribute  $X_i$  by  $X'_i \in [0, 1]$ , where  $X'_i = 1$  if  $X_i \in [l_i, r_i]$  and  $X'_i = 0$  otherwise. Non-query attributes  $X_{m+1}, \dots, X_n$  are removed.
- In terms of these new attributes, our query now becomes  $\varphi' = \bigwedge_{i=1}^m \phi'_i(X'_i = 1)$ . Note that  $\varphi$  is true if and only if  $\varphi'$  is true.

We can view this redefinition of the problem as a discretization of the possible values of the attributes into 2 bins, i.e.,  $X_i \in [l_i, r_i]$  and  $X_i \notin [l_i, r_i]$ . We must also define the costs and probabilities in this redefined problem, which, of course, depend on the input ranges  $R_1, \dots, R_n$  of OPTSEQUENTIAL:

- As in the original problem, the cost for acquiring each  $X'_i$  is equal to  $C_i$ , if  $R_i = [1, K_i]$ , and 0 otherwise.
- The joint probability for the new attributes  $X'_1, \dots, X'_m$   $P(X'_1, \dots, X'_m)$  is initialized to be the joint probability of the truth value of the query predicates, given the input ranges:  $P(X'_1, \dots, X'_m \mid R_1, \dots, R_m)$ .

Due to lack of space, we defer the development of an algorithm for finding optimal sequential plans for arbitrary queries to the full version of the paper [8].

#### 4.1.3 Greedy Sequential Plan (GreedySeq)

This heuristic was initially proposed by Munagala et al. [20]. The running time complexity of finding the optimal sequential plan is  $O(m2^m)$ , which makes it impractical for queries with large number of predicates (the problem of finding optimal sequential plan for conjunctive queries is NP-Hard [20]). For queries with large number of predicates, we can instead use the greedy heuristic proposed by [20]. Briefly, this heuristic proceeds by choosing the predicates to be applied in order:

1. Let  $P_a$  be the set of predicates that have already been chosen (set to be empty at the beginning).
2. For each of the predicates  $\phi_i$ ,  $1 \leq i \leq m$ , let  $p_i$  be the probability that the predicate is satisfied *given* that all the predicates in  $P_a$  have already been satisfied.
3. Choose the predicate  $\phi_j$  that minimizes  $\frac{C_j}{1 - p_j}$  to be evaluated, where  $C_j$  is the cost of acquiring the corresponding attribute.
4. Add  $\phi_j$  to  $P_a$ .
5. Repeat Step 2 until all predicates have been chosen.

This algorithm can be shown to be 4-approximate [20], and as demonstrated in [20], it outperforms *Naïve* since it takes into account correlations in the data.

```

GREEDYSPLIT( $\varphi, R_1, \dots, R_n$ )
// Greedily find the best attribute to split on, assuming that the optimal sequential
// plan is used after this split.
LET  $C_{min} \leftarrow \infty$ .
FOR  $i \leftarrow 1$  TO  $n$ :
// If the  $i$ th attribute has not been observed yet, pay its observation cost.
IF RANGE OF  $X_i$  IS  $[1, K_i]$ , THEN:
LET  $C' \leftarrow C_i$ .
ELSE:
LET  $C' \leftarrow 0$ .
IF  $C' < C_{min}$ , THEN:
// Iterate through each value of the  $i$ th attribute, computing the expected
// cost according to Equation 6.
FOR  $x_i \leftarrow a + 1$  TO  $b$ , WHERE  $R_i = [a, b]$ :
LET  $[\hat{C}_{<x_i}, \hat{P}_{<x_i}] \leftarrow \text{OPTSEQUENTIAL}(\varphi,$ 
 $R_1, \dots, [a, x_i - 1], \dots, R_n)$ .
LET  $P_{<x_i} \leftarrow P(X_i \in [a, x_i - 1] \mid R_1, \dots, R_n)$ .
LET  $C' \leftarrow C' + P_{<x_i} \hat{C}_{<x_i}$ .
IF  $C' < C_{min}$ , THEN:
LET  $[\hat{C}_{\geq x_i}, \hat{P}_{\geq x_i}] \leftarrow \text{OPTSEQUENTIAL}(\varphi,$ 
 $R_1, \dots, [x_i, b], \dots, R_n)$ .
LET  $P_{\geq x_i} \leftarrow (1 - P_{<x_i})$ .
LET  $C' \leftarrow C' + P_{\geq x_i} \hat{C}_{\geq x_i}$ .
// If a lower cost split is found, the plan is modified to include new
// split.
IF  $C' < C_{min}$ , THEN:
LET  $C_{min} \leftarrow C'$ .
LET  $T \leftarrow T(X_i \geq x_i)$ ,  $\hat{P}_{<} \leftarrow \hat{P}_{<x_i}$ ,  $\hat{P}_{\geq} \leftarrow \hat{P}_{\geq x_i}$ .
RETURN  $[C_{min}, T, \hat{P}_{<}, \hat{P}_{\geq}]$ .

```

Figure 6: Greedy selection of binary splitting point.

## 4.2 Greedy conditional planning algorithm

Next we present our greedy conditional planning algorithm that chooses conditioning predicates greedily, by making locally optimal decisions.

