Technical Report: Graph-Structured Sparse Optimization for Connected Subgraph Detection

Baojian Zhou
Computer Science Department
University at Albany – SUNY
Albany, USA
bzhou6@albany.edu

Feng Chen
Computer Science Department
University at Albany – SUNY
Albany, USA
fchen5@albany.edu

Abstract—Structured sparse optimization is an important and challenging problem for analyzing high-dimensional data in a variety of applications such as bioinformatics, medical imaging, social networks, and astronomy. Although a number of structured sparsity models have been explored, such as trees, groups, clusters, and paths, connected subgraphs have been rarely explored in the current literature. One of the main technical challenges is that there is no *structured sparsity-inducing norm* that can directly model the space of connected subgraphs, and there is no exact implementation of a *projection oracle* for connected subgraphs due to its NP-hardness. In this paper, we explore efficient approximate projection oracles for connected subgraphs, and propose two new efficient algorithms, namely, GRAPH-IHT and GRAPH-GHTP, to optimize a generic nonlinear objective function subject to connectivity constraint on the support of the variables. Our proposed algorithms enjoy strong guarantees analogous to several current methods for sparsity-constrained optimization, such as Projected Gradient Descent (PGD), Approximate Model Iterative Hard Thresholding (AM-IHT), and Gradient Hard Thresholding Pursuit (GHTP) with respect to convergence rate and approximation accuracy. We apply our proposed algorithms to optimize several well-known graph scan statistics in several applications of connected subgraph detection as a case study, and the experimental results demonstrate that our proposed algorithms outperform state-of-the-art methods.

I. INTRODUCTION

In recent years, structured sparse methods have attracted much attention in many domains such as bioinformatics, medical imaging, social networks, and astronomy [2], [4], [14], [16], [17]. Structured sparse methods have been shown effective to identify latent patterns in high-dimensional data via the integration of prior knowledge about the structure of the patterns of interest, and at the same time remain a mathematically tractable concept. A number of structured sparsity models have been well explored, such as the sparsity models defined through trees [14], groups [17], clusters [16], and paths [2]. The generic optimization problem based on a structured sparsity model has the form

$$\min_{\mathbf{x} \in \mathbb{R}^n} f(\mathbf{x}) \quad s.t. \quad \text{supp}(\mathbf{x}) \in \mathbb{M} \quad (1)$$

where $f : \mathbb{R}^n \to \mathbb{R}$ is a differentiable cost function, the sparsity model $\mathbb{M}$ is defined as a family of structured supports: $\mathbb{M} = \{S_1, S_2, \cdots, S_L\}$, where $S_i \subseteq [n]$ satisfies a certain structure property (e.g., trees, groups, clusters), $[n] =$ $\{1, 2, \cdots, n\}$, and the support set $\text{supp}(\mathbf{x})$ refers to the set of indexes of non-zero entries in $\mathbf{x}$. For example, the popular k-sparsity model is defined as $\mathbb{M} = \{S \subseteq [n] \mid |S| \leq k\}$.

Existing structured sparse methods fall into two main categories: 1) **Sparsity-inducing norms based.** The methods in this category explore structured sparsity models (e.g., trees, groups, clusters, and paths) [4] that can be encoded as structured sparsity-inducing norms, and reformulate Problem (1) as a convex (or non-convex) optimization problem

$$\min_{\mathbf{x} \in \mathbb{R}^n} f(\mathbf{x}) + \lambda \cdot \Omega(\mathbf{x}) \quad (2)$$

where $\Omega(\mathbf{x})$ is a structured sparsity-inducing norm of $\mathbb{M}$ that is typically non-smooth and non-Euclidean and λ is a trade-off parameter. 2) **Model-projection based.** The methods in this category rely on a projection oracle of $\mathbb{M}$:

$$\mathrm{P}(\mathbf{b}) = \arg\min_{\mathbf{x} \in \mathbb{R}^n} \|\mathbf{b} - \mathbf{x}\|_2^2 \quad s.t. \quad \text{supp}(\mathbf{x}) \in \mathbb{M}, \quad (3)$$

and decompose the problem into two sub-problems, including unconstrained minimization of $f(\mathbf{x})$ and the projection problem $\mathrm{P}(\mathbf{b})$. Most of the methods in this category assume that the projection problem $\mathrm{P}(\mathbf{b})$ can be solved **exactly**, including the forward-backward algorithm [42], the gradient descent algorithm [38], the gradient hard-thresholding algorithms [6], [18], [40], the projected iterative hard thresholding [5], [7], and the Newton greedy pursuit algorithm [41]. However, when an exact solver of $\mathrm{P}(\mathbf{b})$ is unavailable and we have to apply approximate projections, the theoretical guarantees of these methods do not hold any more. We note that there is one recent approach named as GRAPH-COSAMP that admits inexact projections by assuming "head" and "tail" oracles for the projections, but is only applicable to compressive sensing or linear regression problems [15].

We consider an underlying graph $\mathbb{G} = (\mathbb{V}, \mathbb{E})$ defined on the coefficients of the unknown vector $\mathbf{x}$, where $\mathbb{V} = [n]$ and $\mathbb{E} \subseteq \mathbb{V} \times \mathbb{V}$. We focus on the sparsity model of connected subgraphs that is defined as

$$\mathbb{M}(\mathbb{G}, k) = \{S \subseteq \mathbb{V} \mid |S| \leq k, S \text{ is connected}\}, \quad (4)$$

where k refers to the allowed maximum subgraph size. There are a wide array of applications that involve the search of interesting or anomalous connected subgraphs in networks. The connectivity constraint ensures that subgraphs reflect changes due to localized in-network processes. We describe

a few applications below.

- *Detection in sensor networks,* e.g., detection of traffic bottlenecks in road networks or airway networks [1]; crime hot spots in geographic networks [22]; and pollutions in water distribution networks [27].
- *Detection in digital signals and images,* e.g., detection of objects in images [15].
- *Disease outbreak detection,* e.g., early detection of disease outbreaks from information networks incorporating data from hospital emergency visits, ambulance dispatch calls and pharmacy sales of over-the-counter drugs [36].
- *Virus detection in a computer network,* e.g., detection of viruses or worms spreading from host to host in a computer network [24].
- *Detection in genome-scale interaction network,* e.g., detection of significantly mutated subnetworks [21].
- *Detection in social media networks,* e.g., detection and forecasting of societal events [8], [9].

To the best of our knowledge, there is no existing approach to Problem (1) for $\mathbb{M}(\mathbb{G}, k)$ that is computationally tractable and provides performance bound. **First**, there is no known structured sparsity-inducing norm for $\mathbb{M}(\mathbb{G}, k)$. The most relevant norm is fused lasso norm [39]: $\Omega(\mathbf{x}) = \sum_{(i,j) \in \mathbb{E}} |x_i - x_j|$, where x_i is the i-th entry in $\mathbf{x}$. This norm is able to enforce the smoothness between neighboring entries in $\mathbf{x}$, but has limited capability to recover all the possible connected subsets as described by $\mathbb{M}(\mathbb{G}, k)$ (See further discussions in Section V Experiments). **Second**, there is no exact solver for the projection oracle of $\mathbb{M}(\mathbb{G}, k)$:

$$P(\mathbf{x}) = \arg \min_{\mathbf{x} \in \mathbb{R}^n} \|\mathbf{b} - \mathbf{x}\|_2^2 \quad s.t. \quad \text{supp}(\mathbf{x}) \in \mathbb{M}(\mathbb{G}, k), \quad (5)$$

as this projection problem is NP-hard due to a reduction from classical Steiner tree problem [19]. As most existing model-projection based methods require an exact solution to the projection oracle $P(\mathbf{x})$, these methods are inapplicable to the problem studied here. To the best of our knowledge, there is only one recent approach named as GRAPH-COSAMP that admits inexact projections for $\mathbb{M}(\mathbb{G}, k)$ by assuming "head" and "tail" oracles for the projections, but is only applicable to compression sensing and linear regression problems [15]. The main contributions of our study are summarized as follows:

- **Design of efficient approximation algorithms.** Two new algorithms, namely, GRAPH-IHT and GRAPH-GHTP, are developed to approximately solve Problem (1) that has a differentiable cost function and a sparsity model of connected subgraphs $\mathbb{M}(\mathbb{G}, k)$. GRAPH-GHTP is required to minimize $f(\mathbf{x})$ over a projected subspace as an intermediate step, which could be too costly in some applications and GRAPH-IHT could be considered as a fast variant of GRAPH-GHTP.
- **Theoretical guarantees and connections.** The convergence rate and accuracy of our proposed algorithms are analyzed under a smoothness condition of $f(\mathbf{x})$ that is more general than popular conditions such as Restricted Strong Convexity/Smoothness (RSC/RSS) and Stable

Mode Restricted Hessian (SMRH). We prove that under mild conditions our proposed GRAPH-IHT and GRAPH-GHTP enjoy rigorous theoretical guarantees.
- **Compressive experiments to validate the effectiveness and efficiency of the proposed techniques.** Both GRAPH-IHT and GRAPH-GHTP are applied to optimize a variety of graph scan statistic models for the connected subgraph detection task. Extensive experiments on a number of benchmark datasets demonstrate that GRAPH-IHT and GRAPH-GHTP perform superior to state-of-the-art methods that are designed specifically for this task in terms of subgraph quality and running time.

Reproducibility: The implementation of our algorithms and the data sets is open-sourced via the link [11].

The remaining parts of this paper are organized as follows. Sections II introduces the sparsity model of connected subgraphs and statement of the problem. Sections III presents two efficient algorithms and their theoretical analysis. Section IV discusses applications of our proposed algorithms to graph scan statistic models. Experiments on several real world benchmark datasets are presented in Section V. Section VI discusses related work and Section VII concludes the paper and describes future work.

II. PROBLEM FORMULATION

Given an underlying graph $\mathbb{G} = (\mathbb{V}, \mathbb{E})$ defined on the coefficients of the unknown vector $\mathbf{x}$, where $\mathbb{V} = [n]$, $\mathbb{E} \subseteq \mathbb{V} \times \mathbb{V}$, and n is typically large (e.g., $n > 10,000$). The sparsity model of connected subgraphs in $\mathbb{G}$ is defined in (4), and its projection oracle $P(\mathbf{x})$ is defined in (5). As this projection oracle is NP-hard to solve, we first introduce efficient approximation algorithms for $P(\mathbf{x})$ and then present statement of the problem that will be studied in the paper.

A. Approximation algorithms for the projection oracle $P(\mathbf{x})$

There are two nearly-linear time approximation algorithms [15] for $P(\mathbf{x})$ that have the following properties:

- **Tail approximation** ($T(\mathbf{x})$): Find a $S \subseteq \mathbb{V}$ such that

$$\|\mathbf{x} - \mathbf{x}_S\|_2 \leq c_T \cdot \min_{S' \in \mathbb{M}(\mathbb{G}, k_T)} \|\mathbf{x} - \mathbf{x}_{S'}\|_2, \quad (6)$$

where $c_T = \sqrt{7}$, $k_T = 5k$, and $\mathbf{x}_S$ is the restriction of $\mathbf{x}$ to indices in S: we have $(\mathbf{x}_S)_i = x_i$ for $i \in S$ and $(\mathbf{x}_S)_i = 0$ otherwise.
- **Head approximation** ($H(\mathbf{x})$): Find a $S \subseteq \mathbb{V}$ such that

$$\|\mathbf{x}_S\|_2 \geq c_H \cdot \max_{S' \in \mathbb{M}(\mathbb{G}, k_H)} \|\mathbf{x}_{S'}\|_2, \quad (7)$$

where $c_H = \sqrt{1/14}$ and $k_H = 2k$.

