

Comprehensive Optimization of Declarative Sensor Network Queries

Ixent Galpin, Christian Y.A. Brenninkmeijer, Farhana Jabeen,
Alvaro A.A. Fernandes, and Norman W. Paton

University of Manchester, United Kingdom
{ixent,brenninc,jabeen,alvaro,norm}@cs.man.ac.uk

Abstract. We present a sensor network query processing architecture that covers all the query optimization phases that are required to map a declarative query to executable code. The architecture is founded on the view that a sensor network truly is a distributed computing infrastructure, albeit a very constrained one. As such, we address the problem of how to develop a comprehensive optimizer for an expressive declarative continuous query language over acquisitional streams as one of finding extensions to classical distributed query processing techniques that contend with the peculiarities of sensor networks as an environment for distributed computing.

1 Introduction

This paper addresses the problem of optimizing the evaluation of declarative queries over sensor networks (SNs) [1]. Throughout, by *sensor networks* we mean ad-hoc, wireless networks whose nodes are energy-constrained sensors endowed with general-purpose computing capability. We believe there is broad, continued interest [2,3,4] in exploring whether techniques developed in the context of classical database environments, such as query optimization, are also applicable and beneficial in non-classical ones, of which SNs are a comparatively recent example. Viewed as a distributed computing infrastructure, SNs are constrained to an unprecedented extent, and it is from such constraints that the challenges we have addressed arise. Addressing these challenges is important because of the prevailing expectations for the wide applicability of SNs [5]. These expectations imply that, as an infrastructural platform, SNs are bound to become more heterogeneous over time than previous work has assumed. Moreover, integrating SN data with data stemming from other networks (sensor-based or not) will be easier if the query processing technology used is less tied to specific execution environments.

We explore the hypothesis that the classical *two-phase optimization* approach [6] from distributed query processing (DQP) can be adapted to be effective and efficient over SNs, as initially proposed in [7]. Two-phase optimization is well established in the case of robust networks (e.g., the Internet, or the interconnect of a parallel machine), and involves the decomposition of query optimization

into a *Single-Site* phase, and a subsequent *Multi-Site* phase, each of which is further decomposed into finer-grained decision-making steps. We aim to reuse well established optimization components where possible, and to identify steps that are necessarily different in SNs. We demonstrate that extending a classical DQP optimization architecture allows an expressive query language to be supported over resource-constrained networks, and that the resulting optimization steps yield good performance, as demonstrated through empirical evaluation.

Related Work. There have been many proposals in the so-called SN-as-database approach (including [2,8,3,4]). Surprisingly, none have fully described an approach to query optimization founded on a classical DQP architecture. Cougar papers [8] propose this idea but no publication describes its realization. SNQL [9] follows the idea through but no precise description (as provided by our algorithms) of the decision-making process has been published. Indeed, few publications provide systematic descriptions of complete query optimization architectures for SN query processors: the most comprehensive description found was for TinyDB [3], in which optimization is limited to operator reordering and the use of cost models to determine an appropriate acquisition rate given a user-specified lifetime. Arguably as a result of this, SN-as-database proposals have tended to limit the expressiveness of the query language. For example, TinyDB focuses on aggregation and provides limited support for joins. In many cases, assumptions are made that constrain the generality of the approach (e.g., Presto [2] focuses on storage-rich networks).

There has also been a tendency to address the optimization problem in a piecewise manner. For example, the use of probabilistic techniques to address the trade-off between acquiring data often and the cost of doing so is proposed in BBQ [10]; the trade-off between energy consumption and time-to-delivery is studied in WaveScheduling [11]; efficient and robust aggregation is the focus of several publications [12,13,14]; Bonfils [15] proposes a cost-based approach to adaptively placing a join which operates over distributed streams; Zadorozhny [16] uses an algebraic approach to generate schedules for the transmission of data in order to maximize the number of concurrent communications. However, these individual results are rarely presented as part of a fully-characterized optimization and evaluation infrastructure, giving rise to a situation in which research at the architecture level seems less well developed than that of techniques that might ultimately be applied within such query processing architectures.

This paper aims to provide a comprehensive, top-to-bottom approach to the optimization problem for expressive declarative continuous queries over potentially heterogeneous SNs. In comparison with past proposals, ours is broader, in that there are fewer compromises with respect to generality and expressiveness, and more holistic, in that it provides a top-to-bottom decomposition of the decision-making steps required to optimize a declarative query into a query execution plan (QEP).

Approach and Main Results. A key aspect of our approach is that it supports the optimization and evaluation of SNEEqL (for Sensor Network Engine query language), a comprehensive stream query language, inspired by classical stream languages such as CQL [17]. Thus, we do not start from the assumption that a query language for use with resource-constrained devices must, as a consequence, provide limited functionality. We then take a classical DQP optimizer architecture as a starting point, and adapt it to consider how the optimization and evaluation of queries for SNEEqL differs if, rather than targeting classical distributed computational resources, we target SNs. We identify what assumptions in classical DQP are violated in the SN case, and propagate the consequences of such violations into the design and engineering of a DQP architecture for SN data management. The differences we consider here that have led to adaptations and extensions, are: (i) *the acquisitional nature of the query processing task*: data is neither lying ready in stores nor is it pushed (as in classical streams), but requested; (ii) *the energy-constrained nature of the sensors*: preserving energy becomes a crucial requirement because it is a major determinant of network longevity, and requires the executing code to shun energy-hungry tasks; (iii) *the fundamentally different nature of the communication links*: wireless links are not robust, and often cannot span the desired distances, so the data flow topology (e.g., as embodied in a query operator tree) needs to be overlaid onto some query-specific network topology (e.g., a routing tree of radio-level links) for data to flow from sensors to clients, and the two trees are not isomorphic (e.g., some network nodes act as pure relay nodes); and (iv) *the need to run sensor nodes according to data-dependent duty cycles*: each element in the computational fabric must act in accordance with an agenda that co-ordinates its activity with that of other elements on which it is dependent or that depend on it, thereby enabling energy management (e.g., by sending devices to energy-saving states until the next activity). We note that these points also distinguish our approach from infrastructures for stream query processing (e.g., [17]), which do not operate in resource constrained environments.

Summary of Contributions. The body of the paper describes: (1) a user-level syntax (Section 2) and algebra (Section 3.1) for SNEEqL, an expressive language for querying over acquisitional sensor streams; (2) an architecture for the optimization of SNEEqL, building on well-established DQP components where possible, but making enhancements or refinements where necessary to accommodate the SN context (Section 3); (3) algorithms that instantiate the components, thereby supporting integrated query planning that includes routing, placement and timing (Section 3.2); and (4) an evaluation of the resulting infrastructure, demonstrating the impact of design and optimization decisions on query performance, and the increased empowering to the user (Section 4) who is able to trade-off different Qualities-of-Service (Qos) according to application needs.

2 Query Language

SNEEqL [18] is a declarative query language for SNs inspired by expressive classical stream query languages such as CQL [17]. A rich language is used even

```

Schema: outflow (id, time, temp, turbidity, pressure)          sources: {0, 2, 4}
        inflow  (id, time, temp, pressure, ph)                sources: {4, 5, 7}

Q1: SELECT RSTREAM id, pressure
     FROM inflow[NOW]
     WHERE pressure > 500;

Q2: SELECT RSTREAM AVG(pressure)
     FROM inflow[NOW]

Q3: SELECT RSTREAM o.id, i.id, o.time, o.pressure, i.pressure
     FROM outflow[NOW] o, inflow[FROM NOW - 1 TO NOW - 1 MINUTES] i
     WHERE o.pressure < i.pressure AND i.pressure > 500

QoS: {ACQUISITION RATE = 3s ; DELIVERY TIME = 5s}

```

Fig. 1. Schema Metadata, Example queries in SNEEqI and QoS Expectations

though our target delivery platform consists of limited devices because: (i) the results of queries written using inexpressive query languages may require off-line post-processing over data that has to be delivered using costly wireless communication; and (ii) sensor applications require comprehensive facilities for correlating data sensed in different locations at different times (e.g., [19,20]).

