

# Spatio-Temporal Sensor Graphs (STSG): A Data Model for the Discovery of Spatio-Temporal Patterns\*

Betsy George<sup>†</sup>  
Dept. of Computer Science  
University of Minnesota  
Minneapolis, MN  
bgeorge@cs.umn.edu

James M. Kang  
Dept. of Computer Science  
University of Minnesota  
Minneapolis, MN  
jkang@cs.umn.edu

Shashi Shekhar  
Dept. of Computer Science  
University of Minnesota  
Minneapolis, MN  
shekhar@cs.umn.edu

## ABSTRACT

Developing a model that facilitates the representation and knowledge discovery on sensor data presents many challenges. With sensors reporting data at a very high frequency, resulting in large volumes of data, there is a need for a model that is memory efficient. Since sensor data is spatio-temporal in nature, the model must also support the time dependence of the data. Balancing the conflicting requirements of simplicity, expressiveness and storage efficiency is challenging. The model should also provide adequate support for the formulation of efficient algorithms for knowledge discovery. Though spatio-temporal data can be modeled using time expanded graphs, this model replicates the entire graph across time instants, resulting in high storage overhead and computationally expensive algorithms. In this paper, we propose Spatio-Temporal Sensor Graphs (STSG) to model sensor data, which allow the properties of edges and nodes to be modeled as a time series of measurement data. Data at each instant would consist of the measured value and the expected error. Also, we present several case studies illustrating how the proposed STSG model facilitates methods to find interesting patterns (e.g., growing hotspots) in sensor data.

## 1. INTRODUCTION

Finding novel and interesting spatio-temporal patterns in the ever increasing collection of sensor data is an important problem in several scientific domains. Many of these scientific domains collect sensor data in outdoor environments with underlying physical interactions. For example, in environmental science, a timely response to anticipated watershed/in-plant events (e.g., chemical spill, terrorism, etc.) to maintain water quality is required. Such a case occurred in Milwaukee, WI, in 1993 where a harmful pathogen

(called *Cryptosporidium parvum*) outbreak occurred in the river streams that infected more than 400,000 people with deaths exceeding 100. The estimated total cost for the treatment of outbreak-related illness was \$96.2 million [6]. As was the case in Milwaukee, such failures typically are detected long after the exposure by observed spikes in doctor/hospital visits or sales of certain medicines. In addition to unplanned “natural” events like the *Cryptosporidium* episode, another concern regarding water supplies is an act of terrorism. Clearly, when public health is at stake, waiting for the illnesses and fatalities to arise is much too late and identifying and modeling these spatio-temporal patterns such as hotspots and growing hotspots from sensor graphs is important [1]. Other applications that generate similar sensor data may be traffic road systems where measurements of traffic flow and congestion are important, especially in emergency operations such as evacuations.

A collection of sensors may be represented as a sensor graph where the nodes represent the sensors and the edges represent selected relationships. For example, sensors upstream and downstream in a river may have physical interactions via water flow and related phenomenon such as plume propagation. Relationships can also be geographical in nature, such as proximity between the sensor units. As an example, Figure 1(a) shows a layout of traffic sensors in the Twin Cities, MN. The graph representation of a part of this layout (given in Figure 1(b)) is shown in Figure 1(c). The nodes of the graph represent the sensors and the edges represent the physical relationships between the various sensors. In this example, the edges are based on the proximity between the sensors.


Formulation of a model to represent a sensor graph that supports mining useful information from data poses some significant challenges. Since the volume of data is large, the model used to represent the sensor graph must be storage efficient. It should also provide sufficient support for the design of correct and efficient algorithms for data analysis. Second, the sensor graph characteristics modeled as pairs,  $\langle \text{measured value, error} \rangle$ , can be time-dependent (e.g., the flow rate in a river stream). The model used to represent a time-dependent graph should be able to represent the time-variance, simultaneously maintaining the storage efficiency.

A sensor graph is spatio-temporal in nature since the relative locations of the sensor nodes and the time-dependence of

---

\*This work was supported by the NSF-SEI grant, NSF-IGERT grant, Oak Ridge National Laboratory grant and US Army Corps of Engineers (Topographic Engineering Center) grant. The content does not necessarily reflect the position or policy of the government and no official endorsement should be inferred.

<sup>†</sup>Corresponding Author



**Figure 1: Sensor networks periodically report time-variant traffic volumes on Twin Cities, MN highways (Best viewed in color, Source: Mn/DOT)**

their characteristics are significant. Spatio-Temporal graphs can be modeled as time expanded graphs, where the entire network is replicated for every time instant [16]. The changes in the graph can be very frequent and for modeling such frequent changes, the time expanded networks would require a large number of copies of the original network, thus leading to network sizes that are too memory expensive. Moreover, while modeling sensor graphs that involve no physical flow, a direct application of this model might not be possible.

A Spatio-temporal Sensor Graph (STSG) models the changes in a spatio-temporal graph by collecting the node and edge attributes into a set of time series. The model can also account for the changes in the topology of the network. The edges and nodes can disappear from the network during certain instants of time and new nodes and edges can be added. A Spatio-temporal sensor graph keeps track of these changes through a time series attached to each node and edge that indicates their presence at various instants of time. The stochastic nature of the physical relationships between the sensors (e.g., the flow rate of the river stream that connects the sensors) is accounted for by expressing each element in the attribute time series as a pair of values (i.e.,  $\langle \text{measured value, error} \rangle$ ). Several case studies are also provided to validate the model in the context of discovery of spatio-temporal patterns from sensor data. Analysis shows that this model is less memory expensive and leads to algorithms that are computationally more efficient than those for the time expanded networks.