### 4.2.1 Greedy binary splits

Recall that the exhaustive algorithm described in Section 3 finds the optimal split point for each problem, by considering the optimal values of each resulting subproblem. Instead, we focus on locally optimal greedy splits. The choice of this greedy split point depends on several factors, including the conditional probability distribution after the split, the cost of attributes, and, most importantly, the query at hand. Thus, it is insufficient to use decision tree or histogram building heuristics that simply consider the distribution when making this choice; we must formulate the expected decrease in cost more precisely. In particular, we define the *locally optimal binary split* to be the split point that results in maximum *immediate benefit* over the optimal sequential plan. More formally, for a problem *Subproblem*( $\varphi, R_1, \dots, R_n$ ), the locally optimal binary split,  $\text{GREEDYSPLIT}(\varphi, R_1, \dots, R_n)$  is defined as a greedy version of Equation (5):

$$\begin{aligned}
\text{GREEDYSPLIT}(\varphi, R_1, \dots, R_n) = & \min_i \min_{x_i \in [a_i+1, b_i]} C'_i + \\
& P(X_i \in [a_i, x_i - 1] \mid R_1, \dots, R_n) \times \\
& \hat{J}(R_1, \dots, [a_i, x_i - 1], \dots, R_n) + \\
& P(X_i \in [x_i, b_i] \mid R_1, \dots, R_n) \times \\
& \hat{J}(R_1, \dots, [x_i, b_i], \dots, R_n), \quad (6)
\end{aligned}$$

where  $\hat{J}(R_1, \dots, R_n)$  is the expected cost of the optimal sequential plan starting from the ranges  $R_1, \dots, R_n$ . Figure 6 shows the complete pseudo-code of the algorithm that finds the locally optimally split point, given a sub-routine  $\text{OPTSEQUENTIAL}$  that computes the optimal sequential plan for a subproblem.

Note that instead of using *OptSeq* algorithm to generate the

```

GREEDYPLAN( $\varphi, \text{MAXSIZE}$ )
// Start plan with one leaf containing the optimal sequential plan, and greedily
// add splits using GreedySplit in Figure 6. A priority queue over the leaves  $n_i$ 
// of the current plan determines which leaf to expand: the priority of a leaf
// is the improvement in expected cost of splitting that leaf, versus using the
// optimal sequential plan.
LET  $\hat{\mathcal{P}} \leftarrow \text{OPTSEQUENTIAL}(\varphi, [1, K_1], \dots, [1, K_n])$ .
LET  $[\bar{C}, T(X_j \geq x_j), \hat{P}_{<x_j}, \hat{P}_{\geq x_j}] \leftarrow \text{GREEDYSPLIT}(\varphi,$ 
 $[1, K_1], \dots, [1, K_n])$ .
LET  $n_1 \leftarrow [\hat{\mathcal{P}}, T(X_j \geq x_j), \hat{P}_{<x_j}, \hat{P}_{\geq x_j}, [1, K_1], \dots, [1, K_n]]$ .
LET  $\mathcal{P} \leftarrow \{n_1\}$ .
ADD TO QUEUE  $n_1$  WITH PRIORITY  $C(\hat{\mathcal{P}}) - \bar{C}$ .

// Iterate expanding leaves until plan size reaches limit.
WHILE  $|\mathcal{P}| < \text{MAXSIZE}$ :
REMOVE TOP OF QUEUE:
 $n_i = [\hat{\mathcal{P}}, T(X_j \geq x_j), \hat{P}_{<x_j}, \hat{P}_{\geq x_j}, R_1, \dots, R_n]$ ,
WHERE  $R_j = [a, b]$ .
LET  $[\bar{C}, T(X_u \geq x_u), \hat{P}_{<x_u}, \hat{P}_{\geq x_u}] \leftarrow \text{GREEDYSPLIT}(\varphi,$ 
 $R_1, \dots, [a, x_j - 1], \dots, R_n)$ .
// The leaf  $n_i$  is expanded to two leaves,  $n'_i$  and  $n''_i$  that are added to the
// queue.
LET  $n'_i \leftarrow [\hat{P}_{<x_j}, T(X_u \geq x_u), \hat{P}_{<x_u}, \hat{P}_{\geq x_u},$ 
 $R_1, \dots, [a, x_j - 1], \dots, R_n]$ .
ADD TO QUEUE  $n'_i$  WITH PRIORITY
 $P(R_1, \dots, [a, x_j - 1], \dots, R_n) (C(\hat{P}_{<x_j}) - \bar{C})$ .
LET  $[\bar{C}, T(X_v \geq x_v), \hat{P}_{<x_v}, \hat{P}_{\geq x_v}] \leftarrow \text{GREEDYSPLIT}(\varphi,$ 
 $R_1, \dots, [x_j, b], \dots, R_n)$ .
LET  $n''_i \leftarrow [\hat{P}_{\geq x_j}, T(X_v \geq x_v), \hat{P}_{<x_v}, \hat{P}_{\geq x_v},$ 
 $R_1, \dots, [x_j, b], \dots, R_n]$ .
ADD TO QUEUE  $n''_i$  WITH PRIORITY
 $P(R_1, \dots, [x_j, b], \dots, R_n) (C(\hat{P}_{\geq x_j}) - \bar{C})$ .
// The plan is modified to include new split.
LET  $\mathcal{P} \leftarrow \mathcal{P} \cup \{n_i \rightarrow n'_i, n_i \rightarrow n''_i\}$ .
RETURN  $\mathcal{P}$ .

```

Figure 7: Greedy planning algorithm.

base plans, we could instead use the *GreedySeq* algorithm, or any other sequential planning algorithm, for that purpose. Considering that *OptSeq* is exponential in the number of predicates, this may be required for queries with large number of predicates.