An example schema and queries are given in Fig. 1 motivated by the application scenario (but not actually occurring) in PipeNet, “a system based on wireless SNs [that] aims to detect, localize and quantify bursts and leaks and other anomalies in water transmission pipelines” [20]. The *logical extents* in the schema (i.e., `inflow` and `outflow`) comprise a (possibly overlapping) subset of the acquisitional streams generated by the source nodes, as specified in Fig. 1. The topology of the network is depicted in Fig. 2. *Q1* requests the tuples from `inflow` whose pressure attribute is above a certain threshold; *Q2* requests the average value of the pressure readings in the `inflow` logical extent; and *Q3* obtains instances in which the `outflow` pressure is less than the `inflow` pressure a minute before (as long as the latter was above a certain threshold). All acquire data and return results every 3 seconds, as stated in the QoS parameters. *Q3* returns the `id` corresponding to the `inflow` and `outflow` source nodes, to assist the user with locating a potential leak. The examples illustrate how SNEEqI can express select-project (*Q1*), aggregation (*Q2*) and join (*Q3*) queries. *Q3* is noteworthy as it correlates data from different locations at different times, and cannot be expressed with previous SN query languages.

In SNEEqI, the only structured type is `tuple`. The primitive collection types in SNEEqI are: `relation`, an instance of which is a bag of tuples with definite cardinality; `window`, an instance of which is a relation whose content may implicitly vary between evaluation episodes; and `stream`, an instance of which is a bag of tuples with indefinite cardinality whose content may implicitly vary throughout query evaluation. As in CQL, operations construct windows out of streams and vice-versa. In all the queries, windows are used to convert from streams to relations, relational operators act on those relations, and stream operators add the resulting tuples into the output stream. Window definitions are of the form *WindowDimension* [SLIDE] [Units], where the *WindowDimension* is of the form NOW or FROM *Start* TO *End*, where the former contains all the tuples with the current time stamp, and the latter contains all the tuples that

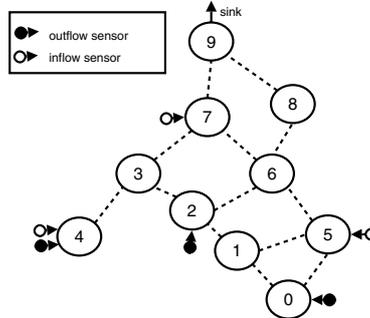


Fig. 2. Connectivities and Modalities

fall within the given range. The *Start* and *End* of a range are of the form *NOW* or *NOW -Literal*, where the *Literal* represents some number of *Units*, which is either *ROWS* or a time unit (*HOURS*, *MINUTES* or *SECONDS*). The optional *SLIDE* indicates the gap in *Units* between the *Start* of successive windows; this defaults to the acquisition rate specified. The results of relational operators are added to the result stream using the relation-to-stream operators from CQL, namely *ISTREAM*, *DSTREAM* and *RSTREAM*, denoting *inserted-only*, *deleted-only* and *all-tuples*, respectively.

In a stream query language, where conceptually data is being consumed, and thus potentially produced, on an ongoing basis, the question exists as to when a query should be evaluated. In SNEEqL, an *evaluation episode* is determined by the acquisition of data, i.e., setting an acquisition rate sets the rate at which evaluation episodes occur. Thus, whenever new data is acquired, it is possible that new results can be derived. In the implementation, however, partial query results may be cached within the network with a view to minimizing network traffic, and different parts of a query may be evaluated at different times, reflecting related, but distinct, QoS expectations, viz., the *acquisition rate* and the *delivery time*, as shown in Fig. 1. Noteworthy features of SNEEqL illustrated in Fig. 1 include: (i) extending CQL stream-to-window capabilities to allow for windows whose endpoint is earlier than now or than the last tuple seen, as happens in *Q3* with the window on *inflow*, which emits tuples that were acquired one minute ago; (ii) allowing sensing capabilities to be logically abstracted in a schema (like Cougar [8], but unlike TinyDB, which assumes that all sensed data comes from a single extent denoted by the keyword *SENSORS*); (iii) allowing the logical streams to stem from more than one set of sensor nodes, and possibly intersecting ones (as is the case in Fig. 1); (iv) expressing joins, where the tuples come from different locations in the SN, as a single query and without using materialization points (as would be required in TinyDB); and (v) allowing QoS expectations to be set for the optimizer, such as acquisition rate and delivery time. A detailed description of the SNEEqL language, including a formal semantics, is given in [18].

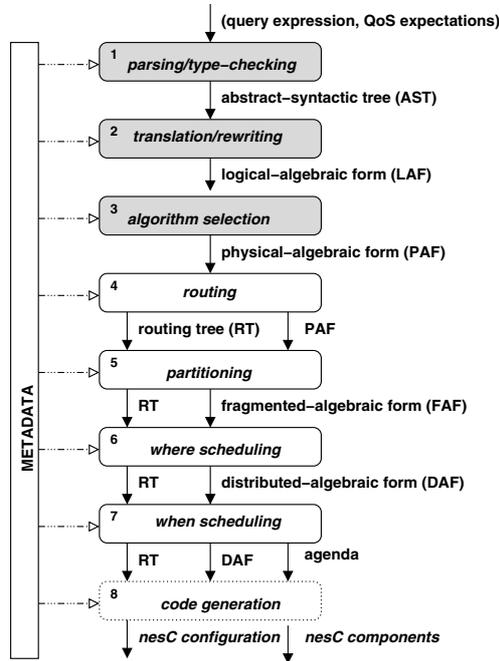


Fig. 3. SNEEQl Compiler/Optimizer Stack

3 Query Compiler/Optimizer

Recall that our goal is to explore the hypothesis that extensions to a classical DQP optimization architecture can provide effective and efficient query processing over SNs. The SNEEQl compilation/optimization stack is illustrated in Fig. 3, and comprises three phases. The first two are similar to those familiar from the two phase-optimization approach, namely *Single-Site* (comprising Steps 1-3, in gray boxes) and *Multi-Site* (comprising steps 4-7, in white, solid boxes). The *Code Generation* phase grounds the execution on the concrete software and hardware platforms available in the network/computing fabric and is performed in a single step, Step 8 (in a white, dashed box), which generates executable code for each site based on the distributed QEP, routing tree and agenda. The optimizer consists of 15K lines of Java.

Classical kinds of metadata (e.g., schematic information, statistics, etc.) are used by the SNEEQl optimizer, but we focus here on the requirement for a more detailed description of the network fabric. Thus, Fig. 2 depicts an example network for the case study in Fig. 1 in terms of a weighted connectivity graph and the sensing modalities available in each site. Dotted edges denote that single-hop communication is possible. Weights normally model energy, but more complex weights could be used. Here, we assume links to have unit costs, but this need not be so. Note that all sites may contribute computing (e.g., processing intermediate results) or communication (e.g., relaying data) capabilities. Metadata

is assumed to be populated by a network management layer, such as Moteworks (<http://www.xbow.com>), which also provides services such as network formation, self-healing, code dissemination etc. Note, however, we have not actually used Moteworks since it does not support simulation, and we have opted to simulate in order to efficiently carry out systematic experimentation.

Throughout the steps in the query stack, size-, memory-, energy- and time-cost models are used by the optimizer to reach motivated decisions. These cannot be described in detail due to space limitations. The cost models are very close in style and level of detail to those in [21], but have been extended and adapted for the SN context.