## 1.1 Application Domain

Modeling spatio-temporal graphs has significant applications in a number of scientific application domains. Discovering knowledge from the large amount of data collected from sensors can be used in predicting trends in domains such as environmental science, thus emphasizing the need for a model. Transportation network flow patterns are being increasingly monitored by sensors. The data can be used to find routes that are frequently congested and can be used in network planning. Since varying levels of congestion can lead to time dependent travel times on road segments, the road network represented by a spatio-temporal graph might give more accurate results for routing queries such as short-

est paths. Accounting for time dependence in transportation networks would make evacuation planning algorithms in emergency planning generate results that are more accurate.

The role of spatio-temporal data mining in the environmental sciences is shown in Figure 2. The figure gives an example of the water flow starting from a water treatment plant to a sensor network collecting data throughout the watershed. The data collected from the sensor network is handled in two parts: (a) the collected data are analyzed for any interesting patterns (e.g., anomalies) and if any are found, a diagnosis and prognosis is made, and (b) the collected data are used for learning of any new and interesting spatio-temporal patterns (e.g., growing hotspots) resulting in the refinement of the diagnosis and prognosis rule base. Based on the prognosis, the treatment plant can readjust any necessary parameters for the plant control system and plant models. Such a spatio-temporal data mining system could be used to issue warnings to recreational users of water resources (e.g., beach closings) and to optimize water treatment conditions to protect the public and sensitive water treatment infrastructure. For example, similar to weather forecasts, a water quality prediction for beaches could be provided (e.g., 30% likelihood that coliform levels will be exceeded).

Similarly, if a spike in pathogen or hazardous chemical concentration is predicted, water intake could be suspended temporarily, processes could be adjusted in real time, or an additional treatment process could be brought online. The discovery of interesting patterns within sensor data for outdoor application domains is often arduous and complex. Many challenges [5, 7, 17, 20] and hurdles need to be overcome. One challenge is to support remote monitoring of sensor networks distributed over an area, to check the overall functioning of the system as well as to detect interesting events related to measurements. Examples include the tasks of identifying malfunctioning sensors or interesting events.

In general, there are two types of outdoor sensor networks (e.g., [2, 5, 7, 10, 12, 17, 20]). The first type is a *wired* sensor network e.g., traffic management center at the Minnesota Department of Transportation [25, 24] that consists of the following: (a) sensors that are wired within the Twin Cities,

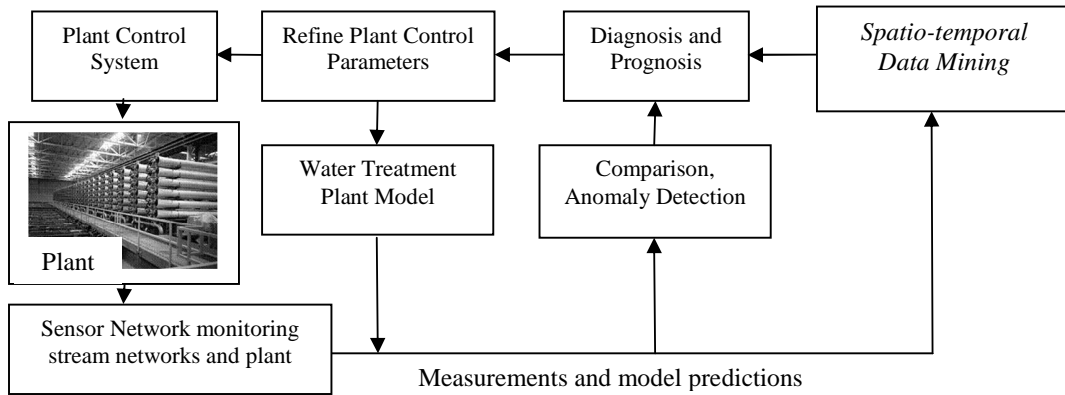


Figure 2: Spatio-temporal Data Mining in Transformative Water Quality

MN highway system and (b) a very large network containing over 4000 sensors across a 20-mile radius in the Twin Cities and sampled every five minutes, and (2) a real live application in a metropolitan area. The second type is a *wireless* sensor network, for example, the one used by the Water Resources community that has been recently deployed at the Minnehaha Creek in Minnesota that consists of (a) a set of sensors placed in the environment communicating wirelessly among each other, (b) an initial deployment of a handful of sensors and a future goal of increasing the number of sensors to monitor the Mississippi River, and (c) a live application in the natural environment [1].

## 1.2 Related Work

Graphs have been used to represent collection of sensors and to formulate algorithms for applications such as routing and location tracking [18, 11]. Recent research in spatial anomaly detection proposed in [25, 24] used sensor datasets structured as graphs. An accurate representation of sensors should include the spatial attributes and time dependent parameters of the graph. Most graph representations ignore the time-dependence of the attributes. Some knowledge discovery algorithms are limited due to the ignoring of spatial relationships [3, 14]. Some work has focused on managing the datasets produced by *wired* and *wireless* sensor networks using ‘spatial time series’ [26].