### 4.2.2 Greedy planning algorithm

Our greedy planning algorithm uses the greedy splits described above to efficiently find a good conditional plan for the query. The algorithm maintains a current decision list plan  $\mathcal{P}$ , which is initially defined to be just a leaf with the root node. Each leaf  $n_i$  in the current plan stores an optimal sequential plan  $\hat{\mathcal{P}}$  for its subproblem, and the optimal greedy split plan defined by Equation (6). The benefit of applying the greedy binary split at  $n_i$  over the optimal sequential plan is given by the expected cost of the sequential plan,  $C(\hat{\mathcal{P}})$  minus the expected cost of the optimal greedy split  $\bar{C}$  computed by Equation (6). We maintain a priority queue over leaves of the current plan. A particular leaf  $n_i$ , whose subproblem is  $R_1, \dots, R_n$ , is inserted in the queue with a priority given by the difference  $C(\hat{\mathcal{P}}) - \bar{C}$  weighed by the probability that our plan will reach the subproblem in leaf  $n_i$ ,  $P(R_1, \dots, R_n)$ , as the expected gain of expanding  $n_i$  is  $P(R_1, \dots, R_n) (C(\hat{\mathcal{P}}) - \bar{C})$ .

Our greedy algorithm chooses the next leaf to expand from the queue based on this priority.

Figure 7 illustrates the complete greedy algorithm. We start with a sequential plan for the root node, and a single split on this node. We then iterate through the queue. We remove the top leaf  $n_i$ , and create two children leaves  $n'_i$  and  $n''_i$ , based on

the greedy split from  $n_i$ . These new leaves are then added to the queue with appropriate priorities.

### 4.2.3 Complexity

The running complexity of this algorithm can be seen to be  $O(nKC_{seq}N + nKK_p)$ , where  $C_{seq}$  denotes the cost of computing sequential plan for a subproblem (which will depend on the sequential planning algorithm used for that purpose),  $N$  denotes the number of splits performed by the algorithm, and  $K_p$  denotes the complexity of computing the probabilities required by this algorithm.

## 4.3 Restricting Choice of Split Points

The final heuristic we present is to reduce the number of candidate split points considered by the conditional planning algorithms. In particular, for continuous variables, we need to choose how to discretize the range of the variable; even for discrete variables with large domains, we may need to restrict the number of split points considered by the conditional planning algorithms. The solution we propose in this paper is to simply divide the domain of the variable into equal sized ranges and consider only the end-points of the ranges. We defer the issue of selecting the candidate split points more intelligently to future work. If  $r_i$  denote the number of split points considered for variable  $X_i$ , we define the *Split Point Selection Factor* (SPSF) of a conditional planning algorithm to be:

$$SPSF = \prod_{1 \leq i \leq n} r_i$$

A small SPSF can be expected to reduce the effectiveness of our algorithms as the choice of split points is too restricted; whereas a SPSF equal to the product of domain sizes allows arbitrary choice of split points.

## 5. Efficient probability computations

Until now, we have ignored the issue of computing the event probabilities required by our planning algorithms. In particular, the algorithms require the computation of conditional probabilities of predicates given a set of predicates that are already known to be satisfied. In this section, we will discuss how these probability computations can be performed efficiently, given a historical dataset of samples.

Consider the case where each required conditional probability is estimated from counts from a dataset  $\mathcal{D}$  of  $d$  tuples as described in Section 2.3. Consider a subproblem generated during the execution of either the EXHAUSTIVEPLAN algorithm or the GREEDYSPLIT algorithm,  $Subproblem(\varphi, R_1 = [a_1, b_1], \dots, R_n = [a_n, b_n])$ , and let the part of the dataset that satisfies the conditions  $X_i \in R_i, \forall i$ , be  $\mathcal{D}(R_1, \dots, R_n)$ .

### 5.1 Probabilities for exhaustive planning

While solving a subproblem, the probabilities that EXHAUSTIVEPLAN (Figure 5) needs to compute are:

$$P_{<x_i} = P(X_i \in [a_i, x_i - 1] \mid R_1, \dots, R_n), \forall i, \forall a_i < x_i \leq b_i.$$

As we can see, all the probabilities required by this step of the algorithm (ignoring any recursive calls for subproblems) are probabilities conditioned on  $(X_1 \in R_1) \wedge \dots \wedge (X_n \in R_n)$ . Note that  $P(x_i \mid R_1, \dots, R_n)$  is simply an independent normalized histogram of  $X_i$  for the data in  $\mathcal{D}(R_1, \dots, R_n)$ . We can compute these histograms for every attribute  $X_1, \dots, X_n$  with one pass over the dataset. Once we have  $P(x_i \mid R_1, \dots, R_n)$ , we can compute the probabilities of the required ranges incrementally by noting that:

$$P_{<(x_i+1)} = P_{<x_i} + P(x_i \mid R_1, \dots, R_n). \quad (7)$$

Thus, we can compute the probabilities of all required ranges in time only  $O(|\mathcal{D}|nK + nK)$ , where  $K = \max_i K_i$ .

We must also consider generating the dataset for each subproblem called recursively from  $R_1, \dots, R_n$ . Here, we can create an index independently for each attribute, for each value of the attribute in time  $O(|\mathcal{D}|nK)$ . We can now use a similar technique to the one used for range probabilities in Equation (7): the set of indices for the range  $[1, x_i]$  is equal to the set of indices for  $[1, x_i - i]$  union with the indices for  $x_i$ . Thus, selecting the dataset for all resulting subproblems is also  $O(|\mathcal{D}|nK + nK)$ . Therefore, the complexity of the exhaustive algorithm using the dataset to estimate probabilities is  $O(|\mathcal{D}|(nK)^2K^{2n})$ .