3.1 Phase 1: Single-Site Optimization

Single-site optimization is decomposed into components that are familiar from classical, centralized query optimizers. We make no specific claims regarding the novelty of these steps, since the techniques used to implement them are well-established. In essence: *Step 1* checks the validity of the query with respect to syntax and the use of types, and builds an abstract syntax tree to represent the query; *Step 2* translates the abstract syntax tree into a logical algebra, the operators of which are reordered to reduce the size of intermediate results; and *Step 3* translates the logical algebra into a physical algebra, which, e.g., makes explicit the algorithms used to implement the operators. Fig. 6 depicts the outcome of Steps 1 to 3 for $Q3$ from Fig. 1, expressed using SNEEqI *physical-algebraic form* (PAF) that is the principal input to multi-site optimization.

Table 1 describes the PAF operators, grouped by their respective input-to-output collection types. A signature has the form OPERATOR_NAME[Parameters](InputArgumentTypes):OutputArgumentTypes, where the argument types are denoted R , S and W , for relation, stream and window respectively, and a vertical bar indicates a choice of one of the types given. ACQUIRE and DELIVER denote data sources and sinks, respectively. The window on `inflow` is represented in the algebra in milliseconds as `TIME_WINDOW[t-60000,t-60000,30000]`, and is relative to `t`, which is bound in turn to the time in which each evaluation episode of the query starts. Note that the ACQUIRE and DELIVER are both *location sensitive*, i.e., there is no leeway as to which node(s) in the SN they may execute on. Furthermore, NL_JOIN is *attribute sensitive*, i.e., in order to carry out partitioned-parallelism the optimizer needs to consider how the input(s) are partitioned in order to preserve operator semantics.

3.2 Phase 2: Multi-site Optimization

For distributed execution, the physical-algebraic form (PAF) is broken up into QEP fragments for evaluation on specific nodes in the network. In a SN, consideration must also be given to routing (the means by which data travels between nodes within the network) and duty cycling (when nodes transition from being switched on and engaged in specific tasks, and being asleep, or in power-saving

Table 1. SNEEq1 Physical Algebra

Stream-to-Stream Operators	
ACQUIRE[AcquisitionInterval, PredExpr, AttrList](S) : S	Take sensor readings every AcquisitionInterval for sources in S and apply SELECT[PredExpr] and PROJECT[AttrList]. <i>LocSen</i> .
DELIVER[](S) : S	Deliver the query results. <i>LocSen</i> .
Stream-to-Window Operators	
TIME_WINDOW[startTime, endTime, Slide](S) : W	Define a time-based window on S from startTime to endTime inclusive and re-evaluate every slide time units.
ROW_WINDOW[startRow, endRow, Slide](S) : W	Define a tuple-based window on S from startRow to endRow inclusive and re-evaluate every slide rows. <i>AttrSen</i> .
Window-to-Stream Operators	
RSTREAM[](W) : S	Emit all the tuples in W .
ISTREAM[](W) : S	Emit the newly-inserted tuples in W since the previous window evaluation.
DSTREAM[](W) : S	Emit the newly-deleted tuples in W since the previous window evaluation.
Window-to-Window or Relation-to-Relation Operators	
NL_JOIN[ProjectList, Expr]($R W, R W$) : $R W$	Pred- Join tuples on PredExpr condition using nested-loop algorithm. <i>AttrSen</i> .
AGGR_INIT[AggrFunction, AttrList]($R W$) : $R W$	Initialize incremental aggregation for attributes in AttrList for type of aggregation specified by AggrFunction. <i>AttrSen</i> .
AGGR_MERGE[AggrFunction, AttrList]($R W$) : $R W$	Merge partial-result tuples of incremental aggregation for attributes in AttrList for type of aggregation specified by AggrFunction. <i>AttrSen</i> .
AGGR_EVAL[AggrFunction, AttrList]($R W$) : $R W$	Evaluate final result of incremental aggregation for attributes in AttrList for type of aggregation specified by AggrFunction. <i>AttrSen</i> .
Any-to-Same-as-input-type Operators	
SELECT[PredExpr]($R S W$) : $R S W$	Eliminate tuples which do not meet PredExpr predicate.
PROJECT[AttrList] ($R S W$) : $R S W$	Generate tuple with AttrList attributes.

modes). Therefore, for Steps 4-7, we consider the case of robust networks and the contrasting case of SNs.

For execution over multiple nodes in robust networks, the second phase is comparatively simple: one step partitions the PAF into fragments and another step allocates them to suitably resourced sites, as in, e.g., [22]. One approach to achieving this is to map the physical-algebraic form of a query to a distributed one in which EXCHANGE operators [23] define boundaries between fragments. An EXCHANGE operator encapsulates all of control flow, data distribution and inter-process communication and is implemented in two parts, referred to as **producer** and **consumer**. The former is the root operator of the upstream fragment, and the latter, a leaf operator of the downstream one. This approach has been successful in DQP engines for the Grid that we developed in previous work [24,25].

However, for the same general approach to be effective and efficient in a SN, a response is needed to the fact that assumptions that are natural in the robust-network setting cease to hold in the new setting and give rise to a different set of challenges, the most important among which are the following: **C1**: *location and time are both concrete*: acquisitional query processing is grounded on the physical world, so sources are located and timed in concrete space and time, and the optimizer may need to respond to the underlying geometry and to synchronization issues; **C2**: *resources are severely bounded*: sensor nodes can be depleted of energy, which may, in turn, render the network useless; **C3**: *communication events are overly expensive*: they have energy unit costs that are typically an order of magnitude larger than the comparable cost for computing and sensing events; and **C4**: *there is a high cost in keeping nodes active for long periods*: because of the need to conserve energy, sensor node components must run tight duty cycles (e.g., going to sleep as soon they become idle).

Our response to this different set of circumstances is reflected in Steps 4-7 in Fig. 3, where rather than a simple partition-then-allocate approach (in which a QEP is first partitioned into fragments, and these fragments are then allocated to specific nodes on the network), we: (a) introduce Step 4, in which the optimizer determines a routing tree for communication links that the data flows in the operator tree can then rely on, with the aim of addressing the issue that paths used by data flows in a query plan can greatly impact energy consumption (a consequence of **C3**); (b) preserve the query plan partitioning step, albeit with different decision criteria, which reflect issues raised by **C1**; (c) preserve the scheduling step (which we rename to *where-scheduling*, to distinguish it from Step 7), in which the decision is taken as to where to place fragment instances in concretely-located sites (e.g., some costs may depend on the geometry of the SN, a consequence of **C1**); and (d) introduce *when-scheduling*, the decision as to when, in concrete time, a fragment instance placed at a site is to be evaluated (and queries being continuous, there are typically many such episodes) to address **C1** and **C4**. **C2** is taken into account in changes throughout the multi-site phase.

For each of the following subsections that describe Steps 4 to 7, we indicate how the proposed technique relates to DQP and to TinyDB, the former because we have used established DQP architectures as our starting point, and the latter because it is the most fully characterized proposal for a SN query processing system. The following notation is used throughout the remainder of this section. Given a query Q , let P_Q denote the graph-representation of the query in physical-algebraic form. Throughout, we assume that: (1) operators (and fragments) are described by properties whose values can be obtained by traditional accessor functions written in dot notation (e.g., P_Q .Sources returns the set of sources in P_Q); and (2) the data structures we use (e.g., sets, graphs, tuples) have functions with intuitive semantics defined on them, written in applicative notation (e.g., for a set S , $\text{ChooseOne}(S)$ returns any $s \in S$; for a graph G , $\text{EdgesIn}(G)$ returns the edges in G); $\text{Insert}((v_1, v_2), G)$ inserts the edge (v_1, v_2) in G .