Traditionally, spatio-temporal networks such as transportation networks have been modeled using time expanded graphs [16, 15, 19]. This method duplicates the original network for each discrete time unit  $t = [0, 1, \dots, T]$  where  $T$  represents the extent of the time horizon. The expanded network has edges connecting a node and its copy at the next instant in addition to the edges in the original network, replicated for every time instant. This approach significantly increases the network size and is very expensive with respect to memory. Because of the increased problem size due to replication of the network, the computations become expensive. In addition, time expanded graphs cannot model sensor networks in cases where there is no flow parameter involved. In such cases, the cross edges that represent a flow from one node to another lose significance.

**What is the novelty in the proposed model?**

We proposed a model called Time Aggregated Graph (TAG) to represent spatio-temporal networks [9, 8]. This model ag-

gregates time dependent parameters on edges and nodes to time-series attributes rather than replicate the entire graph for each time instant. TAG can also model topological changes that could occur in the graph (e.g., disappearance of an edge during a time interval) over time.

**What is the improvement over existing approaches?**

Analysis shows that the model is more memory efficient than time expanded graphs. According to the analysis in [22], the memory requirement for a time expanded network is  $O(nT) + O(n + m)T$ , where  $n$  is the number of nodes and  $m$  is the number of edges in the original graph. The memory requirement for the time-aggregated graphs would be  $O(m + n)T$ . Since the physical relationships among sensors are stochastic in nature, there is a need to model the probabilistic characteristics of edges and nodes, which is not modeled in TAG. Spatio-Temporal Sensor Graph (STSG) represents stochastic parameters by specifying the measurement and the expected error. Each attribute value would be a pair,  $\langle \text{measured value, error} \rangle$ .

**Is the approach general?**

The STSG model has been used in this paper as the basis for the proposed algorithms for hotspot discovery and growing hotspot detection. The basic model (TAG) has been used in formulating routing algorithms transportation networks with time-dependent travel times such as shortest path computation and best start time computation [8].

## 1.3 Contribution

This paper makes the following contributions:

- We propose Spatio-Temporal Sensor Graphs (STSG) to model sensor datasets. Time aggregated graphs are generalized to include probability parameters to incorporate the stochastic nature of sensor graphs.
- We present a case study on anomaly detection and propose a method to generate the STSG model.
- We give a case study on basic hotspots and propose a detection method utilizing the STSG model.
- We illustrate a case study on growing hotspots and propose a detection method using STSG model.

## 1.4 Scope and Organization

The paper presents a model called Spatio-Temporal Sensor Graph to represent sensor data. It presents a case study of

the model using algorithms to find hot spots and growing hot spots using sensor data. Though the model incorporates the stochastic nature of various attributes, the proposed algorithms currently do not account for the probabilistic values of various parameters.

The rest of the paper is organized as follows. Section 2 presents basic concepts to provide a formal model of Spatio-Temporal Sensor Graphs. Section 3 presents several case studies and proposed approaches utilizing the STSG model. Finally, Section 4 concludes this paper with a conclusion and future work.

## 2. BASIC CONCEPTS

Traditionally graphs have been extensively used to model spatial networks [23]; weights assigned to nodes and edges are used to encode additional information. For example, the spatial location of a sensor can be represented using the attribute assigned to the node that represents the sensor and the flow rate of a river stream between two sensors can be represented by an attribute of the edge connecting the nodes. In a real world scenario, it is not uncommon for these parameters to be time-dependent. This section discusses a graph based model that can capture the time-dependence of network parameters. In addition, the model captures the possibility of edges and nodes being absent during certain instants of time.

### 2.1 Spatio-temporal Sensor Graph (STSG)

A graph  $G = (N, E)$  consists of a finite set of nodes  $N$  and edges  $E$  between the nodes in  $N$ . If the pair of nodes that determine the edge is ordered, the graph is directed; if it is not, the graph is undirected. In most cases, additional information is attached to the nodes and the edges. In this section, we discuss how the time dependence of these edge/node parameters are handled in the proposed model, the Spatio-Temporal Sensor Graph.

We define the Spatio-Temporal Sensor Graph as follows.

$$STSG = (N, E, TF, \\ f_1 \dots f_k, g_1 \dots g_l, \\ (nw_1, ne_1) \dots (nw_k, ne_p), \\ (ew_1, ee_1) \dots (ew_p, ee_p), \\ f_i : N \rightarrow \mathbb{R}^{TF}, g_i : E \rightarrow \mathbb{R}^{TF}; \\ nw_i : N \rightarrow \mathbb{R}^{TF}, ne_i : N \rightarrow \mathbb{PD}, \\ ew_i : E \rightarrow \mathbb{R}^{TF}, ee_i : E \rightarrow \mathbb{PD})$$

where  $N$  is the set of nodes,  $E$  is the set of edges,  $TF$  is the length of the entire time interval,  $f_1 \dots f_k$  are the mappings from nodes to the time-series associated with the nodes (for example, the time instants at which the node is present),  $g_1 \dots g_l$  are mappings from edges to the time series associated with the edges and  $(ew_1, ee_1) \dots (ew_p, ee_p)$  indicate the time dependent attribute on the edges.  $\mathbb{PD}$  indicates a probabilistic error. These attributes are the quantitative descriptors of the physical relationship between the nodes. To represent the stochastic nature of the measured values of physical phenomena, each attribute is a pair that represent the measured value and the associated error.