## 5.2 Probabilities for greedy planning

When the GREEDYSPLIT is considering a conditioning predicate  $X_i \in [a_i, x_i - 1]$ , it requires computation of two types of distributions:

- $P_{<x_i} = P(X_i \in [a_i, x_i - 1] \mid R_1, \dots, R_n)$  as above,
- a joint probability distribution over the rediscritized attributes in the query predicates, conditioned on  $X_1 \in R_1 \wedge \dots \wedge X_n \in R_n$  and on the particular choice of splitting attribute and assignment.

The conditional probabilities  $P_{<x_i}$  can be computed incrementally as in Equation (7). The joint distribution over the rediscritized attributes can be similarly computed from a normalized joint histogram over each attribute  $X_i$  and over the rediscritized attributes  $X'_1, \dots, X'_m$ . We can then use an incremental rule similar to the one in Equation (7) to compute the required joint probabilities. Thus, the total computation required in at each split is  $O(|\mathcal{D}|nK)$  to create the histograms and  $O(nK + nKv^m)$  to incrementally update the joint distributions. Using these optimizations, if the number of subproblems solved by GREEDYSPLIT is  $N$ , then the total running time is  $O(NnK(C_{optseq} + v^m))$ , or, for conjunctive queries,  $O(NnKm2^m)$ .

## 6. Evaluation

We now experimentally evaluate the algorithms described in the above sections over various datasets. We demonstrate that our algorithm offers a substantial performance increase over a non-conditional query execution engine for a wide class of queries over real-world data sets. We also show that our GreedySplit heuristic closely approximates the exhaustive algorithm and that our techniques scale with respect to the number of query predicates, the domain size of the attributes in the query, and the amount of historical data. Finally, we show that our heuristic performs well with a small number of splits, corresponding to a small plan size that may be required in memory-constrained environments.

We implemented our algorithms in Java on traditional PC hardware. We evaluate their performance by costing and running plans on this centralized PC; we reserve implementing a plan executor that runs on sensor network hardware for future work. We believe our techniques are readily implementable on sensor networks, as the energy overhead required to execute a conditional plan, especially when compared to the costs of acquiring sensor readings, will be small.

**Datasets:** We present results for two real world data sets collected from TinyOS-based sensors, as well as a synthetic dataset adapted from [2].

- **Lab:** This is a data set containing 400,000 *light*, *temperature*, and *humidity* readings collected every two minutes for several months from about 45 motes. The data set also

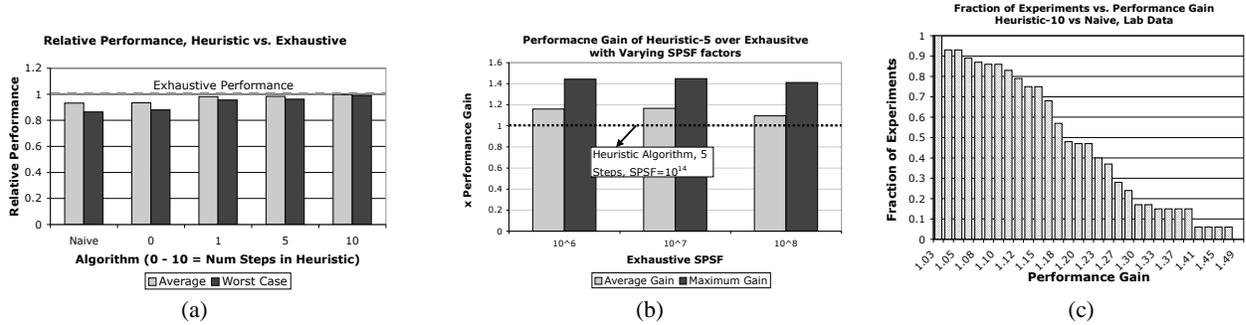


Figure 8: Quality of plans: (a) *Exhaustive* algorithm versus the *Naive* and *Heuristic* algorithms on the lab data set, the three *Heuristic* bars represent different numbers of splits in the *Heuristic* algorithm; (b) impact of using a smaller SPSF's to train the *Exhaustive* algorithm (versus *Heuristic-5*); (c) cumulative frequency of performance gain for experiments over the lab data set, the frequency of at a particular x-coordinate indicates the fraction of experiments that did at least that well.

includes *id* of the generating sensors, *time of the day*, and battery *voltage*. The first three attributes are very expensive to acquire, whereas the rest of the attributes are relatively inexpensive. For simplicity, we model the costs of the first three attributes to be 100 units each, and the costs of acquiring the rest of the attributes to be 1 unit each.

- **Garden-5, Garden-11:** This dataset consists of data collected from a set of 11 motes deployed in a forest. Corresponding to each mote, there are 3 attributes, *temperature*, *voltage*, and *humidity*. We present results for a subset of this dataset containing data from 5 motes (Garden-5) as well as for the entire dataset (Garden-11). The queries are issued against the sensor network as a whole; in essence, we treat these datasets as having 16 or 34 attributes respectively (3 per mote, and time), and query the state of the sensor network. Finally, the costs of acquiring *temperature* and *humidity* are set to be 100 units each, whereas the cost of the rest of the attributes is set to be 1 unit each.
- **Synthetic:** Finally, we present results on a synthetic dataset adapted from [2]. The data generator uses three parameters,  $n$  (*number of attributes*),  $\Gamma$  (*correlation factor*), and  $sel$  (*unconditional selectivity*). The  $n$  attributes are divided into  $\lfloor n/\Gamma + 1 \rfloor$  groups containing  $\Gamma + 1$  attributes each. Each attribute takes exactly two values, 0 and 1. The data is generated so that (1) any two attributes in the same group are positively correlated, and have identical values for 80% of the tuples in the dataset, (2) any two attributes in different groups are distributed independently, and (3) for each attribute, the fraction of the tuples for which that attribute is equal to 1 is approximately  $sel$ . To model cheap correlated attributes, we let one attribute in each group to have cost equal to 1 unit, whereas the rest of the attributes have cost 100 units each. Finally, the query is simply a conjunctive query checking whether all expensive attributes have value = 1. We present results for various settings of the parameters.