Routing. Step 4 in Fig. 3 decides which sites to use for routing the tuples involved in evaluating P_Q . The aim is to generate a routing tree for P_Q which is economical with respect to the total energy cost required to transmit tuples. Let $G = (V, E, d)$ be the weighted connectivity graph for the target SN (e.g., the one in Fig. 2). Let $P_Q.\text{Sources} \subseteq G.V$ and $P_Q.\text{Destination} \in G.V$ denote, respectively, the set of sites that are data sources, and the destination site, in P_Q . The aim is, for each source site, to reduce the total cost to the destination. We observe that this is an instance of the *Steiner tree* problem, in which, given a graph, a tree of minimal cost is derived which connects a required set of nodes (the *Steiner nodes*) using any additional nodes which are necessary [1]. Thus, the SNEEqI-optimal routing tree R_Q for Q is the Steiner tree for G with Steiner nodes $P_Q.\text{Sources} \cup \{P_Q.\text{Destination}\}$. The problem of computing a Steiner tree is NP-complete, so the heuristic algorithm given in [1] (and reputed to perform well in practice) is used to compute an approximation. The resulting routing tree for $Q3$ over the network given in Fig. 2 is depicted in Fig. 7.

Relationship to DQP: The routing step has been introduced in the SN context due to the implications of the high cost of wireless communications, viz., that the paths used to route data between fragments in a query plan have a significant bearing on its cost. Traditionally, in DQP, the paths for communication are solely the concern of the network layer. In a sense, for SNEEqI, this is also a preparatory step to assist *where-scheduling* step, in that the routing tree imposes constraints on the data flows, and thus on where operations can be placed.

Relationship to Related Work: In TinyDB, routing tree formation is undertaken by a distributed, parent-selection protocol at runtime. Our approach aims, given the sites where location-sensitive operators need to be placed, to reduce the distance traveled by tuples. TinyDB does not directly consider the locations of data sources while forming its routing tree, whereas the approach taken here makes finer-grained decisions about which depletable resources (e.g., energy) to make use of in a query. This is useful, e.g., if energy stocks are consumed at different rates at different nodes.

Partitioning. Step 5 in Fig. 3 defines the fragmented form F_Q of P_Q by breaking up selected edges $(child, op) \in P_Q$ into a path $[(child, e_p), (e_c, op)]$ where e_p and e_c denote, respectively, the **producer** and **consumer** parts of an EXCHANGE operator. The edge selection criteria are semantic, in the case of location- or attribute-sensitive operators in which correctness criteria constrain placement, and pragmatic in the case of an operator whose output size is larger than that of its child(ren) in which case placement seeks to reduce overall network traffic. Let **Size** estimate the size of the output of an operator or fragment, or the total output size of a collection of operator or fragment siblings. The algorithm that computes F_Q is shown in Fig. 4. Fig. 8 depicts the distributed-algebraic form (the output of where-scheduling) given the routing tree in Fig. 7 for the physical-algebraic form in Fig. 6. The EXCHANGE operators that define the four fragments shown in Fig. 8 are placed by this step. The fragment identifier **F_n** denotes the fragment number. The assigned set of sites for each fragment (below the fragment identifier) are determined subsequently in where-scheduling.

```

FRAGMENT-DEFINITION( $P_Q$ , Size)
1  $F_Q \leftarrow P_Q$ 
2 while  $\triangleright$  post-order traversing  $F_Q$ ,
    $\triangleright$  let  $op$  denote the current operator
3   do for each  $child \in op.Children$ 
4     do if  $Size(op) > Size(op.Children)$  or  $op.LocationSensitive = \text{yes}$ 
5     or  $op.AttributeSensitive = \text{yes}$ 
6     then  $Delete((child, op), P_Q) ; Insert((child, e_p), P_Q)$ 
7          $Insert((e_p, e_c), P_Q) ; Insert((e_c, op), P_Q)$ 
8 return  $F_Q$ 

```

Fig. 4. The partitioning algorithm

EXCHANGE has been inserted between each ACQUIRE and the JOIN, because the predicate of the latter involves tuples from different sites, and therefore data redistribution is required. Note also that an EXCHANGE has been inserted below the DELIVER, because the latter is (as is ACQUIRE) *location sensitive*, i.e., there is no leeway as to where it may be placed.

Relationship to DQP: This step differs slightly from its counterpart in DQP. EXCHANGE operators are inserted more liberally at edges where a reduction in data flow will occur, so that radio transmissions take place along such edges whenever possible.

Relationship to Related Work: Unlike SNEEQ/DQP, TinyDB does not partition its query plans into fragments. The entire query plan is shipped to sites which are required to participate in it, even if they are just relaying data.

Where-Scheduling. Step 6 in Fig. 3 decides which QEP fragments are to run on which routing tree nodes. This results in the distributed-algebraic form of the query. Creation and placement of fragment instances is mostly determined by semantic constraints that arise from location sensitivity (in the case of ACQUIRE and DELIVER operators) and attribute sensitivity (in the case JOIN and aggregation operators, where tuples in the same logical extent may be traveling through different sites in the routing tree). Provided that location and attribute sensitivity are respected, the approach aims to assign fragment instances to sites, where a reduction in result size is predicted (so as to be economical with respect to the radio traffic generated).

Let G , P_Q and F_Q be as above. Let $R_Q = \text{ROUTING}(P_Q, G)$ be the routing tree computed for Q . The where-scheduling algorithm computes D_Q , i.e., the graph-representation of the query in distributed-algebraic form, by deciding on the creation and assignment of fragment instances in F_Q to sites in the routing tree R_Q . If the size of the output of a fragment is smaller than that of its child(ren) then it is assigned to the deepest possible site(s) (i.e., the one with the longest path to the root) in R_Q , otherwise it is assigned to the shallowest site for which there is available memory, ideally the root. The aim is to reduce radio traffic (by postponing the need to transmit the result with increased size). Semantic criteria dictate that if a fragment contains a location-sensitive operator, then instances of it are created and assigned to each corresponding site (i.e., one

that acts as source or destination in F_Q). Semantic criteria also dictate that if a fragment contains an attribute-sensitive operator, then an instance of it is created and assigned to what we refer to as a confluence site for the operator.

To grasp the notion of a *confluence site* in this context, note that the extent of one logical flow (i.e., the output of a logical operator) may comprise tuples that, in the routing tree, travel along different routes (because, ultimately, there may be more than one sensor feeding tuples into the same logical extent). In response to this, instances of the same fragment are created in different sites, in which case EXCHANGE operators take on the responsibility for data distribution among fragment instances (concomitantly with their responsibility for mediating communication events). It follows that a fragment instance containing an attribute-sensitive operator is said to be effectively-placed only at sites in which the logical extent of its operand(s) has been reconstituted by confluence. Such sites are referred to as confluence sites. For a JOIN, a confluence site is a site through which all tuples from both its operands travel. In the case of aggregation operators, which are broken up into three physical operators (viz., AGGR_INIT, AGGR_MERGE, AGGR_EVAL), the notion of a confluence site does not apply to an AGGR_INIT. For a binary AGGR_MERGE (such as for an AVG, where AGGR_MERGE updates a (SUM, COUNT) pair), a confluence site is a site that tuples from both its operands travel through. Finally, for an AGGR_EVAL, a confluence site is a site through which tuples from all corresponding AGGR_MERGE operators travel. The most efficient confluence site to which to assign a fragment instance is considered to be the deepest, as it is the earliest to be reached in the path to the destination and hence the most likely to reduce downstream traffic.