It can be readily proved that, if $c_T = c_H = 1$, then $T(\mathbf{x}) = H(\mathbf{x}) = P(\mathbf{x})$, which indicates that these two approximations ($T(\mathbf{x})$ and $H(\mathbf{x})$) stem from the fact that $c_T > 1$ and $c_H < 1$.

B. Problem statement

Given a predefined cost function $f(\mathbf{x})$ that is differentiable, the input graph $\mathbb{G}$, and the sparsity model of connected
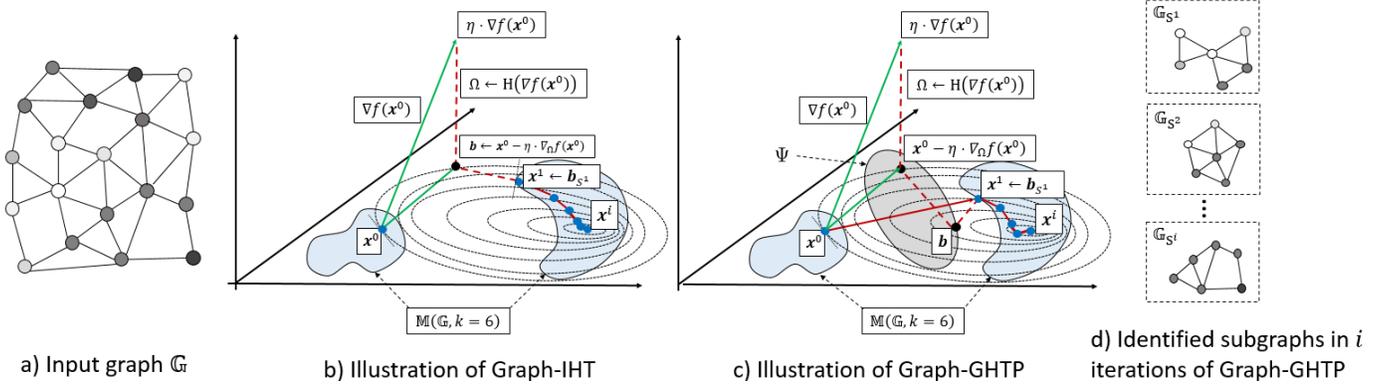
Fig. 1. Illustration of GRAPH-IHT and GRAPH-GHTP on the main steps of each iteration. In this example, the gray scale of each node i encodes the weight of this node $w_i \in \mathbb{R}$, $f(\mathbf{x}) = -\mathbf{w}^\mathsf{T}\mathbf{x} + \frac{1}{2}\|\mathbf{x}\|^2$, and the maximum size of subgraphs is set to $k = 6$, where $\mathbf{w} = [w_1, \cdots, w_n]^T$. The resulting problem tends to find a connected subgraph with the largest overall weight (See discussion about the effect of $\frac{1}{2}\|x\|^2$ in Section IV). In each iteration, S^i is the connected subset of nodes and its induced subgraph is denoted as $\mathbb{G}_{S^i}$. The sequence of intermediate vectors and subgraphs are: $(\mathbf{x}^0, \mathbb{G}_{S^1}), \cdots, (\mathbf{x}^i, \mathbb{G}_{S^i})$.

subgraph $\mathbb{M}(\mathbb{G}, k)$, the problem to be studied is formulated as:

$$\min_{\mathbf{x} \in \mathbb{R}^n} f(\mathbf{x}) \quad s.t. \quad \text{supp}(\mathbf{x}) \in \mathbb{M}(\mathbb{G}, k). \qquad (8)$$

Problem (8) is difficult to solve as it involves decision variables from a nonconvex set that is composed of many disjoint subsets. In this paper, we will develop nearly-linear time algorithms to approximately solve Problem (8). The key idea is decompose this problem to sub-problems that are easier to solve. These sub-problems include an optimization sub-problem of $f(\mathbf{x})$ that is independent on $\mathbb{M}(\mathbb{G}, k)$ and projection approximations for $\mathbb{M}(\mathbb{G}, k)$, including $\mathrm{T}(\mathbf{x})$ and $\mathrm{H}(\mathbf{x})$. We will design efficient algorithms to couple these sub-problems to obtain global solutions to Problem (8) with good trade-off on running time and accuracy.

III. ALGORITHMS

This section first presents two efficient algorithms, namely, GRAPH-IHT and GRAPH-GHTP, and then analyzes their time complexities and performance bounds.

A. Algorithm GRAPH-IHT

The proposed GRAPH-IHT algorithm generalizes the traditional algorithm named as projected gradient descent [5], [7] that requires a exact solver of the projection oracle $\mathrm{P}(\mathbf{x})$. The high-level summary of GRAPH-IHT is shown in Algorithm 1 and illustrated in Figure 1 (b). The procedure generates a sequence of intermediate vectors $\mathbf{x}^0$, $\mathbf{x}^1$, $\cdots$ from an initial approximation $\mathbf{x}^0$. At the i-th iteration, the **first step** (Line 5) first calculates the gradient "$\nabla f(\mathbf{x}^i)$", and then identifies a subset of nodes via head approximation that returns a support set with the head value at least a constant factor of the optimal head value: "$\Omega \leftarrow \mathrm{H}(\nabla f(\mathbf{x}^i))$". The support set Ω can be interpreted as the subspace where the nonconvex set "$\{\mathbf{x} \mid \text{supp}(\mathbf{x}) \in \mathbb{M}(\mathbb{G}, k)\}$" is located, and the projected gradient in this subspace is: "$\nabla_\Omega f(\mathbf{x}^i)$". The **second step** (Line 6) calculates the projected gradient descent at the point $\mathbf{x}^i$ with step-size η: "$\mathbf{b} \leftarrow \mathbf{x}^i - \eta \cdot \nabla_\Omega f(\mathbf{x}^i)$". The **third step** (Line 7)

identifies a subset of nodes via tail approximation that returns a support set with tail value at most a constant times larger than the optimal tail value: "$S^{i+1} \leftarrow \mathrm{T}(\mathbf{b})$". The **last step** (Line 8) calculates the intermediate solution $\mathbf{x}^{i+1}$: $\mathbf{x}^{i+1} = \mathbf{b}_{S^{i+1}}$. The previous two steps can be interpreted as the projection of $\mathbf{b}$ to the nonconvex set "$\{\mathbf{x} \mid \text{supp}(\mathbf{x}) \in \mathbb{M}(\mathbb{G}, k)\}$" using the tail approximation.

Algorithm 1 GRAPH-IHT

1: **Input**: Input graph $\mathbb{G}$, maximum subgraph size k, and step size η (1 by default).
2: **Output**: The estimated vector $\hat{\mathbf{x}}$ and the corresponding connected subgraph $\mathbb{S}$.
3: $i \leftarrow 0$, $\mathbf{x}^i \leftarrow \mathbf{0}$; $S^i \leftarrow \emptyset$;
4: **repeat**
5: $\quad \Omega \leftarrow \mathrm{H}(\nabla f(\mathbf{x}^i))$;
6: $\quad \mathbf{b} \leftarrow \mathbf{x}^i - \eta \cdot \nabla_\Omega f(\mathbf{x}^i)$;
7: $\quad S^{i+1} \leftarrow \mathrm{T}(\mathbf{b})$;
8: $\quad \mathbf{x}^{i+1} \leftarrow \mathbf{b}_{S^{i+1}}$;
9: $\quad i \leftarrow i + 1$;
10: **until** halting condition holds
11: **return** $\hat{\mathbf{x}} = \mathbf{x}^i$ and $\mathbb{S} = \mathbb{G}_{S^i}$;

B. Algorithm GRAPH-GHTP

The proposed GRAPH-GHTP algorithm generalizes the traditional algorithm named as Gradient Hard Threshold Pursuit (GHTP) that is designed specifically for the k-sparsity model: $\mathbb{M} = \{S \subseteq [n] \mid |S| \le k\}$ [40]. The high-level summary of GRAPH-GHTP is shown in Algorithm 2 and illustrated in Figure 1 (c). The **first two steps** (Line 5 and Line 6) in each iteration is the same as the first two steps (Line 5 and Line 6) of GRAPH-IHT, except that we return the support of the projected gradient descent: "$\Psi \leftarrow \text{supp}(\mathbf{x}^i - \eta \cdot \nabla_\Omega f(\mathbf{x}^i))$", in which pursuing the minimization will be most effective. Over the support set S, the function f is minimized to produce an intermediate estimate at the **third** step (Line 7):

Algorithm 2 GRAPH-GHTP

1: **Input**: Input graph $\mathbb{G}$, maximum subgraph size k, and step size η (1 by default).
2: **Output**: The estimated vector $\hat{\mathbf{x}}$ and the corresponding connected subgraph $\mathbb{S}$.
3: $i \leftarrow 0$, $\mathbf{x}^i \leftarrow \mathbf{0}$; $S^i \leftarrow \emptyset$;
4: **repeat**
5:     $\Omega \leftarrow \mathrm{H}(\nabla f(\mathbf{x}^i))$
6:     $\Psi \leftarrow \mathrm{supp}(\mathbf{x}^i - \eta \cdot \nabla_\Omega f(\mathbf{x}^i))$;
7:     $\mathbf{b} \leftarrow \arg\min_{\mathbf{x} \in \mathbb{R}^n} f(\mathbf{x}) \ s.t. \ \mathrm{supp}(\mathbf{x}) \subseteq \Psi$;
8:     $S^{i+1} \leftarrow \mathrm{T}(\mathbf{b})$;
9:     $\mathbf{x}^{i+1} \leftarrow \mathbf{b}_{S^{i+1}}$;
10:     $i \leftarrow i + 1$;
11: **until** halting condition holds
12: **return** $\hat{\mathbf{x}} = \mathbf{x}^i$ and $\mathbb{S} = \mathbb{G}_{S^i}$;

"$\mathbf{b} \leftarrow \arg\min_{\mathbf{x} \in \mathbb{R}^n} f(\mathbf{x}) \ s.t. \ \mathrm{supp}(\mathbf{x}) \subseteq \Omega$". The **fourth and fifth** steps (Line 8 and Line 9) are the same as the last two steps (Line 7 and Line 8) of GRAPH-IHT in each iteration.

C. Relations between GRAPH-IHT and GRAPH-GHTP

These two algorithms are both variants of gradient descent. In overall, GRAPH-GHTP converges faster than GRAPH-IHT as it identifies a better intermediate solution in each iteration by minimizing $f(\mathbf{x})$ over a projected subspace $\{\mathbf{x} \mid \mathrm{supp}(\mathbf{x}) \subseteq \Omega\}$. If the cost function $f(\mathbf{x})$ is linear or has some special structure, this intermediate step can be conducted in nearly-linear time. However, when this step is too costly in some applications, GRAPH-IHT is preferred.