**Test v. Training:** Our plans generated by our algorithms are built using a set of *training* data, and evaluated on a disjoint set of *test* data (for the first Lab and Garden datasets). Readings from these two sets are from non-overlapping time windows, simulating the case where historical data is used to generate a model that is later run for days or weeks within a sensor network.

**Algorithms Compared:** We consider the following algorithms in this study:

- **Naive** (Cf. Section 4.1.1),

- **CorrSeq:** Sequential plan chosen by considering data correlations. If the number of attributes is small (e.g. *Lab* dataset), we use the *OptSeq* algorithm, whereas for the other datasets, we use the *GreedySeq* algorithm for this purpose.
- **Exhaustive:** The conditional plan computed by our exhaustive algorithm.
- **Heuristic-k:** The conditional plan generated by our greedy conditional planning algorithm with at most  $k$  conditional branches allowed. Once again, we use *OptSeq* for choosing the base plans for the first dataset, and *GreedySeq* for choosing the base plans for the second dataset.

## 6.1 Lab Dataset

In our first experiment, we compare the performance of the exhaustive algorithm to our greedy heuristic. Our test queries consist of three-predicate queries over the lab data. We found that the *Naive* is very successful on queries with very low selectivities, as the first few chosen attributes are usually sufficient to answer the query; for this reason, we chose a more challenging setting where most predicates generated for our experiments are satisfied by a large (approximately 50%) portion of the data set. For each query, we select, uniformly and at random, the left endpoint of the range of the query; the width of each predicate is chosen to be two standard deviations of the attribute which it is over. Except where noted, to allow the exhaustive algorithm to run, we use a SPSF of  $10^8$  for the experiments with exhaustive in this section.

Figure 8(a) shows the average performance of *Heuristic* relative to *Exhaustive* when both are running on the dataset with SPSF set to  $10^8$ . We vary the number of splits allowed for *Heuristic* from 0 to 10. Each bar represents the average of 95 different queries. Notice that in all cases, our algorithms outperform *Naive*, and that both the worst case performance and average performance of *Heuristic-10* is very close to the performance of *Exhaustive*. We tried this experiment with a variety of SPSF's and data sets, and obtained similar results; space constraints preclude the inclusion of these graphs.

The space and time complexity of the exhaustive algorithm are very high. On a 2.4GHz Pentium 4 with 1 gigabyte of RAM, the largest problems we could solve were still several orders of magnitude smaller than the smallest of our real-world data sets. We discuss scalability results for all of our algorithms in Section 6.4.

Finally, we compare the performance *Exhaustive* against that

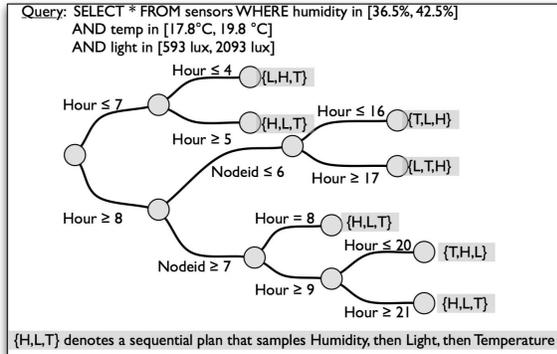


Figure 9: A conditional plan generated by the heuristic algorithm for a query over our lab data set looking for instances that are bright, cool, and dry.

of *Heuristic* for varying SPSF’s. Figure 8(b) shows the results of this experiment, where we compare *Heuristic-5* with a SPSF of  $10^{14}$  to *Exhaustive* with varying SPSF’s. As before, each bar represents the average (or maximum) of 95 trials. Notice that *Exhaustive* with smaller SPSF’s performs substantially worse than *Heuristic* with large SPSF’s.

In general, as these results show, constraining the split point selection too much is probably not appropriate, as it tends to obscure interesting correlations in the data by grouping together parts of the attribute space that may be partially correlated with other attributes, diluting the benefit our algorithms can obtain by exploiting such correlation.

### 6.1.1 Detailed Plan Study

Figure 9 illustrates an example of a real plan produced by our system for one of our test queries from the lab data set. The total performance gain for this plan is about 20% over *Naive*. In this case, the query looks for sensors that are reading cool temperatures, and relatively high light intensities – indicating, perhaps, that someone is working in the lab at night when it is typically cold and dark. None of the predicates in this query have a particularly low marginal selectivity, but the total selectivity of the query will be low, as there are few cases when it is both cold and bright in our lab. Notice that in this case, our algorithm first conditions on the hour in the day: when it is morning (hours 0 - 6), it prefers to sample light first, as very early in the morning the lab is unused and it is dark, so this predicate will fail. In the afternoon, additional conditional predicates on nodeid are introduced. The split at  $\text{nodeid} \leq 6$  represents a group of sensors (1-6) in a common part of the lab that is not used at night, so darkness is highly correlated with time of day. At the other sensors ( $\text{nodeid} \geq 7$ ), the lab is sometimes in use until late into the night. Thus, in this part of the lab, light may not be correlated with time of day. Interestingly, it appears that humidity is correlated with time of day, as the generated plan samples humidity first late at night. This is likely because the temperature control system in our building is turned off at night, and air conditioning and heating tends to keep the humidity low.