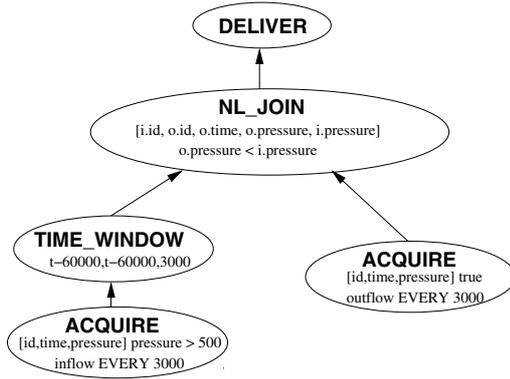
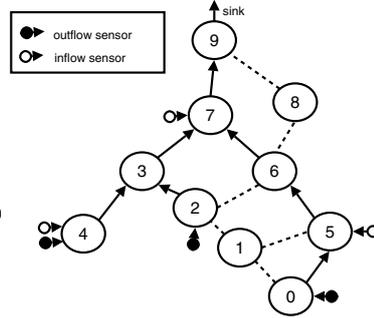
Let P_Q and R_Q be as above. Let $s \Delta op$ be true iff s is the deepest confluence site for op . The algorithm that computes D_Q is shown in Fig. 5. The resulting D_Q for the example query is shown in Fig. 8. It can be observed that instances of **F2** and **F3** have been created at multiple sites, as these fragments contain location-sensitive ACQUIRE operators, whose placement is dictated by the schema definition in Fig. 1. Also, a single instance of attribute-sensitive **F1** has been created and assigned to site 7, the deepest confluence site where tuples from both **F2** and **F3** are available (as it is a non-location-sensitive fragment and has been placed according to its expected output size, to reduce communication). Note also the absence of site 3 in Fig. 8 wrt. Fig. 7. This is because site 3 is only a relay node in the routing tree.

Relationship to DQP: Compared to DQP, here the allocation of fragments is constrained by the routing tree, and operator confluence constraints, which enables the optimizer to make well-informed decisions (based on network topology) about where to carry out work. In classical DQP, the optimizer does not have to consider the network topology, as this is abstracted away by the network protocols. As such, the corresponding focus of where-scheduling in DQP tends to be on finding sites with adequate resources (e.g., memory and bandwidth) available to provide the best response time (e.g., Mariposa [22]).

Relationship to Related Work: Our approach differs from that of TinyDB, since its QEP is never fragmented. In TinyDB, a node in the routing tree either

```

FRAGMENT-INSTANCE-ASSIGNMENT( $F_Q, R_Q, \text{Size}$ )
1  $D_Q \leftarrow F_Q$ 
2 while  $\triangleright$  post-order traversing  $D_Q$ 
    $\triangleright$  let  $f$  denote the current fragment
3   do if  $op \in f$  and  $op.\text{LocationSensitive} = \text{yes}$ 
4     then for each  $s \in op.\text{Sites}$ 
5       do Assign( $f.\text{New}, s, D_Q$ )
6     elseif  $op \in f$  and  $op.\text{AttributeSensitive} = \text{yes}$ 
7       and  $\text{Size}(f) < \text{Size}(f.\text{Children})$ 
8       then while  $\triangleright$  post-order traversing  $R_Q$ ,
9          $\triangleright$  let  $s$  denote the current site
10        do if  $s \Delta op$ 
11          then Assign( $f.\text{New}, s, D_Q$ )
12        elseif  $\text{Size}(f) < \text{Size}(f.\text{Children})$ 
13          then for each  $c \in f.\text{Children}$ 
14            do for each  $s \in c.\text{Sites}$ 
15              do Assign( $f.\text{New}, s, D_Q$ )
16        else Assign( $f.\text{New}, R_Q.\text{Root}, D_Q$ )
17 return  $D_Q$ 
    
```

 Fig. 5. The *where-scheduling* algorithm

 Fig. 6. Q_3 Physical-Algebraic Form

 Fig. 7. Q_3 Routing Tree

(i) evaluates the QEP, if the site has data sources applicable to the query, or (ii) restricts itself to relaying results to its parent from any child nodes that are evaluating the QEP. Our approach allows different, more specific workloads to be placed in different nodes. For example, unlike TinyDB, it is possible to compare results from different sites in a single query, as in Fig. 8. Furthermore, it is also possible to schedule different parts of the QEP to different sites on the basis of the resources (memory, energy or processing time) available at each site. The SNEEQ optimizer, therefore, responds to resource heterogeneity in the fabric. TinyDB responds to excessive workload by shedding tuples, replicating the strategy of stream processors (e.g., STREAM [17]). However, in SNs, since there is a high cost associated with transmitting tuples, load shedding is an undesirable option. As the query processor has control over data acquisition, it

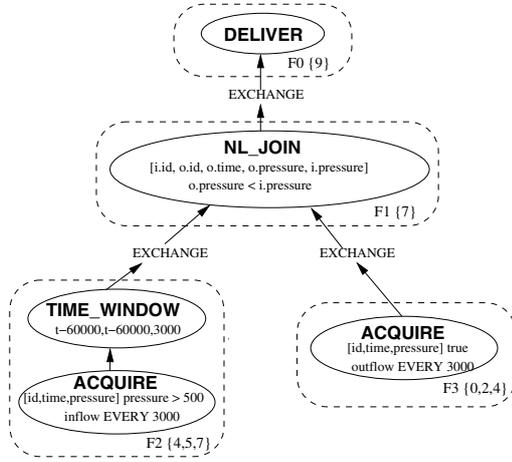


Fig. 8. Q3 Distributed-Algebraic Form

Time (ms)	Sites							
	0	5	6	2	4	3	7	9
0	F3 ₁	F2 ₁		F3 ₁	F2 ₁		F2 ₁	
63				F3 ₁				
3000	F3 ₂	F2 ₂		F3 ₂	F2 ₂		F2 ₂	
3032				tx3		rx2		
3063	tx5	rx0			F3 ₂			
3125					tx3	rx4		
3157		tx6	rx5					
3313			tx7				rx6	
3469						tx7	rx3	
3688							F1	
3719							tx9	rx7
4313								F0
4344								

Fig. 9. Q3 Agenda

seems more appropriate to tailor the optimization process so as to select plans that do not generate excess tuples in the first place.

When-Scheduling. Step 7 in Fig. 3 stipulates execution times for each fragment. Doing so efficiently is seldom a specific optimization goal in classical DQP. However, in SNs, the need to co-ordinate transmission and reception and to abide by severe energy constraints make it important to favor duty cycles in which the hardware spends most of its time in energy-saving states. The approach adopted by the SNEEQ compiler/optimizer to decide on the timed execution of each fragment instance at each site is to build an agenda that, insofar as permitted by the memory available at the site, and given the acquisition rate α and the delivery time δ set for the query, buffers as many results as possible before transmitting. The aim is to be economical with respect to both the time in which a site needs to be active and the amount of radio traffic that is generated.

The agenda is built by an iterative process of adjustment. Given the memory available at, and the memory requirements of the fragment instances assigned to, each site, a candidate buffering factor β is computed for each site. This candidate β is used, along with the acquisition rate α , to compute a candidate agenda. If the candidate agenda *makespan* (i.e., the time that the execution of the last task is due to be completed) exceeds the smallest of the delivery time δ and the product of α and β , the buffering factor is adjusted downwards and a new candidate agenda is computed. The process stops when the makespan meets the above criteria. Let *Memory*, and *Time*, be, respectively, a model to estimate the memory required by, and the execution time of, an operator or fragment. The algorithm that computes the agenda is shown in Fig. 10 and 11.

