Part A
Project Summary

This project will extend aspects of **spectral graph theory** to a **spectral theory for cut structures**, defined as sets of graphs with approximately prescribed cuts. The central goal behind studying spectral properties of sets of graphs rather than fixed graphs is the design of **fast algebraic algorithms** for problems on graphs. Progress will be enabled by a deeper understanding of the generalized spectra of graph pairs. Beyond the algorithmic power of the new theory, the project will also explore its **descriptive power** in explaining the performance of algorithms and other phenomena in **classes of graphs** such as social networks.

The project will disseminate implementations of the new algorithms, and will apply them to **massive data sets**, from **neural and medical imaging** applications. The research and its dissemination will be undertaken in organic integration with **educational activities** carefully designed to exploit the opportunities and meet the challenges of an expanding Computer Science department in the University of ***********, the largest Hispanic-serving institution in the United States.

Intellectual Merit. The discovery of **nearly-optimal solvers** for graph Laplacians and the more general class of symmetric diagonally dominant (SDD) linear systems has enabled the design of faster algorithms for a growing list of problems that now includes the max-flow problem. SDD solvers are thus justifiably considered a **powerful algorithmic primitive**. The near-optimality of SDD solvers extends to algorithms for the approximate computation of a few extreme eigenvectors. However, despite the arguably fundamental nature of eigenvectors, the availability of fast eigensolvers has not so far led to related algorithmic improvements. The reason is that the cut-related combinatorial information contained in the eigenvalues and eigenvectors can be significantly 'loose' for **fixed graphs**. This is not necessarily true for **cut structures**. A cut structure –given in the form of an input graph and a 'slack factor' for the cut sizes– may always contain a graph whose eigenvalues and eigenvectors provide 'tight' information which reflects directly back to the input graph, albeit within an approximation factor. This leads to a **spectral modification** approach, which suggests the transformation of the input graph A to a graph B with similar cut structure, but with the additional property that spectral methods work well for it. This line of reasoning provides the main motivation for the development of the **spectral theory for cut structures**.

It is expected that results of the proposed research will include: **(a)** A significant improvement in the **approximation guarantees in nearly-linear time** via spectral methods for **graph partitioning**. Spectral partitioning alone doesn't guarantee an $f(n)$ approximation ratio for *any* non-trivial function f of the graph size n. The proposed research will yield the first such approximation, with a polylogarithmic approximation ratio being a likely outcome. **(b)** Extensions of the spectral partitioning method to the **generalized sparsest cut problem** that seeks the minimization of the ratio of simultaneous cuts in two graphs A and B; a notable instantiation of the problem is the minimum s-t cut problem. **(c)** A **better understanding** and analysis of the performance of (variants of) known algorithms on **classes of graphs**, such as social networks. Examples include routing algorithms and the recent algorithm for the computation of approximate max flows.

Broader Impacts. The project will disseminate implementations of the new algorithms and will apply them to imaging problems in industry and academia. The project will undertake educational activities in ***********, an institution entirely comprised by Hispanic students, a demographic group with well documented under-represenation problems. New courses will be introduced. The PI will recruit more undergraduate students in small research and development projects. Crucial to this endeavor will be the continuation of the PI's syn-ergistic activities within *********** The PI will work towards the overall strengthening of Computer Science in ***********, by serving as the coordinator of undergraduate research in his department, as a member of the CS-Math graduate committee, and also by influencing the recruitment of new faculty. Activities within *********** will impact other Hispanic and Native American students through active participation in the annual SACNAS conference.

Key words: Spectral graph theory, Graph partitioning, Spectral partitioning, Linear system solvers.

Part C

Project Description

1 Results from Prior NSF Support

1.1 Short Summary

The PI and his collaborators have designed a **near-optimal solver** for symmetric diagonally dominant (SDD) linear systems. The solver is the **key subroutine** in the fastest known approximation algorithms for several problems, including the max-flow problem. In addition, they've presented the first work-efficient parallel SDD solver, a result that implies **work-efficient parallel algorithms** for several problems, including single-source shortest paths, maximum flow, min-cost flow, and approximate max-flow. These algorithms are based on new insights and algorithms for **spectral graph sparsification** and the computation of **low-stretch subgraphs**. They've also developed provably good preconditioners based on fast **multi-way partitioning** algorithms, and studied their connections with **multigrid** solvers, shedding new light into their convergence.

1.2 Broader Impacts

The **source code** of **Combinatorial Multigrid**, an SDD solver based on spectral graph theory, is **publicly available** on the PI's web page. Symmetric diagonally dominant systems are ubiquitous in engineering, including applications in **computer vision**. The PI has initiated a collaboration with Professor Eduardo Rosa-Molinar and his **biological imaging** group at the University of Puerto Rico. The collaboration has led to the development of algorithms and software based on SDD solvers for **neural EM imaging**. On instances generated from these applications our solver is by far the fastest among all publicly available solvers.