## 6.2 Garden Dataset

For the garden dataset, we generated two sets of queries containing identical range predicates over *temperature* and *humidity* over all notes. The predicates used were generated as follows:

- $a \leq \text{temperature}/\text{humidity} \leq b$ , where the range  $\langle a, b \rangle$  was selected randomly so as to cover a specified fraction of the domain size (which was varied between 1.25 to 3.25).
- $\text{not}(a \leq \text{temperature}/\text{humidity} \leq b)$ .

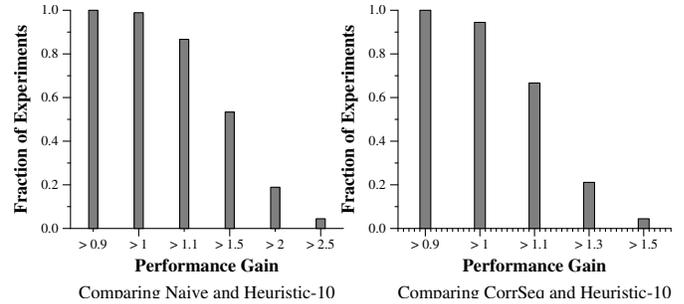


Figure 10: Results for the Garden-5 Dataset

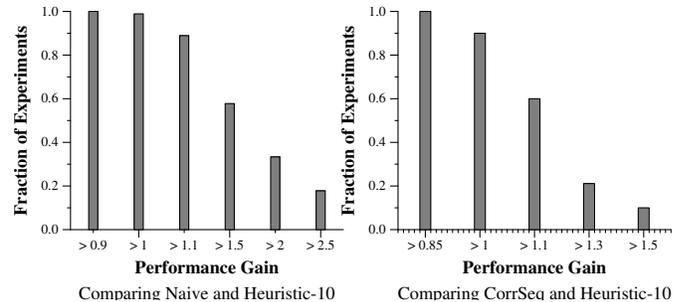


Figure 11: Results for the Garden-11 Dataset

The SPSF for *Heuristic* was set to  $10^n$ , where  $n$  is the number of attributes in the dataset.

Figure 10 shows the results of running 90 queries generated as above on the Garden-5 dataset. The queries in this case consist of 10 predicates each. We plot two graphs, comparing *Heuristic* to both *Naive* and *CorrSeq*. As we can see, *Heuristic* performs significantly better than both *Naive* and *CorrSeq* for a large fraction of queries. Though a *Heuristic* plan will never be worse than *Naive* or *CorrSeq* over the training data, because of the minor differences between the probability distributions of the test and the training datasets, for some of the queries, *Heuristic* actually performs slightly worse than (less than 10%) the other algorithms; as we can see, the penalty in those cases is negligible, whereas the gains for the rest of queries are significantly higher.

Figure 11 shows a similar plot for the Garden-11 dataset (with queries consisting of 22 predicates each). The performance improvement is even more significant in this case, with a factor of 4 improvement over *Naive* for some of the queries.

## 6.3 Synthetic Dataset

Finally, we present results over the synthetic dataset adapted from [2] for four settings of the parameters: (1)  $\Gamma = 1, n = 10$ , (2)  $\Gamma = 3, n = 10$ , (3)  $\Gamma = 1, n = 40$ , and (4)  $\Gamma = 3, n = 40$ . The queries (generated as described above) contain 5, 7, 20, and 30 predicates respectively. We plot the execution costs of running the plans generated by our algorithms vs the unconditional selectivity (*sel*) of the predicates. For *Heuristic*, we present results with at most 5 branches, and with at most 10 branches (*Heuristic-5*, and *Heuristic-10*).

As we can see in Figure 12, in all cases, conditional planning offers significant performance improvements over the plans generated by both *Naive* and *CorrSeq*, in several cases by more than a factor of 2. Note that when  $\Gamma = 1$ , *Naive* and *CorrSeq* produce nearly identical query plans.

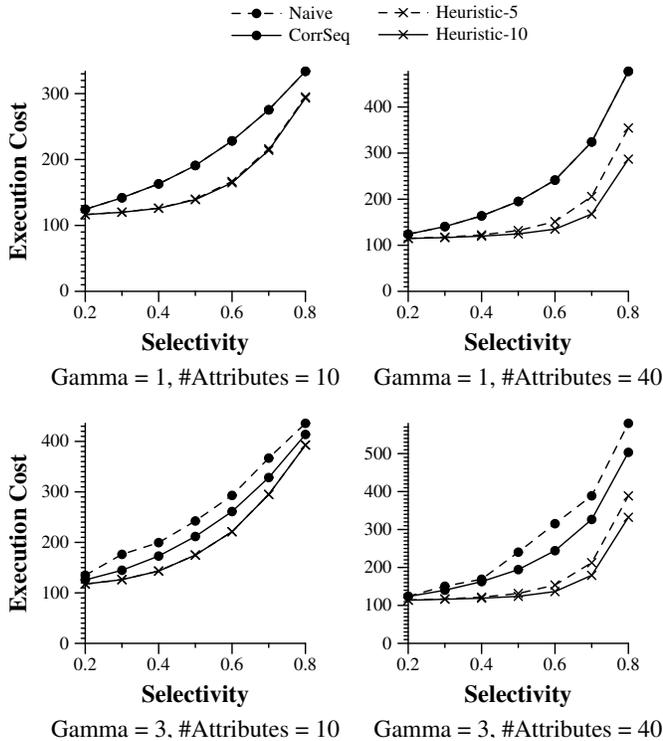


Figure 12: Results for Synthetic Dataset for various parameter settings. Here, *CorrSeq* and *Naive* overlap completely when  $\Gamma = 1$ , whereas *Heuristic-5* and *Heuristic-10* have very similar performance when number of attributes is 10.