**Example:** A graph representation of a network at three instants of time is shown in Figure 3(a), including temporal changes in connectivity and edge properties (e.g., flow rate). Each edge attribute is a pair  $\langle$ measured value, error $\rangle$ . The

first parameter in the pair is the measured value and the second is the expected error. For example, the edge from node N3 to node N4 disappears at the time instant  $t=2$  and reappears at  $t=3$ , and the attribute on the edge N3-N4 changes from  $(1,0.5)$  at  $t=1$ , to  $(4,0.5)$  at  $t=3$ . Figure 3(b) shows the spatio-temporal sensor graph. This is encoded in the time-aggregated graph using the edge attribute series of N3-N4 which is  $[(1,0.5), \infty, (4,0.5)]$ ; the entry ' $\infty$ ' indicates that the edge is absent at the time instant  $t = 2$ .

A time expanded graph would need copies of the entire graph. The number of nodes is larger by a factor of  $T$ , where  $T$  is the number of time instants and the number of edges is also larger in number compared to the STSG. Typically the value of  $T$  is very large in a spatial network such as a sensor graph since the changes in the network are quite frequent. This would result in time expanded networks that are enormously large and would make the computations slow. Though flow based physical phenomena can be modeled using the cross edges between the copies in a time expanded graph, it might not be possible to represent every type of physical relationship between the sensors.


## 3. CASE STUDIES

This section presents three case studies using the spatio-temporal sensor graph (STSG) to discover interesting patterns. First, anomaly detection is presented that generates a set of time series at each sensor node where each time interval is identified as an anomaly. The second case study is the discovery of basic hotspots using the STSG model and the time series that was discovered in the unanticipated anomaly case study. Finally, a case study on the discovery of growing hotspots using the STSG model is presented. The case studies addressing unanticipated anomaly detection, basic hotspot, and growing hotspots are presented in Sections 3.1, 3.2, and 3.3 respectively. The algorithms presented do not consider the probabilistic parameters in discovering and predicting hotspots. Prediction is done assuming that the physical model of the edge parameter is accurate. However, the algorithms can be extended to incorporate the stochastic nature of the edge parameters.

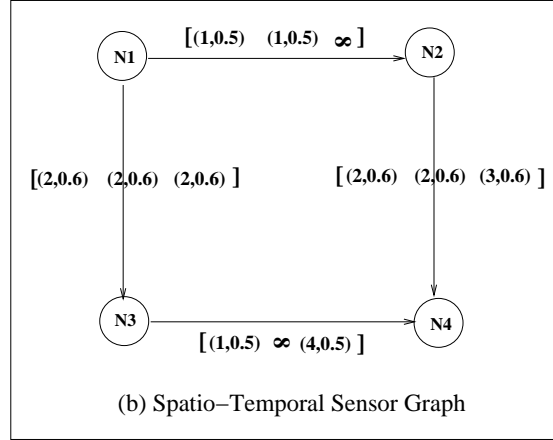
### 3.1 Anomaly Detection

**Definition.** The problem of anomaly detection can be generally described as identifying a set of observations that are inconsistent or generated from a different mechanism than the rest of the dataset [23]. In general, anomaly detection can be categorized in five different groups. (1) global anomalies where a data point may be different from the entire dataset and is often used in several traditional aggregate approaches, (2) spatial anomalies where a data point may be different when comparing against its physical neighbors, (3) temporal anomalies which are similar to spatial anomalies except that the neighborhood region is based on its temporal neighbors, (4), spatio-temporal anomalies where a data point is different based on the neighborhood of both space and time, and (5) network anomalies where a data point is different based on the neighboring connected nodes.

**Application.** Several interesting applications can utilize anomaly detection such as the monitoring of watersheds and in-plant systems using a sensor network. However, challenges arise in such situations. For example, in a watershed,



(a) Snapshots of a Network at time instants 1, 2, and 3



(b) Spatio-Temporal Sensor Graph

Figure 3: Spatio-temporal Sensor Graph and Snapshots at various instants

a level of normality needs to be defined to determine the types of anomalies. Normality may be based on domain models (e.g., physics based differential equation governing fluid flow, see [4, 21] for more information) using some form of an aggregate on historical data gathered by the sensor network. The historical dataset may be broken into several slices (or by month) based on the seasonality. Another challenge includes change detection between sensors in a watershed. For example, suppose we have sensors in place within the length of a river with a train track running along the side of it. Suppose a train carrying harmful contaminants causes a spill into the river stream. The sensors within this part of the river will start observing measurements that are significantly different from domain models as well as in the upstream section of the river. Thus, an anomaly can be detected to warn the water treatment plant.

**Method.** Algorithm 1 presents the pseudocode to detect spatio-temporal anomalies from a sensor network. There are three types of input to this approach. The first is a set of nodes within the sensor network where each node contains information about the measurements gathered at a time interval. The second is a model defined by the domain scientists to determine the predicted output under normal conditions. Third is the maximum number of time intervals recorded within the sensor network. The output of this approach consists of a set of time series for each node where an anomaly was detected. These time series will be needed

---

**Algorithm 1** Pseudo code for Anomaly Detection

---

```

1: Function ANOMALY(set NodeIds, model m, int
   maxTime)
   {Phase I: Domain Science Model}
2: for each node n ∈ NodeIds do
3:   for time t = 1 to maxTime do
4:     Calculate predicted output o using m for
       n.measurements at t
5:     successorNode.predicted.t = o
6:   end for
7: end for
   {Phase II: Identify Anomalies}
8: for node n = 1 to NodeIds.size() do
9:   for time t = 1 to maxTime do
10:    if n.measurements ≠ n.predicted at t then
11:      n.anomaly = n.anomaly ∪ t
12:    end if
13:   end for
14: end for

```

---

as inputs for the case studies discussed in Section 3.2 and Section 3.3.