D. Theoretical Analysis of GRAPH-IHT

In order to demonstrate the accuracy of estimates using Algorithm 1, we require that the cost function $f(\mathbf{x})$ satisfies the Weak Restricted Strong Convexity (WRSC) condition as follows:

Definition III.1 (Weak Restricted Strong Convexity Property (WRSC))**.** A function $f(\mathbf{x})$ has the $(\xi, \delta, \mathbb{M})$-model-WRSC if $\forall \mathbf{x}, \mathbf{y} \in \mathbb{R}^n$ and $\forall S \in \mathbb{M}$ with $\mathrm{supp}(\mathbf{x}) \cup \mathrm{supp}(\mathbf{y}) \subseteq S$, the following inequality holds for some $\xi > 0$ and $0 < \delta < 1$:

$$\|\mathbf{x} - \mathbf{y} - \xi \nabla_S f(\mathbf{x}) + \xi \nabla_S f(\mathbf{y})\|_2 \le \delta \|\mathbf{x} - \mathbf{y}\|_2. \qquad (9)$$

The WRSC is weaker than the popular *Restricted Strong Convexity/Smoothness* (RSC/RSS) conditions that are used in theoretical analysis of convex optimization algorithms [40]. The RSC condition basically characterizes cost functions that have quadratic bounds on the derivative of the objective function when restricted to model-sparse vectors. The RSC/RSS conditions imply condition WRSC, which indicates that WRSC is no stronger than RSC/RSS [40]. In the special case where $f(\mathbf{x}) = \|\mathbf{y} - A\mathbf{x}\|_2^2$ and $\xi = 1$, the condition $(\xi, \delta, \mathbb{M})$-model-WRSC reduces to the well known Restricted Isometry Property (RIP) condition in compressive sensing.

Theorem III.1. *Consider the sparsity model of connected subgraphs $\mathbb{M}(\mathbb{G}, k)$ for some $k \in \dot{\mathbb{N}}$ and a cost function $f : \mathbb{R}^n \to \mathbb{R}$ that satisfies the $(\xi, \delta, \mathbb{M}(\mathbb{G}, 5k))$-model-WRSC*

condition. If $\eta = c_H(1 - \delta) - \delta$ then for any $\mathbf{x} \in \mathbb{R}^n$ such that $supp(\mathbf{x}) \in \mathbb{M}(\mathbb{G}, k)$, with $\eta > 0$ the iterates of Algorithm 2 obey

$$\|\mathbf{x}^{i+1} - \mathbf{x}\|_2 \le \alpha \|\mathbf{x}^i - \mathbf{x}\|_2 + \beta \|\nabla_I f(\mathbf{x})\|_2 \qquad (10)$$

where

$$\alpha_0 = c_H(1 - \delta) - \delta, \beta_0 = \delta(1 + c_H),$$

$$\alpha = \frac{\sqrt{2}(1 + c_T)}{1 - \delta} \left(\sqrt{1 - \alpha_0^2} + \left((2 - \frac{\eta}{\xi})\delta + 1 - \frac{\eta}{\xi} \right) \right),$$

$$\beta = \frac{1 + c_T}{1 - \delta} \left((1 + 2\sqrt{2})\xi + (2 - 2\sqrt{2})\eta + \frac{\sqrt{2}\beta_0}{\alpha_0} + \frac{\sqrt{2}\alpha_0\beta_0}{\sqrt{1 - \alpha_0^2}} \right),$$

and $I = \arg\max_{S \in \mathbb{M}(\mathbb{G}, 8k)} \|\nabla_S f(x)\|_2$

Before we prove this result, we give the following two lemmas III.2 and III.3.

Lemma III.2. *[40] Assume that f is a differentiable function. If f satisfies condition $(\xi, \delta, \mathbb{M})$-WRSC, then $\forall \mathbf{x}, \mathbf{y} \in \mathbb{R}^n$ with $supp(\mathbf{x}) \cup supp(\mathbf{y}) \subset S \in \mathbb{M}$, the following two inequalities hold*

$$\frac{1 - \delta}{\xi} \|\mathbf{x} - \mathbf{y}\|_2 \le \|\nabla_S f(\mathbf{x}) - \nabla_S f(\mathbf{y})\|_2 \le \frac{1 + \delta}{\xi} \|\mathbf{x} - \mathbf{y}\|_2$$

$$f(\mathbf{x}) \le f(\mathbf{y}) + \langle \nabla f(\mathbf{y}), \mathbf{x} - \mathbf{y} \rangle + \frac{1 + \delta}{2\xi} \|\mathbf{x} - \mathbf{y}\|_2^2$$

Lemma III.3. *Let $\alpha_0 = c_H(1 - \delta) - \delta$, $\beta_0 = \xi(1 + c_H)$, $\mathbf{r}^i = \mathbf{x}^i - \mathbf{x}$, and $\Omega = H(\nabla f(\mathbf{x}^i))$. Then*

$$\|\mathbf{r}_{\Gamma^c}^i\|_2 \le \sqrt{1 - \alpha_0^2} \|\mathbf{r}^i\|_2 + \left[\frac{\beta_0}{\alpha_0} + \frac{\alpha_0\beta_0}{\sqrt{1 - \alpha_0^2}} \right] \|\nabla_I f(\mathbf{x})\|_2$$

where $I = \arg\max_{S \in \mathbb{M}(\mathbb{G}, 8k)} \|\nabla_S f(\mathbf{x})\|_2$. We assume that c_H and δ are such that $\alpha_0 > 0$.

Proof: Denote $\Phi = \mathrm{supp}(\mathbf{x}) \in \mathbb{M}(\mathbb{G}, k)$, $\Omega = H(\nabla f(\mathbf{x}^i)) \in \mathbb{M}(\mathbb{G}, 2k)$, $\mathbf{r}^i = \mathbf{x}^i - \mathbf{x}$, and $\Lambda = \mathrm{supp}(\mathbf{r}^i) \in \mathbb{M}(\mathbb{G}, 6k)$. The component $\|\nabla_\Gamma f(\mathbf{x}^i)\|_2$ can be lower bounded as

$$
\begin{aligned}
\|\nabla_\Gamma f(\mathbf{x}^i)\|_2 &\ge c_H \|\nabla_\Phi f(\mathbf{x}^i)\|_2 \\
&\ge c_H (\|\nabla_\Phi f(\mathbf{x}^i) - \nabla_\Phi f(\mathbf{x})\|_2 - \|\nabla_\Phi f(\mathbf{x})\|_2 \\
&\ge \frac{c_H(1 - \delta)}{\xi} \|\mathbf{r}^i\|_2 - c_H \|\nabla_I f(\mathbf{x})\|_2,
\end{aligned}
$$

where the first inequality follows from the definition of head approximation and the last inequality follows from Lemma III.2 of our paper. The component $\|\nabla_\Gamma f(\mathbf{x}^i)\|_2$ can also be upper bounded as

$$
\begin{aligned}
\|\nabla_\Gamma f(\mathbf{x}^i)\|_2 &\le \frac{1}{\xi} \|\xi \nabla_\Gamma f(\mathbf{x}^i) - \xi \nabla_\Gamma f(\mathbf{x})\|_2 + \|\nabla_\Gamma f(\mathbf{x})\|_2 \\
&\le \frac{1}{\xi} \|\xi \nabla_\Gamma f(\mathbf{x}^i) - \xi \nabla_\Gamma f(\mathbf{x}) - \mathbf{r}_\Gamma^i + \mathbf{r}_\Gamma^i\|_2 + \\
&\quad \|\nabla_\Gamma f(\mathbf{x})\|_2 \\
&\le \frac{1}{\xi} \|\xi \nabla_{\Gamma \cup \Omega} f(\mathbf{x}^i) - \xi \nabla_{\Gamma \cup \Omega} f(\mathbf{x}) - \mathbf{r}_{\Gamma \cup \Omega}^i\|_2 + \\
&\quad \|\mathbf{r}_\Gamma^i\|_2 + \|\nabla_\Gamma f(\mathbf{x})\|_2 \\
&\le \frac{\delta}{\xi} \cdot \|\mathbf{r}^i\|_2 + \frac{1}{\xi} \|\mathbf{r}_\Gamma^i\|_2 + \|\nabla_I f(\mathbf{x})\|_2,
\end{aligned}
$$

where the fourth inequality follows from condition $(\xi, \delta, \mathbb{M}(\mathbb{G}, 8k))$-WRSC and the fact that $\mathbf{r}_{\Gamma \cup \Omega}^i = \mathbf{r}^i$. Combining the two bounds and grouping terms, we obtain

the inequality:
$$\|\mathbf{r}_\Gamma^i\| \geq \alpha_0 \|\mathbf{r}^i\|_2 - \xi(1+c_H)\|\nabla_I f(\mathbf{x})\|_2.$$
We have $\|\mathbf{r}_\Gamma^i\| \geq \alpha_0 \|\mathbf{r}^i\|_2 - \beta_0 \|\nabla_I f(\mathbf{x})\|_2$. After a number of algebraic manipulations, we obtain the inequality
$$\|\mathbf{r}_{\Gamma^c}^i\|_2 \leq \sqrt{1-\alpha_0^2}\|\mathbf{r}^i\|_2 + \left[\frac{\beta_0}{\alpha_0} + \frac{\alpha_0\beta_0}{\sqrt{1-\alpha_0^2}}\right]\|\nabla_I f(\mathbf{x})\|_2,$$
which proves the lemma. ∎

We give the formal proof of III.1.

Proof: From the traingle inequality, we have
$$
\begin{aligned}
\|\mathbf{r}^{i+1}\|_2 &= \|\mathbf{x}^{i+1} - \mathbf{x}\|_2 \\
&= \|\mathbf{b}_\Psi - \mathbf{x}\|_2 \\
&\leq \|\mathbf{b} - \mathbf{x}\|_2 + \|\mathbf{b} - \mathbf{b}_\Psi\|_2 \\
&\leq (1+c_T)\|\mathbf{b} - \mathbf{x}\|_2 \\
&= (1+c_T)\|\mathbf{x}^i - \eta\nabla_\Omega f(\mathbf{x}^i) - \mathbf{x}^i\|_2 \\
&= (1+c_T)\|\mathbf{r}^i - \eta\nabla_\Omega f(\mathbf{x}^i)\|_2,
\end{aligned}
$$
where $\nabla_\Omega f(\mathbf{x}^i)$ is the projected vector of $f(\mathbf{x}^i)$ in which the entries outoside Ω are set to zero and the entries in Ω are unchanged. $\|\mathbf{r}^i - \eta\nabla_\Omega f(\mathbf{x}^i)\|_2$ has the inequalities
$$\|\mathbf{r}^i - \eta\nabla_\Omega f(\mathbf{x}^i)\|_2 = \|\mathbf{r}_{\Omega^c}^i + \mathbf{r}_\Omega^i - \eta\nabla_\Omega f(\mathbf{x}^i)\|$$
$$\leq \|\mathbf{r}_{\Omega^c}^i\|_2 + \|\mathbf{r}_\Omega^i - \eta\nabla_\Omega f(\mathbf{x}^i) + \eta\nabla_\Omega f(\mathbf{x}) - \eta\nabla_\Omega f(\mathbf{x})\|$$
$$\leq \|\mathbf{r}_{\Omega^c}^i\|_2 + \|\mathbf{r}_\Omega^i - \eta\nabla_\Omega f(\mathbf{x}^i) + \eta\nabla_\Omega f(\mathbf{x})\| + \|\eta\nabla_\Omega f(\mathbf{x})\|$$
$$\leq \|\mathbf{r}_{\Omega^c}^i\|_2 + \|\mathbf{r}_\Omega^i - \xi\nabla_\Omega f(\mathbf{x}^i) + \xi\nabla_\Omega f(\mathbf{x})\| +$$
$$(\xi - \eta)\|\nabla_\Omega f(\mathbf{x}^i) - \nabla_\Omega f(\mathbf{x})\|_2 + \|\eta\nabla_\Omega f(\mathbf{x})\|_2$$
$$\leq \|\mathbf{r}_{\Omega^c}^i\|_2 + (1 - \eta/\xi + (2 - \eta/\xi)\delta)\|\mathbf{r}^i\|_2 + \eta\|\nabla_I f(\mathbf{x})\|_2$$
where the last inequality follows from condition $(\xi,\delta,\mathbb{M})$-WRSC and Lemma III.2. From Lemma III.3, we have
$$\|\mathbf{r}_{\Gamma^c}^i\|_2 \leq \sqrt{1-\alpha_0^2}\|\mathbf{r}^i\|_2 + \left[\frac{\beta_0}{\alpha_0} + \frac{\alpha_0\beta_0}{\sqrt{1-\alpha_0^2}}\right]\|\nabla_I f(\mathbf{x})\|_2$$