The agenda can be conceptualized as a matrix, in which the rows, identified by a relative time point, denote concurrent tasks in the sites which identify the columns. For Fig. 8, the computed agenda is shown in Fig. 9, where $\alpha = 3000\text{ms}$,

```

WHEN-SCHEDULING( $D_Q, R_Q, \alpha, \delta, \text{Memory}, \text{Time}$ )
1  while  $\triangleright$  pre-order traversing  $R_Q$ ,
       $\triangleright$  let  $s$  denote the current site
2      do  $\text{reqMem}_e \leftarrow \text{reqMem}_f \leftarrow 0$ 
3          for each  $f \in s.\text{AssignedFragments}$ 
4              do  $x \leftarrow \text{Memory}(f.\text{EXCHANGE})$ 
5                   $\text{reqMem}_f \leftarrow + \text{Memory}(f) - x$ 
6                   $\text{reqMem}_e \leftarrow + x$ 
7           $\beta^*[s] \leftarrow \lfloor \frac{s.\text{AvailableMemory} - \text{reqMem}_f}{\text{reqMem}_e} \rfloor$ 
8   $\beta \leftarrow \min(\beta^*)$ 
9  while  $\text{agenda}.\text{Makespan} > \min(\alpha * \beta, \delta)$ 
10     do  $\text{agenda} \leftarrow \text{BUILD-AGENDA}(D_Q, R_Q, \alpha, \beta, \text{Time})$ 
11         decr( $\beta$ )
12 return  $\text{agenda}$ 

```

Fig. 10. Computing a SNEEqI Execution Schedule

$\beta = 2$ and $\delta = 5000\text{ms}$. Thus, a non-empty cell (t, s) with value a , denotes that task a starts at time t in site s . In an agenda, there is a column for each site and a row for each time when some task is started. Thus, if cell $(t, s) = a$, then at time t in site s , task a is started. A task is either the evaluation of a fragment (which subsumes sensing), denoted by Fn in Fig. 9, where n is the fragment number, or a communication event, denoted by $tx\ n$ or $rx\ n$, i.e., respectively, tuple transmission to, or tuple reception from, site n . Note that leaf fragments **F2** and **F3** are annotated with a subscript, as they are evaluated β times in each agenda evaluation. Blank cells denote the lack of a task to be performed at that time for the site, in which case, an OS-level power management component is delegated the task of deciding whether to enter a energy-saving state.

In SNEEqI (unlike TinyDB), tuples from more than one evaluation time can be transmitted in a single communication burst, thus enabling the radio to be switched on for less time, and also saving the energy required to power it up and down. This requires tuples between evaluations to be buffered, and results in an increase in the time-to-delivery. Therefore, the buffering factor is constrained by both the available memory and by user expectations as to the delivery time. Note that, query evaluation being continuous, the agenda repeats. The period with which it does so is $p = \alpha\beta$, i.e., $p = 3000 * 2 = 6000$ for the example query. Thus, the acquisition rate α dictates when an ACQUIRE executes; α and the buffering factor β dictate when a DELIVER executes. In this example, note that the agenda makespan is 4344ms . This is calculated by summing the duration of tasks in the agenda (taking into account whether each task has been scheduled sequentially, or concomitantly, in relation to other tasks). Therefore, the delivery time specified in Fig. 1 is met by the example agenda.

Relationship to DQP: The time-sensitive nature of data acquisition in SNs, the delivery time requirements which may be expressed by the user, the need for wireless communications to be co-ordinated and for sensor nodes to duty-cycle, all make the timing of tasks an important concern in the case of SNs. In DQP this is not an issue, as these decisions are delegated to the OS/network layers.

```

BUILD-AGENDA( $D_Q, R_Q, \alpha, \beta, \text{Time}$ )
   $\triangleright$  schedule leaf fragments first
1  for  $i \leftarrow 1$  to  $\beta$ 
2    do for each  $s \in R_Q.\text{Sites}$ 
3      do  $\text{nextSlot}[s] \leftarrow \alpha * (i - 1)$ 
4      while  $\triangleright$  post-order traversing  $D_Q$ 
            $\triangleright$  let  $f$  denote the current fragment
5        do if  $f.\text{isLeaf} = \text{yes}$ 
6          then  $s.f.\text{ActAt} \leftarrow []$ 
7            for each  $s \in f.\text{Sites}$ 
8              do  $s.f.\text{ActAt}.\text{Append } \text{nextSlot}[s]$ 
9                 $\text{nextSlot}[s] \leftarrow + \text{Time}(s.f)$ 
   $\triangleright$  schedule non-leaf fragments next
10 while  $\triangleright$  post-order traversing  $R_Q$ ,
         $\triangleright$  let  $s$  denote the current site
11   do while  $\triangleright$  post-order traversing  $D_Q$ 
            $\triangleright$  let  $f$  denote the current fragment
12     do if  $f \in s.\text{AssignedFragments}$ 
13       then  $f.\text{ActAt} \leftarrow \text{nextSlot}[s]$ 
14            $\text{nextSlot}[s] \leftarrow + \text{Time}(f) * \beta$ 
   $\triangleright$  schedule comms between fragments
15    $s.\text{TX}.\text{ActAt} \leftarrow \max(\text{nextSlot}[s], \text{nextSlot}[s.\text{Parent}])$ 
16    $s.\text{Parent}.\text{RX}(s).\text{ActAt} \leftarrow s.\text{TX}.\text{ActAt}$ 
17    $\text{nextSlot}[s] \leftarrow + \text{Time}(s.\text{TX})$ 
18    $\text{nextSlot}[s.\text{Parent}] \leftarrow + s.\text{Parent}.\text{RX}$ 
19 return agenda

```

Fig. 11. The agenda construction algorithm

Relationship to Related Work: In TinyDB, cost models are used to determine an acquisition rate to meet a user-specified lifetime. The schedule of work for each site is then determined by its level in the routing tree and the acquisition rate, and tuples are transmitted downstream following every acquisition without any buffering. In contrast, our approach allows the optimizer to determine an appropriate level of buffering, given the delivery time constraints specified by the user, which results in significant energy savings as described in Section 4 without having to compromise the acquisition interval. Note that this differs from the orthogonal approach proposed in TiNA [26], which achieves energy savings by not sending a tuple if an attribute is within a given threshold with respect to the previous tuple. It would not be difficult to incorporate such a technique into the SNEEqI optimizer for greater energy savings. Zadorozhny [16] addresses a subset of the when-scheduling problem; an algebraic approach to generating schedules with as many non-interfering, concurrent communications as possible, is proposed. It is functionally similar to the proposed BUILD-AGENDA algorithm, although it only considers the scheduling of communications, and not computations as we do.

3.3 Phase 3: Code Generation

Step 8 in Fig. 3 generates executable code for each site based on the distributed QEP, routing tree and agenda. The current implementation of SNEEqI generates nesC [27] code for execution in TinyOS [28], a component-based, event-driven

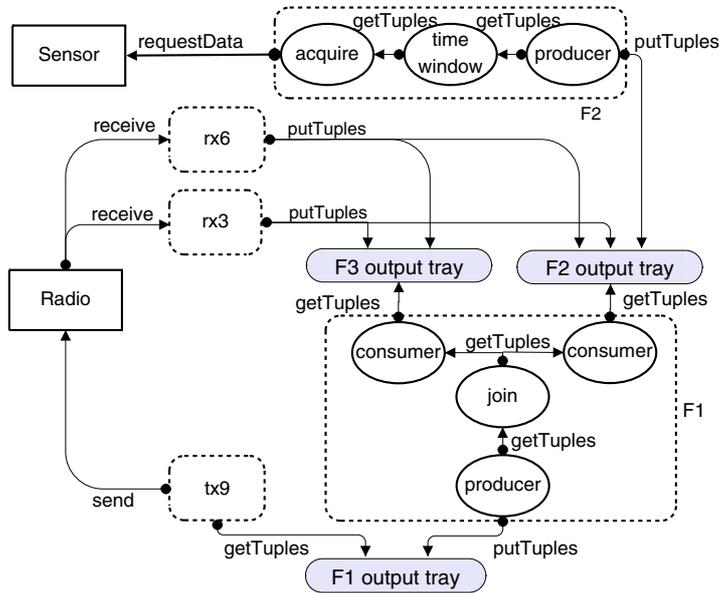


Fig. 12. Generated Code for Site 7 in Fig. 9

runtime environment designed for wireless SNs. nesC is a C-based language for writing programs over a library of TinyOS components (themselves written in nesC). Physical operators, such as those described in this and the previous section, are implemented as nesC template components. The code generator uses these component templates to translate the task-performing obligations in a site into nesC code that embodies the computing and communication activity depicted in abstract form by diagrams like the one in Fig. 12. The figure describes the activity in site 7, where the join (as well as sensing) is performed. In the figure, arrows denote component interaction, the black-circle end denoting the initiator of the interaction. The following kinds of components are represented in the figure: (i) square-cornered boxes denote software abstractions of hardware components, such as the sensor board and the radio; (ii) dashed, round-cornered boxes denote components that carry out agenda tasks in response to a clock event, such as a communication event or the evaluation of a QEP fragment; (iii) ovals denote operators which comprise fragments; note the the correspondence with Fig. 8 (recall that an EXCHANGE operator is typically implemented in two parts, referred to as *producer* and *consumer*, with the former communicating with the upstream fragment, and the latter, the downstream one); and (iv) shaded, round-cornered boxes denote (passive) buffers onto which tuples are written/pushed and from which tuples are read/pulled by other components.