Algorithm 1 consists of two phases to detect anomalies within a sensor network. In Phase I, a domain science model is used to calculate the predicted output under normal conditions for each node (Lines 2-7 of Algorithm 1). The input is based on the current node and the predicted output is for the successor nodes, based on a domain model (e.g., [4, 21]). The

predicted output is stored at the successor node (Line 5 of Algorithm 1). This process is computed for all time intervals for each node in the sensor graph.

Phase II of Algorithm 1 identifies the anomalies based on the domain science model and the actual measurements from the sensor graph (Lines 8-14 of Algorithm 1). Every node is investigated for anomalies except the first node because the predicted output generated by the domain science model is calculated based on its predecessor node (Lines 8 of Algorithm 1). An anomaly is identified when the actual measurements at a node and the predicted output for a time interval is different or the difference is greater than some threshold (Line 10 of Algorithm 1). If an anomaly is found, then the time interval  $t$  is assigned to a set of time series for the node (Lines 11 of Algorithm 1). This information is intended for the Spatio-Temporal Sensor Graph (STSG) model which will be used in the Sections( 3.2, 3.3).

**Execution Trace.** Table 1 presents an example dataset containing a set of measurements found in a sensor network. The values depicted in this example uses a simple prediction model, for further information on differential models, see [4, 21]. In this example, we have five nodes running along a river where the node id increases based on the direction of the water flow. In each node, there are three time intervals and a corresponding measurement at that sensor (e.g., concentration of a chemical in the water).

In Phase I of Algorithm 1, the predicted outputs are calculated for each node having a predecessor node. An example of these predicted outputs are shown in Table 1. Notice that there are no predicted values for node 1 because a predecessor node (e.g., node 0) does not exist in the dataset.

In Phase II of Algorithm 1, the anomalies are identified by having a set of time series when the predicted values and the actual sensor measurements differ. Table 1 gives an example where nodes 2 and 3 has anomalies occurred at  $\{t_2, t_3\}$  and nodes 4 and 5 has anomalies at  $\{t_3\}$ . Figure 4(b) gives an illustration of the STSG model with the anomaly time series found in this execution trace.

**Computational Complexity.** Algorithm 1 has two phases: (Phase 1) the predicted output for each node and is calculated for all time intervals having the complexity of  $O(nT)$  where  $n$  is the number of nodes and  $T$  is the total number of time intervals; and (Phase 2) the anomalies are detected by examining  $(n - 1)$  nodes (the boundary nodes are not checked) for all time intervals and having the complexity of  $O((n - 1)T)$ . Thus, the total computational complexity of both phases 1 and 2 is  $O(nT + (n - 1)T) = O(nT)$ .

## 3.2 Basic Hotspot Detection

**Definition.** The problem of hot spot detection is to discover the sensor nodes that display significant differences between observed values and expected ‘standard’ values.

**Application.** In application domains such as river systems where chemical levels are constantly monitored, sensors are deployed to detect the changes. In this context, a hotspot is indicated by a sensor reporting an anomaly (as discussed in Section 3.1). We discuss a method to discover hotspots using

a model called the Spatio-Temporal Sensor Graph (STSG). The nodes in the STSG represent the sensors. An edge is added between the nodes if and only if there is a physical relationship between the nodes. The presence of a hotspot at a node at various time instants is indicated by a node time series. In addition, the time dependence of the physical relationships modeled by the edges can be represented by edge time series attributes. Figure 4 illustrates the graph model for the sensor graph. For the sake of simplicity, edge attributes are not shown in Figure 4. Figure 4(a) shows an example network. The nodes that are active at time instants  $t = 2$  and  $t = 3$  are shown in Figure 4(b) and (c). The Spatio-Temporal Sensor Graph representation is shown in Figure 4(c). The time series attributes on the nodes indicate the hotspots at various time instants. For example, the time series 2,3 on the node N2 indicates that the node is a hot spot at  $t = 2, 3$ .

**Method.** Given a sensor graph called the source node, the hot spot at any time instant is the set of nodes where an anomaly has been detected the given time instant. We use a modified breadth first strategy to find the nodes that indicate the hot spots at any time instant. The pseudo-code is provided in Algorithm 2.

The algorithm finds the hotspots using the Spatio-Temporal Sensor Graph model. Each node has a time series attribute that encodes the information about the time instants at which the node has an anomaly. For example, the time series [2,3] at node N2 in Figure 4(d) indicates that the node is a hotspot at  $t = 2, 3$ . The algorithm searches the graph starting at (any) given node for each value of time  $t$  and finds the hotspots. The search uses a breadth-first strategy, modified to incorporate the fact that each node has a time series that needs to be checked. When each node is visited, the algorithm checks to see whether it is a hot spot by checking the node time series. The node time series is assumed to be sorted. The output of the algorithm is the set of hotspots at every time instant.