Combining the above inequalities, we prove the theorem. ∎

Our proposed GRAPH-IHT generalizes several existing sparsity-constrained optimization algorithms: 1) **Projected Gradient Descent (PGD)** [30]. If we redefine $\mathbf{H}(b) = \text{supp}(b)$ and $\mathbf{T}(b) = \text{supp}(\mathbf{P}(b))$, where $\mathbf{P}(b)$ is the projection oracle defined in Equation (5), then GRAPH-IHT reduces to the PGD method; 2) **Approximated Model-IHT(AM-IHT)** [13]. If the cost function $f(\mathbf{x})$ is defined as the least square cost function $f(\mathbf{x}) = \|\mathbf{y} - \mathbf{A}\mathbf{x}\|_2^2$, then $\nabla f(\mathbf{x})$ has the specific form $-\mathbf{A}^\top(\mathbf{y} - \mathbf{A})$ and GRAPH-IHT reduces to the AM-IHT algorithm, the state-of-the-art variant of IHT for compressive sensing and linear regression problems. In particular, let $\mathbf{e} = \mathbf{y} - \mathbf{A}\mathbf{x}$. The component $\|\nabla f(\mathbf{x}^i)\|_2 = \|\mathbf{A}^\top\mathbf{e}\|_2$ is upper bound by bounded by $\sqrt{1+\delta}\|\mathbf{e}\|_2$ [13], Assume that $\xi = 1$ and $\eta = 1$. Condition $(\xi,\eta,\mathbb{M})$-WRSC then reduces to the RIP condition in compressive sensing. The convergence inequality (10) then reduces to
$$\|\mathbf{x}^{i+1} - \mathbf{x}\|_2 \leq \alpha'\|\mathbf{x}^i - \mathbf{x}\|_2 + \beta'\|\mathbf{e}\|_2, \qquad (11)$$
where $\alpha' = (1+c_T)\left[\delta + \sqrt{1-\alpha_0^2}\right]$ and
$$\beta' = (1+c_T)\left[\frac{(\alpha_0 + \beta_0)\sqrt{1+\delta}}{\alpha_0} + \frac{\alpha_0\beta_0(\sqrt{1+\delta})}{\sqrt{1-\alpha_0^2}}\right].$$

Surprisingly, the above convergence inequality **is identical to the convergence inequality of AM-IHT** derived in [13] based on the RIP condition, which indicates that GRAPH-IHT has the same convergence rate and approximation error as AM-IHT, although we did not make any attempt to explore the special properties of the RIP condition. We note that the convergence properties of GRAPH-IHT hold in fairly general setups beyond compressive sensing and linear regression. As we consider GRAPH-IHT as a fast variant of GRAPH-GHTP, due to space limit we ignore the discussions about the convergence condition of GRAPH-IHT. The theoretical analysis of GRAPH-GHTP to be discussed in the next subsection can be readily adapted to the theoretical analysis of GRAPH-IHT.

E. Theoretical Analysis of GRAPH-GHTP

Theorem III.4. *Consider the sparsity model of connected subgraphs* $\mathbb{M}(\mathbb{G},k)$ *for some* $k \in \mathbb{N}$ *and a cost function* $f : \mathbb{R}^n \to \mathbb{R}$ *that satisfies the* $(\xi,\delta,\mathbb{M}(\mathbb{G},5k))$-*model-WRSC condition. If* $\eta = c_H(1-\delta) - \delta$ *then for any* $\mathbf{x} \in \mathbb{R}^n$ *such that* $\text{supp}(\mathbf{x}) \in \mathbb{M}(\mathbb{G},k)$, *with* $\eta > 0$ *the iterates of Algorithm 2 obey*
$$\|\mathbf{x}^{i+1} - \mathbf{x}\|_2 \leq \alpha\|\mathbf{x}^i - \mathbf{x}\|_2 + \beta\|\nabla_I f(\mathbf{x})\|_2 \qquad (12)$$
where
$$\alpha_0 = c_H(1-\delta) - \delta, \beta_0 = \delta(1+c_H),$$
$$\alpha = \frac{\sqrt{2}(1+c_T)}{1-\delta}\left(\sqrt{1-\alpha_0^2} + \left((2-\frac{\eta}{\xi})\delta + 1 - \frac{\eta}{\xi}\right)\right),$$
$$\beta = \frac{1+c_T}{1-\delta}\left((1+2\sqrt{2})\xi + (2-2\sqrt{2})\eta + \frac{\sqrt{2}\beta_0}{\alpha_0} + \frac{\sqrt{2}\alpha_0\beta_0}{\sqrt{1-\alpha_0^2}}\right),$$
and $I = \arg\max_{S \in \mathbb{M}(\mathbb{G},8k)}\|\nabla_S f(x)\|_2$

Proof: Denote $\Omega = \mathbf{H}(\nabla f(\mathbf{x}^i))$ and $\Psi = \text{supp}(\mathbf{x}^i - \eta \cdot \nabla_\Omega f(\mathbf{x}^i))$. Let $\mathbf{r}^{i+1} = \mathbf{x}^{i+1} - \mathbf{x}$. $\|\mathbf{r}^{i+1}\|_2$ is bounded as
$$
\begin{aligned}
\|\mathbf{r}^{i+1}\|_2 = \|\mathbf{x}^{i+1} - \mathbf{x}\|_2 &\leq \|\mathbf{x}^{i+1} - \mathbf{b}\|_2 + \|\mathbf{x} - \mathbf{b}\|_2 \\
&\leq c_T\|\mathbf{x} - \mathbf{b}\|_2 + \|\mathbf{x} - \mathbf{b}\|_2 \\
&\leq (1+c_T)\|\mathbf{x} - \mathbf{b}\|_2, \qquad (13)
\end{aligned}
$$
where the second inequality follows from the definition of tail approximation. The component $\|(\mathbf{x} - \mathbf{b})_\Psi\|_2^2$ is bounded as
$$\|(\mathbf{x} - \mathbf{b})_\Psi\|_2^2 = \langle \mathbf{b} - \mathbf{x}, (\mathbf{b} - \mathbf{x})_\Psi\rangle$$
$$= \langle \mathbf{b} - \mathbf{x} - \xi\nabla_\Psi f(\mathbf{b}) + \xi\nabla_\Psi f(\mathbf{x}), (\mathbf{b} - \mathbf{x})_\Psi\rangle -$$
$$\langle \xi\nabla_\Psi f(\mathbf{x}), (\mathbf{b} - \mathbf{x})_\Psi\rangle$$
$$\leq \delta\|\mathbf{b} - \mathbf{x}\|_2\|(\mathbf{b} - \mathbf{x})_\Psi\|_2 + \xi\|\nabla_\Psi f(\mathbf{x})\|_2\|(\mathbf{b} - \mathbf{x})_\Psi\|_2,$$
where the second equality follows from the fact that $\nabla_S f(\mathbf{b}) = \mathbf{0}$ since $\mathbf{b}$ is the solution to the problem in the third Step (Line 7) of GRAPH-GHTP, and the last inequality can be derived from condition $(\xi,\delta,\mathbb{M}(\mathbb{G},8k))$-WRSC. After simplification, we have
$$\|(\mathbf{x} - \mathbf{b})_\Psi\|_2 \leq \delta\|\mathbf{b} - \mathbf{x}\|_2 + \xi\|\nabla_\Psi f(\mathbf{x})\|_2.$$
It follows that
$$\|\mathbf{x} - \mathbf{b}\|_2 \leq \|(\mathbf{x} - \mathbf{b})_\Psi\|_2 + \|(\mathbf{x} - \mathbf{b})_{\Psi^c}\|_2$$
$$\leq \delta\|\mathbf{b} - \mathbf{x}\|_2 + \xi\|\nabla_\Psi f(\mathbf{x})\|_2 + \|(\mathbf{x} - \mathbf{b})_{\Psi^c}\|_2.$$

After rearrangement we obtain
$$\|\mathbf{b} - \mathbf{x}\|_2 \leq \frac{\|(\mathbf{b} - \mathbf{x})_{\Psi^c}\|_2}{1-\delta} + \frac{\xi\|\nabla_\Psi f(\mathbf{x})\|_2}{1-\delta}, \qquad (14)$$

where this equality follows from the fact that supp($\mathbf{b}$) $\subseteq S$. Let $\Phi = $ supp($\mathbf{x}$) $\in \mathbb{M}(\mathbb{G}, k)$.

$$\|(\mathbf{x}^i - \eta\nabla_\Omega f(\mathbf{x}^i))_\Phi\|_2 \leq \|(\mathbf{x}^i - \eta\nabla_\Omega f(\mathbf{x}^i))_\Psi\|_2,$$

as $\Psi = $ supp($\mathbf{x}^i - \eta \cdot \nabla_\Omega f(\mathbf{x}^i)$). By eliminating the contribution on $\Phi \cap \Psi$, we derive

$$\|(\mathbf{x}^i - \eta\nabla_\Omega f(\mathbf{x}^i))_{\Phi\setminus\Psi}\|_2 \leq \|(\mathbf{x}^i - \eta\nabla_\Omega f(\mathbf{x}^i))_{\Psi\setminus\Phi}\|_2$$