Fig. 12 corresponds to the site 7 column in the agenda in Fig. 9 as follows. Firstly, the acquisitional fragment F2 executes twice and places the sensed tuples in the F2 output tray. Subsequently, tuples are received from sites 6 and 3, and are placed in the F2 output tray and F3 output tray accordingly. Inside fragment

F1, an exchange consumer gets tuples from F2 and another one gets tuples from F3 for the NL_JOIN. The results are fetched by a producer that writes them to the F1 output tray. Finally, tx9 transmits the tuples to site 9.

4 Experimental Evaluation

The goal of this section is to present experimental evidence we have collected in support of our overall research hypothesis, viz., that the extensions (described in Section 3) to DQP techniques that are proven in the case of robust networks lead to effective and efficient DQP over SNs. The experiments are analytical, i.e., aimed at collecting evidence as to the performance of the code produced by the SNEEqI compiler/optimizer.

Experimental Design. Our experimental design centered around the aim of collecting evidence about the performance of our approach for a range of query constructs and across a range of QoS expectations. The QoS expectations we used were acquisition interval and delivery time. The evidence takes the form of measurements of the following performance indicators: network longevity, average delivery time, and total energy consumption. The queries used are generated from those in Fig. 1 (denoted Q1, Q2 and Q3 in the graphs) as follows: (1) in experiments where the acquisition rate is varied, the acquisition rate actually used in the query expression is not, of course, the one in Fig. 1, but should instead be read off the x-axis; and (2) varying the acquisition rate also requires that the `startTime` parameter in the `TIME_WINDOW` over the `inflow` extent of `Q3` is adjusted accordingly at each point in the x-axis, so that the scope and the acquisition rate are consistent with one another at every such point. In the experiments, we assume the selectivity of every predicate to be 1, i.e., every predicated evaluates to true for every tuple. This is so because it is the worst-case scenario in terms of the dependent variables we are measuring, viz., energy consumption and network longevity, as it causes the maximum amount of data to flow through the QEP (and hence through the routing tree formed for it over the underlying network).

Experimental Set-Up. The Experiments were run using Avrora [29], which simulates the behavior of SN programs at the machine code level with cycle-accuracy, and provides energy consumption data for each hardware component. The simulator ran executables compiled from TinyOS 1.1.15. All results are for 600 simulated seconds. A 10-node SN (depicted in Fig. 2 and the schema in Fig. 1) and 30-node SN are simulated in the experiments. The 30-node SN, not shown in this paper due to space limitations, has the same proportion of sources as the 10-node SN. The sensor nodes we have simulated were [Type = Mica2, CPU = 8-bit 7.3728MHz AVR, RAM = 4K, PM = 128K, Radio = CC1000, Energy Stock = 31320 J (2 Lithium AA batteries)].

Experiment 1: *Impact of acquisition interval on total energy consumption.* SNs often monitor remote or inaccessible areas (e.g., [19]), and a significant

part of the total cost of ownership of a SN deployment is influenced by the energy stock required to keep it running. Results are reported in Fig. 13. The following can be observed: (i) As the acquisition rate α is increased, the total energy consumption decreases, as the query plan is acquiring, processing and transmitting less data overall, and sleeping for a longer proportion of the query evaluation process. (ii) A point is reached when increasing the acquisition rate leads to a marginally lower energy saving, due to the constant overhead incurred by having most of the components in a low-power state. The overhead is constant for the default Mica2 sensor board which is simulated, as it does not have a low power state, and is therefore always on. (iii) The radio energy consumption shrinks disproportionately compared to the CPU, because in relative terms, the energy saving when in a low power state is much greater for the radio than it is for the CPU. (iv) *Q2* has the lowest energy consumption because in the aggregation, the tuples from the source sites are reduced to a single tuple; in contrast, *Q3* joins tuples from the `inflow` and `outflow` extents which comprise all tuples in each acquisition, and therefore consumes the most energy.

Experiment 2: *Impact of acquisition interval on network longevity.* The lifetime of a SN deployment is a vital metric as it indicates how often the SN energy stock will need to be replenished. Note that, here, network longevity is assumed to be the time it takes for the first node in the routing tree to fail, so this is, in a sense, a shortest-time-to-failure metric. Fig. 14 reports the results obtained. It can be observed that as acquisition interval increases, energy savings accrue, and hence network longevity increases, albeit with diminished returns, for the same reasons as in *Experiment 1*.

Experiment 3: *Impact of delivery time on network longevity.* For some applications, network longevity is of paramount importance and a delay in receiving the data may be tolerated (e.g., [19]), whereas in other applications it may be more important to receive the data quickly with less regard to preserving energy (e.g., [20]). Results which report the relationship between delivery time and longevity are shown in Fig. 15. It can be observed that: (i) The optimizer reacts to an increase in tolerable delivery time δ by increasing β , which in turn leads to an increase in network longevity – in other words, the optimizer empowers users, by enabling them to trade-off delivery time for greater network longevity; and (ii) Inflection points occur when increasing the buffering factor does not reduce energy consumption. This is because the when-scheduling algorithm assumes that increasing the buffering factor is always beneficial; this is however not always the case. A higher buffering factor may lead to a number of tuples being transmitted at a time that cannot be packed efficiently into a message, leading to padding in the message payload. As an example, consider the case where a single source node is being queried, $\beta = 4$, the number of tuples/message = 3. Packets will be formed of 3 tuples and then 1 tuple, which is inefficient. As a result, more messages need to be sent overall, leading to a higher energy consumption, and hence, a decreased network lifetime. This demonstrates that in order to ascertain an optimal buffering factor, minimizing the padding of messages is also an

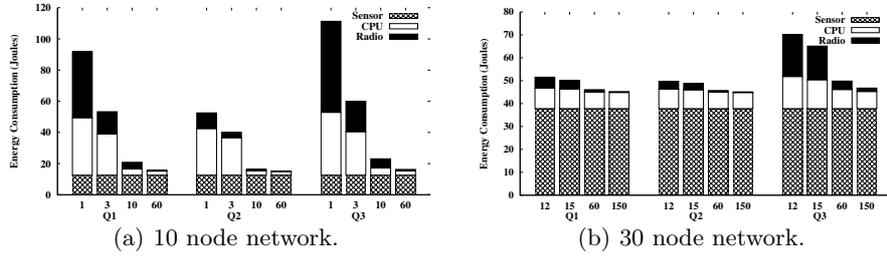


Fig. 13. Energy consumption vs. α (in seconds)

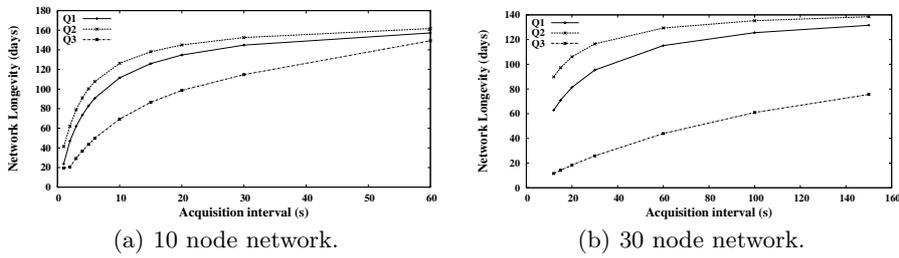


Fig. 14. Network Longevity vs. α

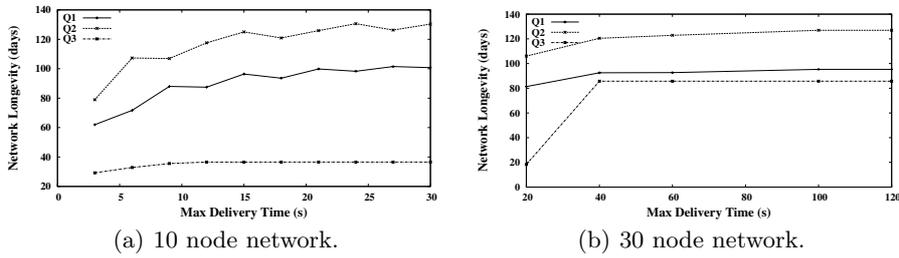


Fig. 15. Network Longevity vs. δ

important consideration. However, we note that maximizing the buffering does lead to an overall improvement in network longevity.