The PI works closely with two undergraduate **Hispanic students** at ***********. The students presented their preliminary findings in SIDIM and PRISM 2011, the two main conferences held annually in Puerto Rico. Newer work will be presented at a **national level** in the 2011 *Neuroscience* conference. Both students spent one month during Summer 2011 at Carnegie Mellon University, where the PI holds an Adjunct Faculty position and serves as a co-PI on an active NSF grant. Finally, the PI and CMU Professor Gary Miller will organize in Puerto Rico an **NSF-sponsored workshop** on spectral graph theory. The target month is February 2012.

1.3 NSF Grants and Publication Details

The PI serves as a co-PI in the following active grant.

****

*CCF-101***** (****,****,****); */*/10

■ [***10] presented a solver for symmetric diagonally dominant (SDD) linear systems. The solver was based on a radically novel approach that eliminated the theoretical and practical limitations of the Spielman and Teng solver. The [KMP10] solver runs in time nearly $O(m \log^2 m)$, where m is the number of non-zeros in the system matrix.

■ [***11] presents an $O(\log m)$ speedup relative to the [***10] solver. The improvement stems from new insights about low-stretch trees and an accelerated algorithm for their computation. The new techniques do not only yield an asymptotic improvement, but also bring the solver closer to practicality.

■ [BGK+11] presents new parallel, work-efficient algorithms for the computation of low-diameter graphs decompositions and low-stretch subgraphs and the solution of SDD linear systems. The results imply the first work-efficient parallel randomized algorithms for several problems, including single-source shortest paths, maximum flow, min-cost flow, and lossy max-flow.
■ The code of the PI's Combinatorial Multigrid Solver (CMG) has been made available publicly on the PI's web page, and is being used by several research groups.
■ [KGLFZ+] presents a spectral graph theory-based approach to problems in neural imaging. The CMG solver is at least 4 times faster relative to all other publicly available solvers on the very large SDD systems generated from this application.

The PI has contributed as a student and postdoctoral scholar in the following grant.

Spectral Graph Theory and Its Applications
*CCF-06****** (***,***,...); */*/07 to */*/10

The grant supported the following work.

■ [**07] gave a linear-work parallel algorithm for solving planar diagonally dominant linear systems. This is the first asymptotically optimal linear system solver for a large class of matrices.

■ [**08] presented the first linear-work parallel algorithm for constructing provably good 'Steiner' preconditioners for the very important in practice class of graphs with constant average degree. The algorithm is remarkably practical.

■ [**09b] presented connections of the Steiner preconditioning with the so-called Algebraic Multigrid algorithms (AMG), for which little theory was known. A new Combinatorial Multigrid Solver, the first practical solver with strict convergence guarantees, was implemented. A related journal publication is in preparation.

■ [***09] demonstrated that our fast linear system solver can be used as an algorithmic primitive, providing the computational core for algorithms in Computer Vision. The list of algorithms includes de-noising, segmentation, and registration. An extended journal version of the paper is about to be published [***11]. ■ [TKI+08] tackled tough problems in the computational analysis of Ophthalmological OCT scans using algorithms based on spectral graph theory.

■ Two CMU patents acknowledging the NSF support have been filed:
[i]. Method and Apparatus for Solving Weighted Planar Graphs, Filed 11/29/2007, #11/998,382.[ii]. Methods and Apparatuses for Solving Graph Laplacians, Filed 6/12/2009, #12/483,750.

CCF-0635257 also supported the PI's research in parameterized algorithms.

■ [***08, **09] presented new algebraic algorithms for a long list of parameterized problems that include subgraph isomorphism problems with applications in the detection of signaling pathways in biological networks. The algebraic algorithms improve exponentially upon algorithms based on the color-coding method.

2 Introduction

The discovery of nearly-linear* time solvers for linear systems on graph Laplacians [ST04] has motivated the design of faster algorithms for a growing list of problems that includes generalized lossy flow problems [SD08], generating random spanning trees [KM09a], optimization problems in Computer Vision [KMT09], and more recently the max-flow problem [CKM$^+$11]. In more recent developments, solvers for Laplacian systems reached a near-optimal complexity [KMP10, KMP11] and they were parallelized [BGK$^+$11], facts that justify even more their characterization as **powerful algorithmic primitive** [Spi, Ten10].