For the right-hand side, we have
$$\|(\mathbf{x}^i - \eta\nabla_\Omega f(\mathbf{x}^i))_{\Psi\setminus\Phi}\|_2 \leq$$
$$\|(\mathbf{x}^i - \mathbf{x} - \eta\nabla_\Omega f(\mathbf{x}^i) + \eta\nabla_\Omega f(\mathbf{x}))_{\Psi\setminus\Phi}\|_2 + \eta\|\nabla_{\Omega\cup\Psi} f(\mathbf{x})\|_2,$$

where the inequality falls from the fact that $\Phi = $ supp($\mathbf{x}$). From the left-hand side, we have
$$\|(\mathbf{x}^i - \eta\nabla_\Omega f(\mathbf{x}^i))_{\Phi\setminus\Psi}\|_2 \leq -\eta\|\nabla_{\Omega\cup\Phi} f(\mathbf{x})\|_2 +$$
$$\|(\mathbf{x}^i - \mathbf{x} - \eta\nabla_\Omega f(\mathbf{x}^i) + \eta\nabla_\Omega f(\mathbf{x}))_{\Phi\setminus\Psi} + (\mathbf{x} - \mathbf{b})_{\Psi^c}\|_2$$

where the inequality follows from the fact that $\mathbf{b}_{\Psi^c} = \mathbf{0}$, $\mathbf{x}_{\Phi\setminus\Psi} = \mathbf{x}_{\Psi^c}$, and $-\mathbf{x}_{\Phi\setminus\Psi} + (\mathbf{x} - \mathbf{b})_{\Psi^c} = \mathbf{0}$. Let $\Phi\Delta\Psi$ be the symmetric difference of the set Φ and Ψ. It follows that

$$\|(\mathbf{b} - \mathbf{x})_{\Psi^c}\|_2$$
$$\leq \sqrt{2}\|(\mathbf{x}^i - \mathbf{x} - \eta\nabla_\Omega f(\mathbf{x}^i) + \eta\nabla_\Omega f(\mathbf{x}))_{\Phi\Delta\Psi}\|_2 + 2\eta\|\nabla_I f(\mathbf{x})\|_2$$
$$\leq \sqrt{2}\|(\mathbf{x}^i - \mathbf{x} - \xi\nabla_\Omega f(\mathbf{x}^i) + \xi\nabla_\Omega f(\mathbf{x}))_{\Phi\Delta\Psi}\|_2 +$$
$$\sqrt{2}(\xi - \eta)\|(\nabla_\Omega f(\mathbf{x}^i) + \nabla_\Omega f(\mathbf{x}))_{\Phi\Delta\Psi}\| + 2\eta\|\nabla_I f(\mathbf{x})\|_2$$
$$\leq \sqrt{2}\|(\mathbf{r}^i_{\Omega^c} + \mathbf{r}^i_\Omega - \xi\nabla_\Omega f(\mathbf{x}^i) + \xi\nabla_\Omega f(\mathbf{x}))_{\Phi\Delta\Psi}\|_2 +$$
$$\sqrt{2}(\xi - \eta)\|(\nabla_\Omega f(\mathbf{x}^i) - \nabla_\Omega f(\mathbf{x}))_{\Phi\Delta\Psi}\|_2 + 2\eta\|\nabla_I f(\mathbf{x})\|_2$$
$$\leq \sqrt{2}\|\mathbf{r}^i_{\Omega^c}\| + \sqrt{2}\|(\mathbf{r}^i_\Omega - \xi\nabla_\Omega f(\mathbf{x}^i) + \xi\nabla_\Omega f(\mathbf{x}))_{\Psi\Delta\Psi}\| +$$
$$\sqrt{2}(\xi - \eta)\|(\nabla_\Omega f(\mathbf{x}^i) - \nabla_\Omega f(\mathbf{x}))_{\Psi\Delta\Psi}\| + 2\eta\|\nabla_I f(\mathbf{x})\|_2$$
$$\leq \sqrt{2}\|\mathbf{r}^i_{\Omega^c}\| + \sqrt{2}\|\mathbf{r}^i - \xi\nabla_{\Omega\cup\Psi\cup\Phi} f(\mathbf{x}^i) + \xi\nabla_{\Omega\cup\Psi\cup\Phi} f(\mathbf{x})\|_2 +$$
$$\sqrt{2}(\xi - \eta)\|(\nabla_{\Omega\cup\Psi\cup\Phi} f(\mathbf{x}^i) - \nabla_{\Omega\cup\Psi\cup\Phi} f(\mathbf{x}))_{\Psi\Delta\Psi}\| + 2\eta\|\nabla_I f(\mathbf{x})\|_2$$
$$\leq \sqrt{2}\|r^i_{\Omega^c}\|_2 + \sqrt{2}\left(\left(2 - \frac{\eta}{\xi}\right)\delta + 1 - \frac{\eta}{\xi}\right)\|r^i\| +$$
$$2\left(\sqrt{2}\xi + (1 - \sqrt{2})\eta\right)\|\nabla_I f(x)\|_2,$$

where the first inequality follows from the fact that

$$\eta\|\nabla_{\Omega\cup\Phi} f(\mathbf{x})\|_2 + \eta\|\nabla_{\Psi\cup\Phi\cup\Omega} f(\mathbf{x})\|_2 \leq 2\eta\|\nabla_I f(\mathbf{x})\|_2,$$

the third inequality follows as $\mathbf{x}^i - \mathbf{x} = \mathbf{r}^i = \mathbf{r}^i_{\Omega^c} + \mathbf{r}^i_\Omega$, the fourth inequality follows from the fact that $\|(\mathbf{r}^i_{\Omega^c})_{\Phi\Delta\Psi}\|_2 \leq \|\mathbf{r}^i_{\Omega^c}\|_2$, the fifth inequality follows as $\mathbf{r}^i \subseteq \Omega \cup \Psi \cup \Phi$, and the last inequality follows from condition $(\xi, \bar{\delta}, \mathbb{M}(\mathbb{G}, 8k))$-WRSC and Lemma III.2. From Lemma III.3, we have
$$\|\mathbf{r}^i_{\Omega^c}\|_2 \leq \sqrt{1 - \eta^2}\|\mathbf{r}^i\|_2 + \left[\frac{\xi(1 + c_H)}{\eta} + \frac{\xi\eta(1 + c_H)}{\sqrt{1 - \eta^2}}\right]\|\nabla_I f(\mathbf{x})\|_2$$

Combining (14) and above inequalities, we prove the theorem. ∎

Theorem III.4 shows the estimator error of GRAPH-GHTP is determined by the multiple of $\|\nabla_S f(\mathbf{x})\|_2$, and the convergence rate is geometric. Specifically, if $\mathbf{x}$ is an uncontrained minimizer of $f(\mathbf{x})$, then $\nabla f(\mathbf{x}) = 0$. It means GRAPH-GHTP is guaranteed to obtain the true $\mathbf{x}$ to arbitrary precision. The estimation error is negligible when $\mathbf{x}$ is sufficiently close to an unconstrained minimizer of $f(\mathbf{x})$ as $\|\nabla_S f(\mathbf{x})\|_2$ is a small value. The parameter

$$\alpha = \frac{\sqrt{2}(1 + c_T)}{1 - \delta}\left(\sqrt{1 - \alpha_0^2} + \left((2 - \frac{\eta}{\xi})\delta + 1 - \frac{\eta}{\xi}\right)\right) < 1,$$

controls the convergence rate of GRAPH-GHTP. Our algorithm allows an exact recovery if $\alpha < 1$. As δ is an arbitrary

constant parameter, it can be an arbitrary small positive value. Let η be ξ and δ be an arbitrary small positive value, the parameters c_H and c_T satisfy the following inequality

$$c_H^2 > 1 - 1/(1 + c_T)^2. \tag{15}$$

It is noted that the head and tail approximation algorithms described in [15] do not meet the inequality (15). Nonetheless, the approximation factor c_H of any given head approximation algorithm can be boosted to any arbitrary constant $c'_H < 1$, which leads to the satisfaction of the above condition as shown in [15]. Boosting the head-approximation algorithm, though strongly suggested by [13], is not empirically necessary.

Our proposed GRAPH-GHTP has strong connections to the recently proposed algorithm named as Gradient Hard Thresholding Pursuit (GHTP) [40] that is designed specifically for the k-sparsity model: $\mathbb{M} = \{S \subseteq [n] \mid |S| \leq k\}$. In particular, if we redefine H($\mathbf{b}$) $= $ supp($\mathbf{b}$) and T($\mathbf{b}$) $= $ supp(P($\mathbf{b}$)), where P($\mathbf{b}$) is the projection oracle defined in Equation (5), and assume that there is an algorithm that solves the projection oracle exactly, in which the sparsity model does not require to be the k-sparsity model. It then follows that the upper bound of $\|\mathbf{r}^i_{\Omega^c}\|_2$ stated in Lemma III.2 in Appendix is updated as $\|\mathbf{r}^i_{\Omega^c}\|_2 \leq 0$, since supp($\mathbf{r}^i$) $= \Omega$ and $\mathbf{r}^i_{\Omega^c} = \mathbf{0}$. In addition, the multiplier $(1 + c_T)$ is replaced as 1 as the first inequality (13) in the proof of Theorem III.4 in Appendix is updated as $\|\mathbf{r}^{i+1}\|_2 \leq \|\mathbf{x} - \mathbf{b}\|_2$, instead of the original version $\|\mathbf{r}^{i+1}\|_2 \leq (1 + c_T)\|\mathbf{x} - \mathbf{b}\|_2$. After these two changes, the shrinkage rate α is updated as

$$\alpha = \frac{\sqrt{2}}{1 - \delta}\left((2 - \frac{\eta}{\xi})\delta + 1 - \frac{\eta}{\xi}\right), \tag{16}$$

which is **the same as the shrinkage rate of GRAPH-GHTP** as derived in [40] specifically for the k-sparsity model. The above shrinkage rate α (16) should satisfy the condition $\alpha < 1$ to ensure the geometric convergence of GRAPH-GHTP, which implies that

$$\eta > ((2\sqrt{2} + 1)\delta + \sqrt{2} - 1)\xi/(\sqrt{2} + \sqrt{2}\delta). \tag{17}$$

It follows that if $\delta < 1/(\sqrt{2} + 1)$, a step-size $\eta < \xi$ can always be found to satisfy the above inequality. This constant condition of δ is analogous to the constant condition of state-of-the-art compressive sensing methods that consider noisy measurements [23] under the assumption of the RIP condition. We derive the analogous constant using the WRSC condition that weaker than the RIP condition.

As discussed above, our proposed GRAPH-GHTP has connections to GHTP on the shrinkage rate of geometric convergence. We note that the shrinkage rate of our proposed GRAPH-GHTP stated in Theorem III.4 is derived based on head and tail approximations of the sparsity model of connected subgraphs $\mathbb{M}(\mathbb{G}, k)$, instead of the k-sparsity model that has an exact projection oracle solver. Our convergence properties hold in fairly general setups beyond k-sparsity model, as a number of popular structured sparsity models such as the "standard" k-sparsity, block sparsity, cluster sparsity, and tree sparsity can be encoded as special cases of $\mathbb{M}(\mathbb{G}, k)$.

Theorem III.5. *Let* $\mathbf{x} \in \mathbb{R}^n$ *such that* $supp(\mathbf{x}) \in \mathbb{M}(\mathbb{G}, k)$, *and* $f : \mathbb{R}^n \to \mathbb{R}$ *be cost function that satisfies condition* $(\xi, \delta, \mathbb{M}(8k, g))$-*WRSC. Assuming that* $\alpha < 1$, GRAPH-GHTP *(or* GRAPH-IHT*) returns a* $\hat{\mathbf{x}}$ *such that,* $supp(\hat{\mathbf{x}}) \in \mathbb{M}(5k, g)$ *and* $\|\mathbf{x} - \hat{\mathbf{x}}\|_2 \leq c\|\nabla_I f(\mathbf{x})\|_2$, *where* $c = (1 + \frac{\beta}{1-\alpha})$ *is a fixed constant. Moreover,* GRAPH-GHTP *runs in time*

$$O\left((T + |\mathbb{E}| \log^3 n) \log(\|\mathbf{x}\|_2 / \|\nabla_I f(\mathbf{x})\|_2)\right), \quad (18)$$

where T *is the time complexity of one execution of the subproblem in Step 6 in* GRAPH-GHTP *(or Step 5 in* GRAPH-IHT*). In particular, if* T *scales linearly with* n, *then* GRAPH-GHTP *(or* GRAPH-IHT*) scales nearly linearly with* n.