To summarize, we can now ascertain that (i) the SNEEqI optimizer exhibits desirable behaviors for a broad range of different scenarios; (ii) SNEEqI allows different qualities of service (e.g., delivery time and lifetime) to be traded-off. This demonstrates that SNEEqI delivers quantifiable benefits vis-à-vis the seminal contribution in the SN query processing area.

5 Conclusions

In this paper we have described SNEEqI, a SN query optimizer based on the two-phase optimization architecture prevalent in DQP. In light of the differences

between SNs and robust networks, we have highlighted the additional decision-making steps which are required, and the different criteria that need to be applied to inform decisions. We have demonstrated that, unlike TinyDB which performs very limited optimization, the staged decision-making approach in SNEEqI offers benefits, including (1) the ability to schedule different workloads to different sites in the network, potentially enabling more economical use of resources such as memory, and to exploit heterogeneity in the SN, and (2) the ability to empower the user to trade-off conflicting qualities of service such as network longevity and delivery time. The effectiveness of the SNEEqI approach of extending a DQP optimizer has been demonstrated through an empirical evaluation, in which the performance of query execution is observed to be well behaved under a range of circumstances. It can therefore be concluded that much can be learned from DQP optimizer architectures in the design of SN optimizer architectures.

Acknowledgements. This work is part of the SemSorGrid4Env project funded by the European Commission's Seventh Framework Programme, and the DIAS-MC project funded by the UK EPSRC WINES programme under Grant EP/C014774/1. We are grateful for this support and for the insight gained from discussions with our collaborators in the project. C.Y.A. Brenninkmeijer thanks the School of Computer Science, and F. Jabeen, the government of Pakistan, for their support.

References

1. Karl, H., Willig, A.: *Protocols and Architectures for Wireless Sensor Networks*. John Wiley, Chichester (2005)
2. Ganesan, D., Mathur, G., et al.: Rethinking data management for storage-centric sensor networks. In: *CIDR*, pp. 22–31 (2007)
3. Madden, S., Franklin, M.J., et al.: Tinydb: An acquisitional query processing system for sensor networks. *ACM Trans. Database Syst.* 30(1), 122–173 (2005)
4. Yao, Y., Gehrke, J.: Query processing in sensor networks. In: *CIDR*, pp. 233–244 (2003)
5. Estrin, D., Culler, D., et al.: Connecting the physical world with pervasive networks. *IEEE Pervasive Computing*, 59–69 (January-March 2002)
6. Kossmann, D.: The state of the art in distributed query processing. *ACM Comput. Surv.* 32(4), 422–469 (2000)
7. Galpin, I., Brenninkmeijer, C.Y.A., et al.: An architecture for query optimization in sensor networks. In: *ICDE*, pp. 1439–1441 (2008)
8. Gehrke, J., Madden, S.: Query processing in sensor networks. In: *IEEE Pervasive Computing*, vol. 3, pp. 46–55. IEEE Computer Society, Los Alamitos (2004)
9. Brayner, A., Lopes, A., et al.: Toward adaptive query processing in wireless sensor networks. *Signal Processing* 87(12), 2911–2933 (2007)
10. Deshpande, A., Guestrin, C., et al.: Model-based approximate querying in sensor networks. *VLDB J.* 14(4), 417–443 (2005)
11. Trigoni, N., Yao, Y., et al.: Wavescheduling: Energy-efficient data dissemination for sensor networks. In: *DMSN*, pp. 48–57. ACM Press, New York (2004)
12. Madden, S., Franklin, M.J.: et al.: Tag: A tiny aggregation service for ad-hoc sensor networks. In: *OSDI*, pp. 131–146 (2002)

13. Manjhi, A., Nath, S.: et al.: Tributaries and deltas: Efficient and robust aggregation in sensor network streams. In: SIGMOD, pp. 287–298 (2005)
14. Trigoni, N., Yao, Y., et al.: Multi-query optimization for sensor networks. In: DCOSS, pp. 307–321 (2005)
15. Bonfils, B.J., Bonnet, P.: Adaptive and decentralized operator placement for in-network query processing. In: Zhao, F., Guibas, L.J. (eds.) IPSN 2003. LNCS, vol. 2634, pp. 47–62. Springer, Heidelberg (2003)
16. Zadorozhny, V.I., Chrysanthis, P.K., et al.: A framework for extending the synergy between query optimization and mac layer in sensor networks. In: DMSN, pp. 68–77. ACM Press, New York (2004)
17. Arasu, A., Babcock, B., et al.: Stream: The stanford stream data manager. IEEE Data Eng. Bull. 26(1), 19–26 (2003)
18. Brenninkmeijer, C.Y.A., Galpin, I., et al.: A semantics for a query language over sensors, streams and relations. In: Gray, A., Jeffery, K., Shao, J. (eds.) BNCOD 2008. LNCS, vol. 5071, pp. 87–99. Springer, Heidelberg (2008)
19. Marshall, I.W., Price, M.C., et al.: Multi-sensor cross correlation for alarm generation in a deployed sensor network. In: Kortuem, G., Finney, J., Lea, R., Sundramoorthy, V. (eds.) EuroSSC 2007. LNCS, vol. 4793, pp. 286–299. Springer, Heidelberg (2007)
20. Stoianov, I., Nachman, L., et al.: Pipenet: a wireless sensor network for pipeline monitoring. In: IPSN, pp. 264–273 (2007)
21. Sampaio, S.F.M., Paton, N.W., et al.: Measuring and modelling the performance of a parallel odmg compliant object database server. CCPE 18(1), 63–109 (2006)
22. Stonebraker, M., Aoki, P.M., et al.: Mariposa: A wide-area distributed database system. VLDB J. 5(1), 48–63 (1996)
23. Graefe, G.: Encapsulation of parallelism in the volcano query processing system. In: SIGMOD, pp. 102–111 (1990)
24. Gounaris, A., Sakellariou, R., et al.: A novel approach to resource scheduling for parallel query processing on computational grids. DPD 19(2-3), 87–106 (2006)
25. Smith, J., Gounaris, A., Watson, P., et al.: Distributed query processing on the grid. In: GRID, pp. 279–290 (2002)
26. Sharaf, M.A., Beaver, J., et al.: Tina: A scheme for temporal coherency-aware in-network aggregation. In: MobiDE, pp. 69–76 (2003)
27. Gay, D., Levis, P., et al.: The nesc language: A holistic approach to networked embedded systems. In: PLDI, pp. 1–11 (2003)
28. Hill, J., Szewczyk, R., et al.: System architecture directions for networked sensors. In: ASPLOS, pp. 93–104 (2000)
29. Titzer, B., Lee, D.K., et al.: Avrora: scalable sensor network simulation with precise timing. In: IPSN, pp. 477–482 (2005)