---

### Algorithm 2 Hotspot Algorithm

---

```

1: Function BASICHOTSPOTS(Graph  $G(N, E)$ , set  $N$ , set
    $E$ , node  $source$ )
2: for  $t = 1, T$  do
3:   mark  $source$  as visited;
4:   enqueue(Q, source);
5:   if  $t$  in node_time_series of source then
6:     hotspots[t] = source;
7:   end if
8:   while Q not empty do
9:      $u = Dequeue()$ ;
10:    For every node  $v$  such that  $uv \in E$  and if  $t$  in
       node_time_series
11:     if  $v$  is not marked then
12:       mark  $v$  as visited.
13:       enqueue(Q,  $v$ );
14:       hotspots[t] = hotspots[t]  $\cup$   $v$ ;
15:     end if
16:   end while
17: end for

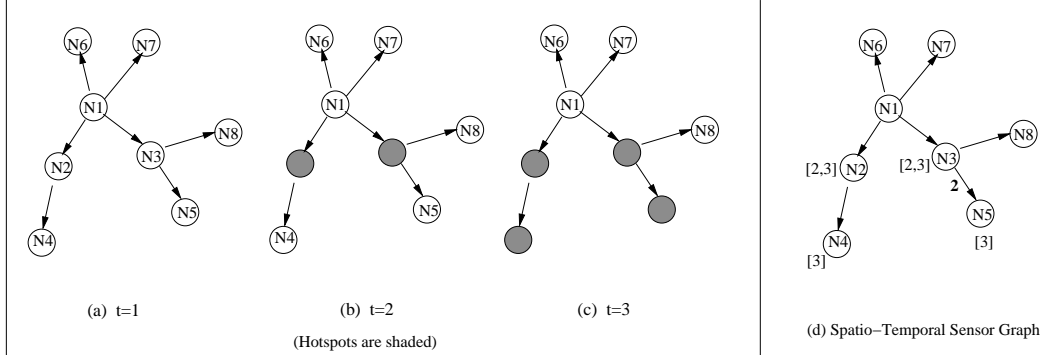
```

---

**Execution Trace.** Table 2 shows the trace of the algorithm for the STSG shown in Figure 4(d). The search starts at node N1 at  $t = 1$  and detects no hotspots. At  $t = 2$ , the

**Table 1: Execution Trace of Anomaly Detection Algorithm**

Nodes	1			2			3			4			5		
Time Slot	1	2	3	1	2	3	1	2	3	1	2	3	1	2	3
Measure	10	20	30	20	50	50	40	50	50	80	100	50	160	200	50
Predicted				20	40	60	40	100	100	80	100	100	160	200	100
Anomaly				$\{t_2, t_3\}$			$\{t_2, t_3\}$			$\{t_3\}$			$\{t_3\}$		


**Figure 4: Spatio-Temporal Sensor Graph Model to Detect Hotspots**

search finds that the nodes N2 and N3 are hotspots based on the entry ‘2’ (indicating the presence of a hotspot at  $t = 2$ ) in their node time series [2,3]. The algorithm performs another iteration for  $t = 3$  and finds the hotspots at N2,N3,N4, and N5. The execution trace is summarized in Table 2.

**Table 2: Execution Trace of the Hotspot Algorithm**

time	t1	t2	t3
Hotspot Nodes	$\emptyset$	$\{N2, N3\}$	$\{N2, N3, N4, N5\}$

### Computational Complexity

The algorithm visits every node of the Spatio-Temporal Sensor Graph at every time instant  $t$  and searches the node time series to detect the hotspots. The algorithm performs  $O(n)$  steps to visit all nodes. At each step, the presence of the node at the given time instant is checked. If the time series is sorted, this look-up is  $O(\log T)$  since the length of the time series is at most  $T$  where  $T$  is the length of the time period. At each node this operation has a complexity of  $degree(node) \cdot \log T$ . The cost over all the nodes is hence,  $O(m \log T)$  where  $m$  is the number of edges in the graph. Since the search is performed at every instant, the computational complexity is  $O((n + m \log T)T)$ , where  $n$  is the number of nodes,  $m$  is the number of edges, and  $T$  is the length of the time period.

### 3.3 Growing Hotspot Detection

**Definition.** The problem of growing hot spot detection is to predict expanding patterns that may lead to significant differences between some observations and the rest of the data set. This problem is different from anomaly detection, where a growing pattern may not be identified as an outlier due to the lack of severity at the early stages of growth, e.g., slow leakage from buried chemical drums to the ground water supply. Domain knowledge (e.g., flow rate, plume model)

is incorporated in spatio-temporal sensor graph representation to predict the growing hot spots.

**Application.** In application domains such as river systems where chemical levels are constantly monitored, sensors may be deployed to detect the changes. In this context, a growing hotspot is indicated by the increase in the number of sensors reporting an anomaly, over time. We discuss a method to discover growing hotspots using a model called the Spatio-Temporal Sensor Graph (STSG), which predicts the spread of hotspots. The nodes in the spatio-temporal sensor graph represent the sensors. The edges in the STSG are quantified with the descriptors that represent the propagation of the cause of anomaly. For example, parameters that model the fluid flow between two sensors can be modeled as an edge attribute. Each attribute is a pair that consists of the measured value and the expected error. If the attribute is time-dependent, STSG will represent the attribute as a time series.