Proof: The i-th iterate of GRAPH-GHTP (or GRAPH-IHT) satisfies

$$\|\mathbf{x} - \mathbf{x}^i\|_2 \leq \alpha^i \|\mathbf{x}\|_2 + \frac{\beta}{1-\alpha} \|\nabla_I f(\mathbf{x})\|_2. \quad (19)$$

After $t = \left\lceil \log\left(\frac{\|\mathbf{x}\|_2}{\|\nabla_I f(\mathbf{x})\|_2}\right) / \log \frac{1}{\alpha} \right\rceil$ iterations, GRAPH-GHTP (or GRAPH-IHT) returns an estimate $\hat{x}$ satisfying $\|\mathbf{x} - \hat{x}\|_2 \leq (1 + \frac{\beta}{1-\alpha})\|\nabla_I f(\mathbf{x})\|_2$. The time complexities of both head approximation and tail approximation are $O(|\mathbb{E}| \log^3 n)$. The time complexity of one iteration in GRAPH-GHTP (or GRAPH-IHT) is $(T + |\mathbb{E}| \log^3 n)$, and the total number of iterations is $\left\lceil \log\left(\frac{\|\mathbf{x}\|_2}{\|\nabla_I f(\mathbf{x})\|_2}\right) / \log \frac{1}{\alpha} \right\rceil$, and the overall time follows. GRAPH-GHTP and GRAPH-IHT are only different in the definition of α and β in this Theorem. ∎

As shown in Theorem 18, the time complexity of GRAPH-GHTP is dependent on the total number of iterations and the time cost (T) to solve the subproblem in Step 6. In comparison, the time complexity of GRAPH-IHT is dependent on the total number of iterations and the time cost T to calculate the gradient $\nabla f(\mathbf{x}^i)$ in Step 5. It implies that, although GRAPH-GHTP converges faster than GRAPH-IHT, the time cost to solve the subproblem in Step 6 is often much higher than the time cost to calculate a gradient $\nabla f(\mathbf{x}^i)$, and hence GRAPH-IHT runs faster than GRAPH-GHTP in practice.

IV. APPLICATIONS ON GRAPH SCAN STATISTICS

In this section, we specialize GRAPH-IHT and GRAPH-GHTP to optimize a number of well-known graph scan statistics for the task of connected subgraph detection, including elevated mean scan (EMS) statistic [29], Kulldorff's scan statistic [25], and expectation-based Poisson (EBP) scan statistic [26]. Each graph scan statistic is defined as the generalized likelihood ratio test (GLRT) statistic of a *specific* hypothesis testing about the distributions of features of normal and abnormal nodes. The EMS statistic corresponds to the following GLRT test: Given a graph $\mathbb{G} = (\mathbb{V}, \mathbb{E})$, where $\mathbb{V} = [n]$ and $\mathbb{E} \subseteq \mathbb{V} \times \mathbb{V}$, each node i is associated with a random variable x_i:

$$x_i = \mu \cdot 1(i \in S) + \epsilon_i, \ i \in \mathbb{V}, \quad (20)$$

where $|\mu|$ represents the signal strength and $\epsilon_i \in \mathcal{N}(0, 1)$. S is some unknown anomalous cluster that forms as a connected subgraph. The task is to decide between the null hypothesis (H_0): $c_i \in \mathcal{N}(0, 1), \forall i \in \mathbb{V}$ and the alternative $(H_1(S))$: $c_i \in \mathcal{N}(\mu, 1), \forall i \in S$ and $c_i \in \mathcal{N}(0, 1), \forall i \notin S$. The EMS statistic

is defined as the GLRT function under this hypothesis testing:

$$F(S) = \frac{\text{Prob}(\text{Data}|H_1(S))}{\text{Prob}(\text{Data}|H_0)} = \frac{1}{\sqrt{|S|}} \sum_{i \in S} c_i. \quad (21)$$

The problem of connected subgraph detection based on the EMS statistic is then formulated as

$$\min_{S \subseteq \mathbb{V}} -\frac{1}{|S|} \left(\sum_{i \in S} c_i\right)^2 \ s.t. \ S \in \mathbb{M}(\mathbb{G}, k), \quad (22)$$

where the **square** of the EMS scan statistic is considered to make the function smooth, and this transformation does not infect the optimum solution. Let the $\{0, 1\}$-vectors form of S be $\mathbf{x} \in \{0, 1\}^n$, such that $supp(\mathbf{x}) = S$. Problem (22) can be reformulated as

$$\min_{x \in \{0,1\}^n} -(\mathbf{c}^\mathsf{T} \mathbf{x})^2 / (\mathbf{1}^\mathsf{T} \mathbf{x}) \ s.t. \ supp(\mathbf{x}) \in \mathbb{M}(\mathbb{G}, k), \quad (23)$$

where $\mathbf{c} = [c_1, \cdots, c_n]^\mathsf{T}$. To apply our proposed algorithms, we relax the input domain of $\mathbf{x}$ and maximize the strongly convex function [3]:

$$\min_{x \in \mathbb{R}^n} -(\mathbf{c}^\mathsf{T} \mathbf{x})^2 / (\mathbf{1}^\mathsf{T} \mathbf{x}) + \frac{1}{2} \mathbf{x}^\mathsf{T} \mathbf{x} \ s.t. \ supp(\mathbf{x}) \in \mathbb{M}(\mathbb{G}, k). \quad (24)$$

The connected subset of nodes can be found as the subset of indexes of positive entries in $\hat{\mathbf{x}}$, where $\hat{\mathbf{x}}$ refers to the solution of the Problem (24). Assume that $\mathbf{c}$ is normalized and $c_i \leq 1$, $\forall i$. Let $\hat{c} = \max\{c_1, \cdots, c_n\}$. The Hessian matrix of the above objective function satisfies the following conditions

$$(1 - \hat{c}^2) \cdot \mathbf{I} \preceq \mathbf{I} - (\mathbf{c} - \frac{\mathbf{c}^\mathsf{T} \mathbf{x}}{\mathbf{1}^\mathsf{T} \mathbf{x}} \mathbf{1})(\mathbf{c} - \frac{\mathbf{c}^\mathsf{T} \mathbf{x}}{\mathbf{1}^\mathsf{T} \mathbf{x}} \mathbf{1})^\mathsf{T} \preceq 1 \cdot \mathbf{I}. \quad (25)$$

According to Lemma 1 (b) in [40]), the objective function $f(\mathbf{x})$ satisfies condition $(\xi, \delta, \mathbb{M}(\mathbb{G}, 8k))$-WRSC that

$$\delta = \sqrt{1 - 2\xi(1 - \hat{c}^2) + \xi^2},$$

for any ξ such that $\xi < 2(1 - \hat{c}^2)$. The geometric convergence of GRAPH-GTHP as shown in Theorem III.4 is guaranteed.

Different from the EMS statistic that is defined for numerical features based on Gaussian distribution, the Kulldorff's scan statistic and Expectation Based Poisson statistic (EBP) are defined for count features based on Poisson distribution. In particular, each node i is associated with a feature c_i, the count of events (e.g., crimes, flu infections) observed at the current time, and a feature b_i, the expected count (or 'baseline') of events by using historical data. Let $\mathbf{c} = [c_1, \cdots, c_n]^\mathsf{T}$ and $\mathbf{b} = [b_1, \cdots, b_n]^\mathsf{T}$. The Kulldorff's scan statistic and EBP scan statistics are described Table I. We note that these two scan statistics do not satisfy the WRSC condition, but as demonstrated in our experiments, our proposed algorithms perform empirically well for all the three scan statistics, and in particular, our proposed GRAPH-GHTP converged in less than 10 iterations in all the settings.

V. EXPERIMENTS

This section evaluates the performance of our proposed methods using four public benchmark data sets for connected subgraph detection. The **experimental code and data sets** are available from the Link [11] for reproducibility.

Score Functions	Definition	Applications
Kulldorff's Scan Statistic [25]	$\mathbf{c}^\mathsf{T}\mathbf{x}\log\frac{\mathbf{c}^\mathsf{T}\mathbf{x}}{\mathbf{b}^\mathsf{T}\mathbf{x}} - \mathbf{1}^\mathsf{T}\mathbf{c}\log\frac{\mathbf{1}^\mathsf{T}\mathbf{c}}{\mathbf{1}^\mathsf{T}\mathbf{b}} + (\mathbf{1}^\mathsf{T}\mathbf{c} - \mathbf{c}^\mathsf{T}\mathbf{x})\log\frac{\mathbf{1}^\mathsf{T}\mathbf{c}-\mathbf{c}^\mathsf{T}\mathbf{x}}{\mathbf{1}^\mathsf{T}\mathbf{b}-\mathbf{b}^\mathsf{T}\mathbf{x}}$	The statistic is used for anomalous pattern detection in graphs with count features, such as detection of traffic bottlenecks in sensor networks [1], [22], detection of anomalous regions in digitals and images [10], detection of attacks in computer networks [24], disease outbreak detection [36], and various others.
Expectation-based Poisson Statistic (EBP) [26]	$\mathbf{c}^\mathsf{T}\mathbf{x}\log\frac{\mathbf{c}^\mathsf{T}\mathbf{x}}{\mathbf{b}^\mathsf{T}\mathbf{x}} + \mathbf{b}^\mathsf{T}\mathbf{x} - \mathbf{c}^\mathsf{T}\mathbf{x}$	This statistic is used for the same applications as above [1], [12], [25], [36], [37], but has different assumptions on data distribution [25].
Elevated Mean Scan Statsitic [29]	$\mathbf{c}^\mathsf{T}\mathbf{x}/\mathbf{1}^\mathsf{T}\mathbf{x}$	This statistic is used for anomalous pattern detection in graphs with numerical features, such as event detection in social networks, network surveillance, disease outbreak detection, biomedical imaging [29], [35]

Dataset	Application	Training & Testing Time Periods	Observed value at node v: c_v	Baseline value at node v: b_v	# of Nodes	# of Edges	# of snapshots
BWSN	Detection of contaminated nodes	Training: Hours 3 to 5 with 0% noise Testing: Hours 3 to 5 with 2%, 4%, $\cdots$, 10% noises	Sensor value (0 or 1)	Average sensor value for EBP Constant '1' for Kulldorff	12,527	14,323	hourly: 3×6
CitHepPh	Detection of emerging research areas	Testing: 1999 to 2002	Count of citations	Average count of citations for EBP Maximum count of citations for Kulldorff	11,895	75,873	yearly: 1×11
Traffic	Detection of most congested subgraphs	Testing: Mar.2014, 5AM to 10PM	$-\log(p^t(v)/\mu)$ [1]	None (EBP and Kulldorff are not applicable)	1,723	5,301	per 15 min: 68×304
ChicagoCrime	Detection of crime hot spots	Testing: Year of 2015	Count of burglaries	Average count of burglaries for EBP Maximum count of burglaries for Kulldorff	46,357	168,020	yearly: 1×15

[1] $p^t(v)$ refers to the statistical p-value of node v at time t that is calculated via empirical calibration based on historical speed values of v from 2013 June. 1st to 2014 Feb. 29; and μ is a significance level threshold and is set to 0.15. The larger this value $-\log(p^t(v)/\mu)$, the more congested in the region near v.