The near-optimality of Laplacian solvers extends easily to algorithms for the approximate computation of a few extreme eigenvectors [ST06]. Thus an immediate by-product of Laplacian solvers was the acceleration of previously known **spectral partitioning** algorithms and heuristics. The impact was great on the relative practical value of several spectral partitioning heuristics (e.g. [TM06]). In algorithms, 'Spectral partitioning works' now in nearly-optimal time for bounded-degree bounded-genus graphs [ST96, Kel04]. In the general case, the Laplacian (eigen)-solver provides a proof that non-trivial (but rather weak) guarantees for the sparsest cut problem are indeed possible in nearly-linear time, via the Cheeger inequality (for example see [Mih89, Chu97]).

However, despite the arguably fundamental nature of eigenvectors, the availability of fast eigensolvers has not so far led to other algorithmic improvements. In particular, the research on spectral partitioning seems to have stagnated [GM98], when several other approaches have produced various trade-offs between running time and polylogarithmic approximation [LR99, ARV04, KRV06, She09, Mąd10]. However, the design of a nearly-linear time algorithm with a polylogarithmic approximation ratio remains an open problem. Spectral methods appear to be a natural candidate approach, but it is clear that a **major advance** in their understanding is required. Is such an advance even possible?

The affirmative answer will come from the proposed extension of spectral graph theory to a **spectral theory for cut structures**, defined as sets of graphs with approximately prescribed cuts. This conceptual leap leads naturally to a notion of **spectral modification** which suggests the transformation of the input graph A to a graph B with similar cuts, but with the additional property that spectral partitioning works well for it. In fact, graph B should in some sense be a **spectral maximizer**; its eigenvalues must be as high as possible. Thus, in contrast to the **spectral sparsification** problem [ST04, SS08, BSS09] which strives to preserve the spectrum of a graph, the **spectral maximization** problem tries to alter the spectrum as much as possible while maintaining the cuts. The concept of spectral modification and the proof that every cut-structure has a 'good' spectral maximizer will draw from connections between spectral graph theory and the Räcke's hierarchical graph decompositions for oblivious routing [Räc02, BKR03].

The existence of a spectral maximizer for a cut structure will induce a natural measure of **spectral distance** within a given cut structure. It will be shown that spectral distance is a notion with significant **descriptive value**. For instance, it will be shown that on **fixed graphs** the **Cheeger gap** is a function of the spectral distance, not only for the sparsest cut but —via generalized Cheeger inequalities– to **cuts generalized** on pairs of graphs, including the *s*-*t* cut and the uniform sparsest cut. This realization will enable the development of partitioning algorithms that make use of prior information about the location of the cut, improving –for example– **medical imaging**. In later stages, the project will examine the the descriptive quality of the spectral distance for the performance of algorithms and other phenomena in interesting classes of graphs. Examples include routing algorithms, navigability, or information spreading in social or biological networks.

The project will strive to find algorithms that –besides being asymptotically near-optimal– are elegant and practical. The key for efficiency in many algorithms will be SDD solvers, which will be developed further. The project will disseminate the results of the research in the form of publications and software.

*We will be using the term 'near-linear in n' to mean $O(n \log^c n)$ for some constant c.

3 Proposed Research

In this Section we explain the proposed research problems and outline potential methods for solving them. Statements whose proofs are known to the PI and his collaborators will be called Claims.

3.1 Definitions and notation

We start with a review of some basic background material and definitions. Given a weighted graph $A = (V,E,w)$ the Laplacian L_A of A is the matrix defined by:

- $L(i,j) = L(j,i) = -w_{i,j}$
- $L(i,i) = \sum_{j \neq i} w_{i,j}$

For any vector x, we have

$$x^T Lx = \sum_{ij} w_{i,j}(x_i - x_j)^2.$$

For $S \subset V$ we let $cap_A(S,V-S)$ denote the total weight crossing the cut between S and $V-S$ that is

$$cap_A(S,V-S) = \sum_{i \in S, j \in V-S} w_{i,j}$$

Let y_S denote the 0-1 vector indicating S, i.e. $y_S(i) = 1$ if $i \in S$ and $y_S(i) = 0$ otherwise. Then it can be seen that

$$cap_A(S,V-S) = y_S^T L_A y_S. \quad (1)$$

Given two Laplacians L_A and L_B, a generalized eigenvalue λ and the corresponding generalized eigenvector x are a pair that satisfy

$$L_A x = \lambda L_B x.$$

A Laplacian pair (L_A, L_B) has $n-1$ positive generalized eigenvalues corresponding to eigenvectors different than the constant vector. Among these $n-1$ generalized eigenvalues let $\lambda_{\min}(L_A,L_B), \lambda_{\max}(L_A,L_B)$ denote the minimum and maximum one respectively. We have

$$\lambda_{\min}(L_A,L_B) = \min_x \frac{x^T L_A x}{x^T L_B x} \quad \text{and} \quad \lambda_{\max}(L_A,L_B) = \max_x \frac{x^T L_A x}{x^T L_B x}$$

The condition number $\kappa(L_A,L_B)$ is defined by

$$\kappa(L_A,L_B) = \lambda_{\max}(L_A,L_B)/\lambda_{\min}(L_A,L_B)$$

Intuitively, the condition number is a measure of how similar the two graphs are, in a spectral sense. We say that B *dominates* A when $\lambda_{\min}(L_B,L_A) \geq 1$ and we denote this event by $L_A \preceq L_B$.