## 6.4 Scalability

We also ran experiments studying the scaling of the exhaustive and heuristic algorithms versus number of predicates, attribute domain sizes, and the amount of historical data. Recall from our discussion above that we expect the performance of our heuristic algorithm to scale linearly in the size of the data set, linearly with domain size, and exponentially (base 2) in the number of query variables. The exhaustive algorithm is also linear in the size of the data set, but is polynomial in the size of the domain size and exponential in the number of query variables, where the base of the exponent is the domain size. Our scalability experiments verify that our implementation obeys the complexity bounds given above. Due to space limitations, we omit these experiments from this paper.

## 7. Applications and extensions

This paper focuses on building the basic data-structures and algorithms required for obtaining good conditional plans for complex acquisitional queries. Next, we outline several extensions built on this basic framework that we are planning to pursue in future.

**Approximate answers:** In a related work, we have proposed an approach to exploit cross-mote attribute correlations in sensor network query processing to efficiently answer probabilistic queries [9]. The algorithms developed in that paper, however, only generate *linear* plans; we are planning to explore how conditional plans can be used instead to avoid unnecessary acquisitions in that scenario as well.

**Complex acquisition costs:** We have focused on a very simple cost function: the independent cost of acquiring a single variable. In general, however, we may have significantly more complex cost functions. For example, motes have sensor boards with multiple sensors that are powered up simultane-

ously. Thus, the cost of acquiring a reading can be decomposed as the high cost of powering up the board, plus a low cost for a reading of each sensor in the board. This can be simulated in our planning algorithms by making the costs of acquiring attributes themselves conditional on the attributes acquired so far. Acquisition costs (latencies) in web querying may also vary depending on the time of the day and other phenomenon. How to capture such cost functions, and how to use them efficiently, is an interesting open question.

Another scenario where complex acquisition costs arise is when a query refers to attributes that are themselves distributed (e.g., if we are querying the state of the sensor network as a whole). The acquisition costs of the attributes vary drastically depending on which node is currently executing a predicate. The dynamic nature of sensor networks poses interesting challenges in estimating and using such cost functions.

**Graphical Models:** The methods presented in the previous section, that use historical data to estimate conditional probabilities, suffer from two issues: First, these algorithms are linear in the size of the dataset, which can be very large. Second, after each split on a predicate, each subproblem will be consistent with at most half of the data. Thus, the amount of data available to estimate probabilities decreases exponentially with the number of splits. After several splits, our probability estimates will thus have very high variance. This can result in choosing arbitrary plans, that may turn out to be significantly worse in reality than on the training data.

An alternative is to build a compact model of the data. *Probabilistic graphical models* provide a compact representation of complex, exponentially-large distributions, by exploiting conditional independencies between attributes. These models offer two significant advantages: First, there are several algorithms that allow us to compute the conditional probabilities efficiently by exploiting structure in the graphical model [6]. Second, by exploiting conditional independencies, we can often represent the required joint distribution with a polynomial number of parameters. Thus, this representation is significantly less susceptible to the overfitting observed when estimating probabilities from a dataset directly. Due to lack of space, we refer the reader to the book by Cowell *et al.* [6] for details.

**Query processing in other environments:** We have largely focused on the application of conditional plans in sensor networks. They are equally applicable in other wide-area environments; for example, on the web, the latency to acquire individual data items can be quite high, and the data may exhibit correlations that can be exploited using conditional plans. Similarly, in compressed databases [4], the cost of acquiring attributes may include the cost of decompression, which can be very high. Conditional plans can reduce the amount of decompression required to execute a query.

Our techniques can also be applied to traditional database query processing. For example, our techniques generalize easily to star queries containing only key-foreign key join predicates, can be thought of as expensive “selections” on the relation at the center of the star (commonly referred to as the *fact* table), and conditional plans can be used to exploit correlations between the *dimension* tables.

**Generalization to other types of queries:** Thus far, we have focused our attention on conjunctive queries. Other interesting types of queries in sensor network queries are existential queries, or queries that use the “LIMIT” clause. For example, we may only be interested in finding out if there exists a sensor that is recording high values of light and temperature. We can use conditional plans to significantly reduce the number of acquisitions

made by determining which of the sensors are most likely to satisfy the predicates.

**Queries over data streams:** In many practical settings, we may be performing a query over a continuous stream of data. If the data distribution does not change over time, we can use conditional plans as described in this paper to evaluate such queries. However, in many settings, the data distribution may change slowly over time. In such cases, we can modify our algorithms to slowly change the plan to adapt to the changing distribution. Specifically, our methods for computing probabilities from a data set in Section 5 can be modified to compute probabilities incrementally over a sliding window of data. As the probabilities change, we can modify our greedy algorithm to re-evaluate the plan, and consider (greedy) modifications. Such adaptive approach could efficiently tackle fast, continuous streams of data.

## 8. Related work

Our idea of conditional plans is quite similar to *parametric query optimization* [17, 11, 5, 10], where part of the query optimization process is postponed until the runtime. Typically, these techniques choose a *set* of query plans at query optimization, and identify a set of conditions that are used to select one of those plans at runtime. This earlier work differed substantially from ours in two essential ways: First, in these traditional approaches, the plan chosen at the runtime is used for executing the query over the entire dataset; thus, even if correlations were taken into account by these approaches, per-tuple variations, which we have seen to be prevalent and widely exploitable, could not be accounted for. Secondly, these approaches did not exploit data correlations while generating the plans.