A. Experiment Design

Datasets: 1) **BWSN Dataset.** A real-world water network is offered in the Battle of the Water Sensor Networks (BWSN) [28]. That has 12,527 nodes and 14,323 edges. In order to simulate a contaminant sub-area, 4 nodes with chemical contaminant plumes, which were distributed in this sub-area, were generated. We use the water network simulator EPANET [31] that was employed in BWSN for a period of 3 hours to simulate the spreads for contaminant plumes on this graph. If a node is polluted by the chemical, then its sensor reports 1, otherwise, 0, in each hour. To test the tolerance of noise of our methods, $K \in \{2, 4, 6, 8, 10\}$ percent vertices were selected randomly, and their sensor reports were set to 0 if their original reports were 1 and vice versa. Each hour has a graph snapshot. The snapshots corresponding to the 3 hours that have 0% noise are considered for training, and the snapshots that have $2\%, \cdots, 10\%$ noise reports for testing. The goal is to detect a connected subgraph that is corresponding to the contaminant sub-area. 2) **CitHepPh Dataset.** We downloaded the high energy physics phenomenology citation data (CitHepPh) from Stanford Network Analysis Project (SNAP) [20]. This citation graph contains 11,897 papers corresponding to graph vertices and 75,873 edges. An undirected edge between two vertices (papers) exists , if one paper is cited by another. The period of these papers published is from January 1992 to April 2002. Each vertex has two attributes for each specific year ($t = 1992, \cdots, t = 2002$). We denote the number of citations of each specific year as the first attribute and the average citations of all papers in that year as the second attribute. The goal is to detect a connected subgraph where the number of citations of vertices (papers) in this subgraph are abnormally high compared with the citations of vertices that are not in this subgraph. This connected subgraph is considered as a potential emerging research area. Since the training data

is required for some baseline methods, the data before 1999 is considered as the training data, and the rest from 1999 to 2002 as the testing data. 3) **Traffic Dataset**. Road traffic speed data from June 1, 2013 to Mar. 31, 2014 in the arterial road network of the Washington D.C. region is collected from the INRIX database (http://inrix.com/publicsector.asp), with 1,723 nodes and 5,301 edges. The database provides traffic speed for each link at a 15-minute rate. For each 15-minute interval, each method identies a connected subgraph as the most congested region. 4) **ChicagoCrime Dataset.** We collected crime data from City of Chicago [https://data.cityofchicago.org/Public-Safety/Crimes-2001-to-present/ijzp-q8t2] from Jan. 2001 to Dec. 2015. There are 46,357 nodes (census blocks) and 168,020 edges (Two census blocks are connected with each other if they are neighbours). Specifically, we collected all records of burglaries from 2001 to 2015. The data covers burglaries in the period from Jan. 2001 to Dec. 2015. Each vertex has an attribute denoting the number of burglaries in sepcific year and average number of burglaries over 10 years. We aim to detect connected census areas which has anomaly high burglaries accidents. The data before 2010 is considered as training data, and the data from 2011 to 2015 is considered as testing data.

Graph Scan Statistics: As shown in Table I, three graph scan statistics were considered as the scoring functions of connected subgraphs, including Kulldorff's scan statistic [25], expectation-based Poisson (EBP) scan statistic [26], and elevated mean scan (EMS) statistic [29]. The first two statistic functions require that each vertex v has a count c_v representing the count of events observed at that vertex, and an expected count ('baseline') b_v. For EMS statistic, only c_v is used. We need to normalize c_v for EMS as it is defined based on the assumptions of standard normal distribution for normal values and shifted-mean normal distribution for abnormal values.

Table II provides details about the calculations of c_v and b_v for each data set.

Comparison Methods: We compared our proposed methods with four state-of-the-art baseline methods that are designed specifically for connected subgraph detection, namely, `GraphLaplacian` [33], `EventTree` [32], `DepthFirstGraphScan` [36] and `NPHGS` [8]. `DepthFirstGraphScan` is an exact search algorithm based on depth-first search and takes weeks to run on graphs that have more than 1000 nodes. We imposed a maximum limit on the depth of the search to 10 to reduce its time complexity.

The basic ideas of these baseline methods are summarized as follows: `NPHGS` starts from random seeds (nodes) as initial candidate clusters and gradually expends each candidate cluster by including its neighboring nodes that could help improve its BJ statistic score until no new nodes can be added. The candidate cluster with the largest BJ statistic score is returned. `DepthFirstGraphScan` adopts a similar strategy to `NPHGS` but expands the initial clusters based on depth-first search. `GraphLaplacian` uses a graph Laplacian penalty function to replace the connectivity constraint and converts the problem to a convex optimization problem. `EventTree` reformulates the connected subgraph detection problem as a prize-collecting steiner tree (PCST) problem [19] and apply the Goemans-Williamson (G-W) algorithm for PCST [19] to detect anomalous subgraphs. We also implemented the generalized fused lasso model (`GenFusedLasso`) for graph scan statistics using the framework of alternating direction method of multipliers (ADMM). `GenFusedLasso` method solves the following minimization problem

$$\min_{\mathbf{x}\in\mathbb{R}^n} -f(\mathbf{x}) + \lambda \sum_{(i,j)\in\mathbb{E}} |x_i - x_j|, \quad (26)$$

where $f(\mathbf{x})$ is a predefined graph scan statistic and the trade-off parameter λ controls the degree of smoothness of neighboring entries in $\mathbf{x}$. We applied the heuristic rounding step proposed in [29] to $\mathbf{x}$ to generate connected subgraphs.

Parameter Tunning: We strictly followed strategies recommended by authors in their original papers to tune the related model parameters. Specifically, for `EventTree`, we tested the set of λ values: $\{0.02, 0.04, \cdots, 2.0, 3.0, \cdots, 20\}$. For `Graph-Laplacian`, we tested the set of λ values: $\{0.001, 0.003, 0.01, 0.03, 0.1, 0.3, 1\}$ and returned the best result. For `GenFusedLasso`, we tested the set of λ values: $\{0.02, 0.04, \cdots, 2.0, 3.0, \cdots, 20\}$. For `NPHGS`, we set the suggested parameters by the authors: $\alpha_{max} = 0.15$ and $K = 5$. Our proposed methods GRAPH-IHT and GRAPH-GHTP have a single parameter k, an upper bound of the subgraph size. We tested the set of k values: $\{50, 100, \cdots, 1000\}$. As the BWSN dataset has the ground truth about the contaminated nodes, we identified the best parameter for each method that has the largest F-measure in the training data. For the other data sets, as we do not have ground truth about the true subgraphs, for each specific scan statistic, we identified the best parameter for each method that was able to identify the connected subgraph with the largest statistic score.

Performance Metrics: 1) Optimization Power. The over-all scores of the three graph scan statistic functions of the connected subgraphs returned by the comparison methods are compared and analyzed. The objective is to identify methods that could find the connected subgraphs with the largest graph scan statistic scores. **2) Precision, Recall, and F-Measure.** For the BWSN dataset, as the true anomalous subgraphs are known, we use F-measure that combines precision and recall to evaluate the quality of detected subgraphs by different methods. **3) Run Time.** The running times of different methods are compared.

B. Evolving Curves of Graph Scan Statistics

Figure 3 compares our methods `Graph-IHT` and GRAPH-GHTP with `GenFusedLasso` on the scores of two graph scan statistics (Kulldorff's scan statistic and elevated mean scan statistic (EMS)) based on the best connected subgraphs identified by both methods in different iterations. Note that, a heuristic rounding process as proposed in [29] is applied to continuous vector $\mathbf{x}^i$ estimated by `GenFusedLasso` at each iteration i in order to identify the best connected subgraph at the current iteration. As the setting of the parameter λ will influence the quality of the detected connected subgraph, the results under different λ values are also shown in Figure 2. The results indicate that our method GRAPH-GHTP converges in less than 10 iterations, and GRAPH-IHT converges in more steps. The qualities (scan statistic scores) of the connected subgraphs identified at different iterations by our two methods are consistently higher than those returned by `GenFusedLasso`.

C. Optimization Power

The comparisons between our method and the other baseline methods are shown in Table III and Table IV. The scores of the three graph scan statistics based on the connected subgraphs returned by these methods are reported in these two tables. The results in indicate that our method outperformed all the baseline methods on the three graph scan statistics, except that `EventTree` achieved the highest Kulldorff score (16738.43) on the CitHepPh dataset, but is only 2.71% larger than the returned score of our method GRAPH-GHTP. We note `EventTree` is a heuristic algorithm and does not provide theoretical guarantee on the quality of the connected subgraph returned, as measured by the scan statistic scores.

D. Water Pollution Detection

Figure 4 shows the precision, recall, and F-measure of all the comparison methods on the detection of polluted nodes in the water distribution network in BWSN with respect to different noise ratios. The results indicate that our proposed method GRAPH-GHTP and `DepthFirstGraphScan` were the best methods on all the three measures for most of the settings. However, `DepthFirstGraphScan` spent 5929 seconds to finish, and GRAPH-GHTP spent only 166 seconds, 35.8 times faster than `DepthFirstGraphScan`. `EventTree` achieved high recalls but low precisions consistently in different settings. In contrast, `GraphLaplacian` and `NPHGS` achieved high precisions but low recalls in most settings.
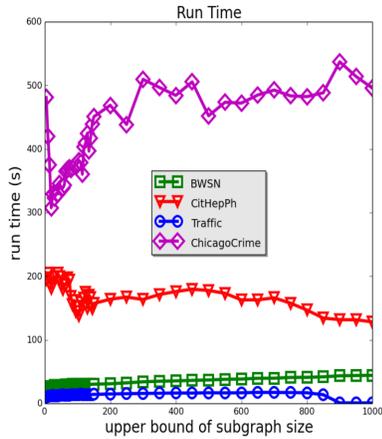
Fig. 2. Scability of GRAPH-GHTP with respect to k (the upper bound of subgraph size).
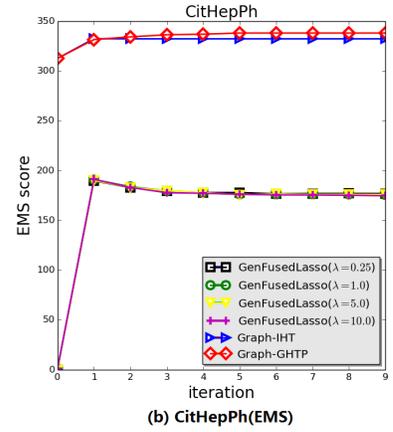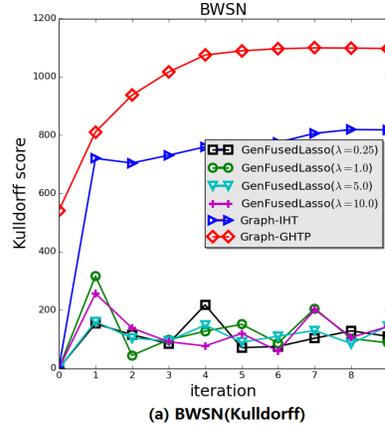
Fig. 3. Evolving curves of graph scan statistic scores between our methods(GRAPH-IHT and GRAPH-GHTP) and `GenFusedLasso` in different iterations.

(a) BWSN(precision)   (b) BWSN(recall)   (c) BWSN(fmeasure)
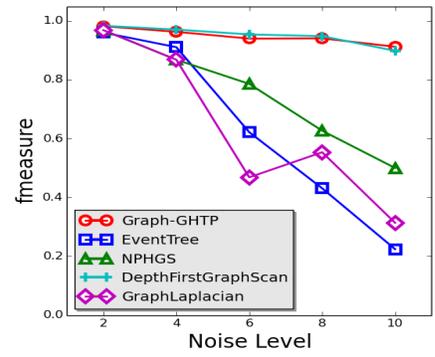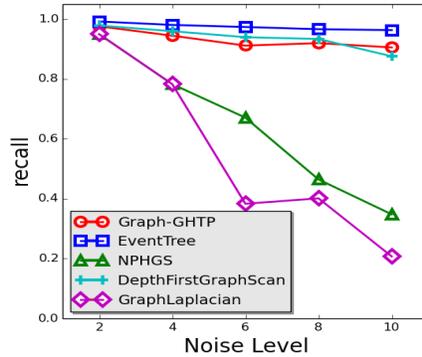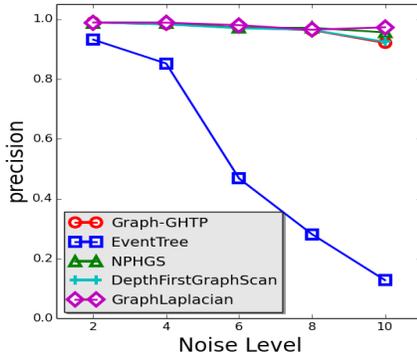
Fig. 4. Precision, Recall, and F-measure curves for water pollution detection in BWSN with respect to different noise ratios.