Figure 5 shows a spatio-temporal sensor graph that represents a collection of sensors. Nodes represent the sensors and the edges indicate that they are connected (for example, by fluid flow). The edge parameters represent the characteristics of the underlying connection. In this example, for the sake of simplicity, the edge parameters are shown to be constants, though spatio-temporal sensor graphs can represent time-varying edge attributes. Figure 5(a) shows the state of the sensors at  $t = 1$ , where the node N1 is active. Figure 5(b) and 5(c) show the predicted active sensor nodes at  $t = 2, 3$ . For example, at  $t = 2$ , node N2 is predicted to be a hot spot and at  $t = 3$ , N11 is expected to be active. The predicted hot spots are shown as shaded circles in the figure.

**Method.** Given a set of sensor nodes (called the source nodes), the hot spots at any future time instant can be predicted from the spatio-temporal sensor graph using the edge




Figure 5: Detecting Growing Hotspots

attributes that describe the propagation between the sensor nodes. The algorithm finds the nodes that are reachable from the source nodes within the interval between two time instants. These nodes are discovered based on the edge attribute values that quantify the physical relationship between the nodes. Pseudocode of the algorithm is provided in Algorithm 3. Though the model represents each attribute as a pair of values that indicate the measured value and the associated error, this algorithm does not consider the error parameter in its computations. The error values are omitted from the Figure 5, for the sake of simplicity.

---

**Algorithm 3** Growing Hotspot Algorithm

---

```

1: Function SPREADINGHOTSPOTS(Graph  $G(N, E)$ , set  $N$ ,
   set  $E$ , set  $source\_nodes$  )
2: Enqueue( $Q, S$ );
3: for  $t=1, T$  do
4:   time_elapsed[ $u$ ] =  $\Delta T$  for all nodes;
5:   while Queue  $Q$  not Empty do
6:      $u = Dequeue()$ ;
7:     For every node  $v$  such that  $uv \in E$ 
8:       if  $v$  is not marked &&  $v$  is reachable based on  $d(u, v)$ 
9:         then
10:          mark  $v$  as visited.
11:          enqueue( $Q, v$ );
12:          decrement  $time\_elapsed[v]$ ;
13:          Add  $v$  to hotspots[ $t$ ];
14:          return hotspots;
15:       end if
16:   end while
17: end for

```

---

**Execution Trace.** Table 3 gives the trace of the algorithm for a stream network shown in Figure 5. For the sake of simplicity, it is assumed that every segment (represented by an edge) has the same length. The edge attributes are the flow rates. The algorithm lists the nodes reachable from every node for each time instant. At  $t = 1$ , node N1 is a hot spot. Since the flow rate on edge N1-N2 is 1 unit, flow is predicted to reach N2 at  $t = 2$  and N2 is added to the list of hot spots at  $t = 2$ . Similarly, flow is predicted to reach N3 at  $t = 2$ . Since the rate is faster, the flow moves through N3 and reaches N5 at  $t = 2$ , adding both N3 and N5 to the list at  $t = 2$ .

Table 3: Trace of the Growing Hot Spot Algorithm

time	t1	t2	t3
Predicted	N1	N1,N2,N3,N4	N1,N2,N3,N4,N5
Hot Spots		N7,N9	N7,N9,N10,N11

**Computational Complexity**

This algorithm finds the nodes that are reachable from a source node at every time instant. The algorithm performs  $O(n)$  steps for each source node. The worst case complexity for each time instant would hence be  $O(n^2)$ . Since the search is performed for each time step, the computational complexity is  $O(n^2T)$ , where  $n$  is the number of nodes and  $T$  is the length of the time period.

**4. CONCLUSION AND FUTURE WORK**

The discovery of spatio-temporal patterns in a sensor graph raises several important questions that need to be answered before further analysis. First, how can we model space and time on a sensor graph? Second, can we determine any new or unusual patterns in the sensor graph? Third, which areas in the sensor graph have similar behavior? Finally, how can we predict which nodes in the sensor graph will have similar behavior?

In this paper, we proposed a spatio-temporal sensor graph (STSG) model to represent sensor data to answer all of these questions. We presented three case studies utilizing the STSG model to discover different types of patterns. First, we proposed an anomaly detection algorithm that generates a set of time series of where anomalies occur within the STSG. Second, a basic hotspot discovery algorithm is proposed using STSG to identify centralized locations at each time interval. Third, a growing hotspot method is proposed using STSG to predict nodes that may be hotspots at future time intervals.

The sensor graphs studied in this paper emphasized physical networks such as road or river networks. However, we expect that the Spatio-Temporal Sensor Graph might be applicable to other types of sensor graphs and this aspect will be investigated. Though the model can incorporate the stochastic nature of sensor graphs, the algorithms proposed

in the paper do not consider the probabilities in computations. We plan to incorporate appropriate modifications in the algorithms to account for the stochastic nature of the attributes.

### Acknowledgements

We are particularly grateful to the members of the Spatial Database Research Group at the University of Minnesota for their helpful comments and valuable discussions. We would also like to express our thanks to Kim Koffolt for improving the readability of this paper. We thank the reviewers for their suggestions that helped in bringing out the novelty of the content relative to the existing approaches.

This work is supported by the National Science Foundation (NSF-SEI grant), Oak Ridge National Laboratory, Topographical Engineers Corps and the Minnesota Department of Transportation.