Adaptive query processing techniques [14] attempt to reoptimize query execution plans during query execution itself. We believe that the idea of conditional plans that we propose is both orthogonal and complementary to adaptive query processing. If sufficiently accurate information about the data is available (as we assume in this work), then conditional plans can reap many of the benefits of adaptive query processing techniques *a priori* (by choosing different query plans for different parts of data). However, in many cases, such information may not be available, and adaptive techniques must be used. [2] address the problem of adaptively ordering *pipelined filters* (*et al.*, selection predicates) that may have correlations. Their focus is on finding good *sequential* plans (that may change with time), and they do not consider conditional plans.

Earlier work on expensive predicates [13, 3, 2, 20] talks about how to optimize queries with expensive predicates. All these techniques however produce a single sequential plan in the end. Shivakumar *et al.*, [23] propose using low-cost predicates to avoid evaluating expensive predicates. Their approach also constructs a sequential plan in the end, and the final query output may contain false positives, or may miss certain answers. Our approach, on the other hand, guarantees correct execution of the original query in all cases.

In prior work [19], we propose the idea of acquisitional query processing where the cost of acquiring attributes is explicitly modeled, though our focus there was entirely on sensor network query processing. In this paper, we have generalized this basic idea, and have proposed an approach to significantly speed up query processing in such environments. Web querying is another domain where the idea of acquisitional query processing, and the techniques we propose in this paper, can be very useful. Chen *et al.*, [4] propose techniques to perform query optimization in compressed database systems, and also model the cost of acquiring attributes explicitly. The techniques we propose in this paper can be extended in straightforward manner to their scenario.

## 9. Conclusions

In this paper, we showed how to exploit correlations between attributes in a database system by modifying the query optimizer to produce *conditional* plans that significantly outperform plans produced by traditional database optimizers. We showed specifically how these correlations can be used to optimize the performance of multi-predicate selection queries with attributes that have high acquisition costs [19] which frequently occur in distributed systems such as sensor networks and the Internet. We developed planning algorithms for generating such conditional plans given a historical data set with correlations. Our experimental results demonstrate the significant performance gains if the query optimizer is modified to take into account these correlations.

## 10. References

- [1] Planetlab. <http://www.planet-lab.org>.
- [2] S. Babu, R. Motwani, K. Munagala, I. Nishizawa, and J. Widom. Adaptive ordering of pipelined stream filters. In *SIGMOD*, 2004.
- [3] S. Chaudhuri and K. Shim. Optimization of queries with user-defined predicates. *TODS*, 24(2):177–228, 1999.
- [4] Z. Chen, J. Gehrke, and F. Korn. Query optimization in compressed database systems. In *ACM SIGMOD*, 2001.
- [5] R. Cole and G. Graefe. Optimization of dynamic query evaluation plans. In *SIGMOD*, 1994.
- [6] R. Cowell, P. Dawid, S. Lauritzen, and D. Spiegelhalter. *Probabilistic Networks and Expert Systems*. Springer, N.Y., 1999.
- [7] I. Crossbow. Wireless sensor networks (mica motes). [http://www.xbow.com/Products/Wireless\\_Sensor\\_Networks.htm](http://www.xbow.com/Products/Wireless_Sensor_Networks.htm).
- [8] A. Deshpande, C. Guestrin, W. Hong, and S. Madden. Exploiting correlated attributes in acquisitional query processing. Technical report, Intel-Research, Berkeley, 2004.
- [9] A. Deshpande, C. Guestrin, S. Madden, J. Hellerstein, and W. Hong. Model-driven data acquisition in sensor networks. In *VLDB*, 2004.
- [10] S. Ganguly. Design and analysis of parametric query optimization algorithms. In *VLDB*, 1998.
- [11] G. Graefe and K. Ward. Dynamic query evaluation plans. In *SIGMOD*, 1989.
- [12] R. Greiner, R. Hayward, and M. Molloy. Optimal depth-first strategies for and-or trees. In *AAAI/IAAI*, 2002.
- [13] J. M. Hellerstein. Optimization techniques for queries with expensive methods. *TODS*, 23(2):113–157, 1998.
- [14] J. M. Hellerstein, M. J. Franklin, S. Chandrasekaran, A. Deshpande, K. Hildrum, S. Madden, V. Raman, and M. Shah. Adaptive query processing: Technology in evolution. *IEEE Data Engineering Bulletin*, 23(2):7–18, 2000.
- [15] J. Hill, R. Szewczyk, A. Woo, S. Hollar, D. Culler, and K. Pister. System architecture directions for networked sensors. In *ASPLOS*, November 2000.
- [16] L. Hyafil and R. Rivest. Constructing optimal binary decision trees is np-complete. *IPL*, 1976.
- [17] Y. E. Ioannidis, R. T. Ng, K. Shim, and T. K. Sellis. Parametric query optimization. In *VLDB*, 1992.
- [18] R. Krishnamurthy, H. Boral, and C. Zaniolo. Optimization of nonrecursive queries. In *VLDB*, 1986.
- [19] S. Madden, M. J. Franklin, J. M. Hellerstein, and W. Hong. The design of an acquisitional query processor for sensor networks. In *ACM SIGMOD*, 2003.
- [20] K. Munagala, S. Babu, R. Motwani, and J. Widom. The pipelined set cover problem. In *ICDT*, 2005.
- [21] J. Polastre. Design and implementation of wireless sensor networks for habitat monitoring. Master’s thesis, UCB, 2003.
- [22] G. Pottie and W. Kaiser. Wireless integrated network sensors. *Communications of the ACM*, 43(5):51 – 58, May 2000.
- [23] N. Shivakumar, H. Garcia-Molina, and C. Chekuri. Filtering with approximate predicates. In *VLDB*, 1998.