TABLE III

Comparison on scores of the three graph scan statistics based on connected subgraphs returned by comparison methods. EMS and EBP are Elevated Mean Scan Statistic and Expectation-Based Poisson Statistic, respectively.

	BWSN				CitHepPh			
	Kulldorff	EMS	EBP	Run Time	Kulldorff	EMS	EBP	Run Time
GRAPH-GHTP	**1097.15**	**21.56**	**79.71**	**165.86**	16296.40	**337.90**	9342.94	155.74
GraphLaplacian	474.96	14.89	49.91	55315.94	2585.44	202.38	2305.05	22424.24
EventTree	834.59	20.25	32.13	441.74	**16738.43**	335.34	9061.56	**124.28**
DepthFirstGraphScan	735.85	20.41	79.30	5929.00	9531.19	260.06	5561.66	12183.88
NPHGS	541.13	16.90	58.59	256.91	11965.14	326.23	9098.22	175.08

E. Scalability Analysis

Table III and Table IV also show the comparison between our proposed method GRAPH-GHTP and other baseline methods on the running time. The results indicate that our proposed method GRAPH-GHTP ran faster than all the baseline methods in most of the settings, except for `EventTree`. `EventTree` was the fastest method but was unable to detect subgraphs with high qualities. As our method has a parameter on the upper bound (k) of the subgraph returned, we also conducted the scalability of our method with respect to different values of k as shown in Figure 2. The results indicate that the running time of our algorithm is insensitive to the setting of k, which is consistent with the time complexity analysis of GRAPH-GHTP as discussed in Theorem III.5.

VI. RELATED WORK

A. Structured sparse optimization. The methods in this category have been briefly reviewed in the introduction section. The most relevant work is by Hegde et al. [15]. The authors present GRAPH-COSAMP, a variant of COSAMP [23], for compressive sensing and linear regression problems based on head and tail approximations of $\mathbb{M}(\mathbb{G}, k)$.

B. Connected subgraph detection. Existing methods in this category fall into three major categories:1) Exact algorithms. The most recent method is a brunch-and-bounding algorithm `DepthFirstGraphScan` [36] that runs in exponential time in the worst case; 2) Heuristic algorithms. The most recent methods in this category include `EventTree` [32], `NPHGS` [8], `AdditiveScan` [37],

TABLE IV

Comparison on scores of the three graph scan statistics based on connected subgraphs returned by comparison methods. EMS and EBP are Elevated Mean Scan Statistic and Expectation-Based Poisson Statistic, respectively. `GraphLaplacian` failed to run on ChicagoCrime due to out-of-memory error.

	Traffic		ChicagoCrime			
	EMS	Run Time	Kulldorff	EMS	EBP	Run Time
GRAPH-GHTP	**20.45**	22.25	**6386.08**	**5.45**	**5172.54**	3177.60
GraphLaplacian	5.40	291.75	-	-	-	-
EventTree	12.40	5.02	4388.42	4.91	3965.96	**226.50**
DepthFirstGraphScan	8.13	47.73	1123.49	2.56	1094.21	12133.50
NPHGS	6.28	**0.22**	966.70	2.43	948.23	701.40

`GraphLaplacian` [33], and `EdgeLasso` [34]; 3) Approximation algorithms that provide performance bounds. The most recent method is presented by Qian et al.. The authors reformulate the connectivity constraint as linear matrix inequalities (LMI) and present a semi-definite programming algorithm based on convex relaxation of the LMI [18, 19] with a performance bound. However, this method is not scalable to large graphs ($\geq$ 1000 nodes). Most of the above methods are considered as baseline methods in our experiments and are briefly summarized in Section V-A.

VII. CONCLUSION

This paper presents, GRAPH-IHT and GRAPH-GHTP, two efficient algorithms to optimize a general nonlinear optimization problem subject to connectivity constraint on the support of variables. Extensive experiments demonstrate the effectiveness and efficency of our algorithms. For the future work, we plan to explore graph-structured constraints other than connectivity constraint and extend our proposed methods such that good theoretical properties of the cost functions that do not satisfy the WRSC condition can also be analyzed.

REFERENCES

[1] B. Anbaroğlu, T. Cheng, and B. Heydecker. Non-recurrent traffic congestion detection on heterogeneous urban road networks. *TRANSPORTMETRICA*, 11(9):754–771, 2015.

[2] M. Asteris and e. a. Kyrillidis. Stay on path: Pca along graph paths. In *ICML*, pages 1728–1736, 2015.

[3] F. Bach. Learning with submodular functions: A convex optimization perspective. *arXiv:1111.6453*, 2011.

[4] F. Bach, R. Jenatton, J. Mairal, G. Obozinski, et al. Structured sparsity through convex optimization. *Stat Sci*, 27(4):450–468, 2012.

[5] S. Bahmani, P. T. Boufounos, and B. Raj. Learning model-based sparsity via projected gradient descent. *IT*, 2016.

[6] S. Bahmani, B. Raj, and P. T. Boufounos. Greedy sparsity-constrained optimization. *JMLR*, 14(1):807–841, 2013.

[7] T. Blumensath. Compressed sensing with nonlinear observations and related nonlinear optimization problems. *IT*, 59(6):3466–3474, 2013.

[8] F. Chen and D. B. Neill. Non-parametric scan statistics for event detection and forecasting in heterogeneous social media graphs. In *ACM SIGKDD*, pages 1166–1175, 2014.

[9] F. Chen and D. B. Neill. Human rights event detection from heterogeneous social media graphs. *Big Data*, 3(1):34–40, 2015.

[10] J. W. Coulston and K. H. Riitters. Geographic analysis of forest health indicators using spatial scan statistics. *EM*, 31(6):764–773, 2003.

[11] Dataset and code. https://github.com/baojianzhou/Graph-GHTP.

[12] W. L. Gorr and Y. Lee. Early warning system for temporary crime hot spots. *Journal of Quantitative Criminology*, 31(1):25–47, 2015.

[13] C. Hegde, P. Indyk, and L. Schmidt. Approximation-tolerant model-based compressive sensing. In *SODA*, pages 1544–1561. SIAM, 2014.

[14] C. Hegde, P. Indyk, and L. Schmidt. A fast approximation algorithm for tree-sparse recovery. In *ISIT*, pages 1842–1846. IEEE, 2014.

[15] C. Hegde, P. Indyk, and L. Schmidt. A nearly-linear time framework for graph-structured sparsity. In *ICML*, pages 928–937, 2015.

[16] J. Huang, T. Zhang, and D. Metaxas. Learning with structured sparsity. *JMLR*, 12:3371–3412, 2011.

[17] L. Jacob, G. Obozinski, and J.-P. Vert. Group lasso with overlap and graph lasso. In *ICML*, pages 433–440. ACM, 2009.

[18] P. Jain, A. Tewari, and P. Kar. On iterative hard thresholding methods for high-dimensional m-estimation. In *NIPS*, pages 685–693, 2014.

[19] D. S. Johnson, M. Minkoff, and S. Phillips. The prize collecting steiner tree problem: theory and practice. In *SODA*, pages 760–769, 2000.

[20] J. e. a. Leskovec. Graphs over time: densification laws, shrinking diameters and possible explanations. In *KDD*, pages 177–187, 2005.

[21] J. Mairal and B. Yu. Supervised feature selection in graphs with path coding penalties and network flows. *JMLR*, 14(1):2449–2485, 2013.

[22] R. Modarres and G. Patil. Hotspot detection with bivariate data. *Journal of Statistical planning and inference*, 137(11):3643–3654, 2007.

[23] D. Needell and J. A. Tropp. Cosamp: Iterative signal recovery from incomplete and inaccurate samples. *ACHA*, 26(3):301–321, 2009.

[24] J. Neil and H. et al. Scan statistics for the online detection of locally anomalous subgraphs. *Technometrics*, 55(4):403–414, 2013.

[25] D. B. Neill. An empirical comparison of spatial scan statistics for outbreak detection. *IJHG*, 8(1):1, 2009.

[26] D. B. Neill. Fast subset scan for spatial pattern detection. *JRSS: Series B (Statistical Methodology)*, 74(2):337–360, 2012.

[27] D. Oliveira and D. P. et al. Detection of patterns in water distribution pipe breakage using spatial scan statistics for point events in a physical network. *JCCE*, 25(1):21–30, 2010.

[28] A. Ostfeld and J. G. e. a. Uber. The battle of the water sensor networks (bwsn): A design challenge for engineers and algorithms. *JWRPM*, 134(6):556–568, 2008.

[29] J. Qian, V. Saligrama, and Y. Chen. Connected sub-graph detection. In *AISTATS*, 2014.

[30] R. T. Rockafellar. Monotone operators and the proximal point algorithm. *SIAM journal on control and optimization*, 14(5):877–898, 1976.

[31] L. A. Rossman and W. Supply. Epanet 2 users manual, 2000.

[32] P. Rozenshtein, A. Anagnostopoulos, A. Gionis, and N. Tatti. Event detection in activity networks. In *SIGKDD*, pages 1176–1185, 2014.

[33] J. Sharpnack, A. Rinaldo, and A. Singh. Changepoint detection over graphs with the spectral scan statistic. *arXiv:1206.0773*, 2012.

[34] J. Sharpnack, A. Rinaldo, and A. Singh. Sparsistency of the edge lasso over graphs. In *AISTATS*, pages 1028–1036, 2012.

[35] J. L. Sharpnack, A. Krishnamurthy, and A. Singh. Near-optimal anomaly detection in graphs using lovász extended scan statistic. In *NIPS*, pages 1959–1967, 2013.

[36] S. Speakman, E. McFowland Iii, and D. B. Neill. Scalable detection of anomalous patterns with connectivity constraints. *JCGS*, 0(ja):00–00, 0.

[37] S. Speakman, Y. Zhang, and D. B. Neill. Dynamic pattern detection with temporal consistency and connectivity constraints. In *ICDM*, pages 697–706. IEEE, 2013.

[38] A. Tewari and R. et al. Greedy algorithms for structurally constrained high dimensional problems. In *NIPS*, pages 882–890, 2011.

[39] B. Xin and K. et al. Efficient generalized fused lasso and its application to the diagnosis of alzheimers disease. In *AAAI*, pages 2163–2169, 2014.

[40] X. Yuan, P. Li, and T. Zhang. Gradient hard thresholding pursuit for sparsity-constrained optimization. In *ICML*, pages 127–135, 2014.

[41] X.-T. Yuan and Q. Liu. Newton greedy pursuit: A quadratic approximation method for sparsity-constrained optimization. In *CVPR*, pages 4122–4129. IEEE, 2014.

[42] T. Zhang. Adaptive forward-backward greedy algorithm for sparse learning with linear models. In *NIPS*, pages 1921–1928, 2009.