## 5. REFERENCES

- [1] Personal communication with Prof. Miki Hondzo, Saint Anthony Falls Laboratory, University of Minnesota.
- [2] A. Arasu, M. Cherniack, E. Galvez, D. Maier, A. S. Maskey, E. Ryvkina, M. Stonebraker, and R. Tibettis. Linear Road: A Stream Data Management Benchmark. In *Proceedings of the 30th International Conference on Very Large Data Bases, Trondheim, Norway*, pages 480–491, Sept, 2005.
- [3] V. Barnett and T. Lewis. *Outliers in Statistical Data*. John Wiley, 3rd edition, 2003.
- [4] H. Bravo, J. Gulliver, and M. Hondzo. Development of a commercial code-based two-fluid model for bubble plumes. In *Environmental Modelling and Software*, volume 22, pages 536–547, 2007.
- [5] C. Chong and S. P. Kumar. Sensor networks: Evolution, opportunities, and challenges. *Proceedings of the IEEE*, 91(8):1247–1256, 2003.
- [6] P.S. Corso, M.H. Kramer, K.A Blair, D.G. Addiss, J.P. Davis, and A.C. Haddix. Cost of illness in the 1993 waterborne cryptosporidium outbreak. In *Emerging Infectious Diseases*, volume 9, pages 426–431, 2003.
- [7] D. Estrin, R. Govindan, J. Heidemann, and S. Kumar. Next century challenges: Scalable coordination in sensor networks. In *Proceedings of the 5th annual ACM/IEEE international conference on Mobile computing and networking*, pages 263–270, New York, NY, USA, 1999. ACM Press.
- [8] B. George, S. Kim, and S. Shekhar. Spatio-temporal Network Databases and Routing Algorithms: A Summary of Results. *To appear in the Proceedings of International Symposium on Spatial and Temporal Databases (SSTD'07)*, July 2007.
- [9] B. George and S. Shekhar. Time-aggregated Graphs for Modeling Spatio-Temporal Networks - An Extended Abstract. *Proceedings of Workshops at International Conference on Conceptual Modeling (ER 2006)*, November 2006.
- [10] D. Goldin. Spatial Queries over Sensor Networks, University of Connecticut, Department of Computer Science and Engineering Technical Report, March 2003.
- [11] C. Gotsman and Y. Koren. Distributed Graph Layout for Sensor Networks. *Journal of Graph Algorithms and Applications*, 9(1), 2005.
- [12] R. Govindan, J. M. Hellerstein, W. Hong, S. Madden, M. Franklin, and S. Shenker. The Sensor Network as a Database. *University of Southern California Technical Report*, (02-771), September 2002.
- [13] R. M. Hozalski, M. Hondzo, N. Jindal, M. J. Semmens, S. Shekhar, and G. L. Emmert. EFRI-ARESCI Preliminary Proposal: Integrated water quality management and treatment enabled by Cyberinfrastructure, 2006.
- [14] R. Johnson. *Applied Multivariate Statistical Analysis*. Prentice Hall, 1992.
- [15] D.E. Kaufman and R.L. Smith. Fastest Paths in Time-Dependent Networks for Intelligent Vehicle Highway Systems Applications. *IVHS Journal*, 1(1):1–11, 1993.
- [16] E. Kohler, K. Langtau, and Skutella M. Time-expanded graphs for flow-dependent transit times. *Proc. 10th Annual European Symposium on Algorithms*, pages 599–611, 2002.
- [17] K. Lorincz, D. J. Malan, T. R. F. Fulford-Jones, A. Nawoj, A. Clavel, V. Shnyder, G. Mainland, S. Moulton, and M. Welsh. Sensor Networks for Emergency Response: Challenges and Opportunities. In *IEEE Pervasive Computing, Special Issue on Pervasive Computing for First Response*, 2004.
- [18] M. Mauve, J. Widmer, and H. Hartenstein. A Survey on Position-based Routing in Mobile Ad-hoc Networks. *IEEE Network*, 15, 2001.
- [19] S. Pallottino and M.G. Scutella. Shortest Path Algorithms in Transportation Models: Classical and Innovative Aspects. *Equilibrium and Advanced transportation Modelling (Kluwer)*, pages 245–281, 1998.
- [20] K. Römer, O. Kasten, and F. Mattern. Middleware Challenges for Wireless Sensor Networks. *ACM Mobile Computing and Communication Review*, 6(4):59–61, October 2002.
- [21] J.C. Rutherford. *River Mixing*. John Wiley and Sons, Inc., 1994.
- [22] D. Sawitzki. Implicit Maximization of Flows over Time. *Technical report, University of Dortmund*, 2004.
- [23] S. Shekhar and S. Chawla. *Spatial Databases: A Tour*. Prentice Hall, 2003.
- [24] S. Shekhar, C. T. Lu, and P. Zhang. Detecting Graph-based Spatial Outliers: Algorithms and Applications. In *Proceedings of the 7th ACM International Conference on Knowledge Discovery and Data Mining*, 2001.
- [25] S. Shekhar, C. T. Lu, and P. Zhang. A Unified Approach to Spatial Outliers Detection. *Geoinformatica, An International Journal on Advances of Computer Science for Geographic Information System*, 7(2), June 2003.
- [26] P. Zhang, Y. Huang, S. Shekhar, and V. Kumar. Exploiting Spatial Autocorrelation to Efficiently Process Correlation-Based Similarity Queries. In *Proceedings of the 8th Symposium on Spatial and Temporal Databases*, July 25-27, 2003.