3.2 Spectral Partitioning: Cut Problems as Discrete Generalized Eigenvalues

A main focus of this project will be the following simultaneous cut problem:

The generalized sparsest cut problem. Given two graphs A and B let

$$\phi(A,B) = \min_{S \subseteq V} \frac{cap_A(S,V-S)}{cap_B(S,V-S)}.$$

We call the set $S \subseteq V$ that achieves $\phi(A,B)$ the **generalized sparsest cut** for the pair (A,B). Given a set $S' \subseteq V$ such that

$$\frac{cap_A(S', V - S')}{cap_B(S', V - S')} = \rho\phi(A,B)$$

we say that S' is a ρ-approximation to the generalized sparsest cut.
We now discuss how three well known problems are instantiations of the generalized sparsest cut problem.

The uniform sparsest cut problem. Let K be the complete graph with uniform weights equal to $1/n$. Then

$$\phi(A,K) = \min_{S \subseteq V} \frac{cap_A(S, V - S)}{|S||V - S|/n}$$

which is within a factor of 2 of

$$\tilde{\phi}(A,K) = \min_{S \subseteq V} \frac{cap_A(S, V - S)}{\min\{|S|, |V - S|\}}.$$

The sparsest cut problem. Let $vol(S)$ denote the total weight incident to $S \subseteq V$. Let P be a weighted complete graph, with the weight $w'_{i,j}$ of edge (i,j) given by the product

$$w'_{i,j} = vol(i)vol(j)/vol(V).$$

Then it can be shown that

$$\phi(A,P) = \min_{S \subseteq V} \frac{cap_A(S, V - S)}{vol(S)vol(V - S)/vol(V)}$$

which is within a factor of 2 of

$$\tilde{\phi}(A,P) = \min_{S \subseteq V} \frac{cap_A(S, V - S)}{\min\{vol(S), vol(V - S)\}}.$$

Both $\phi(A,P)$ and $\tilde{\phi}(A,P)$ are often referred to as the **conductance** of the graph A. We will denote the conductance of A by $\phi(A)$.

The minimum s-t cut problem. Let E_{st} be the graph consisting of only one edge between the nodes s and t. Then $\phi(A, E_{st})$ is the value of the min s-t cut in A, and the set $S \subseteq V$ that achieves $\phi(A, E_{st})$ provides the actual min s-t cut.

It is clear that given two graphs A and B on n nodes we have

$$\phi(A,B) = \min_{x \in \{0,1\}^n} \frac{x^T L_A x}{x^T L_B x}.$$

In particular, $\phi(A,B)$ can be considered as a **discrete generalized eigenvalue** of the pair (L_A, L_B) as it is evident that relaxing the definition of $\phi(A,B)$ over the real numbers yields the definition of $\lambda_{\min}(L_A, L_B)$.

Spectral partitioning attempts to solve the discrete generalized eigenvalue problem by computing approximations to the actual generalized eigenvalue and eigenvectors, and 'rounding' the eigenvector to a discrete vector indicating a cut which is hopefully a good approximation to the generalized sparsest cut. The spectral partitioning method has been used mostly for the sparsest cut problem. The generalized eigenproblem (A, E_{st}) also appears implicitly in the recent max-flow and s-t cut algorithms of [CKM$^+$11] where it is viewed as an equivalent linear system. The linear systems produces electrical flows that are crude approximations of the combinatorial flows. As it will be explained in Section 3.7, the (in some sense) dual view of the electrical flow as an eigenvector yields additional insights.

3.3 Cut Structures and Spectral Maximization

For a given graph $A = (V, E, w)$ with $|V| = n$, let the **cut structure** $\mathcal{C}_A$ be the set of graphs that have the same cuts with A, within some factor c. Concretely, for some number $c \geq 1$ let

$$\mathcal{C}_A(c) = \{B : \phi(B,A) \geq 1 \;\text{ and } \phi(A,B) \geq 1/c\}.$$

We say that a graph $B \in \mathcal{C}_A$ is a d-**spectral maximizer** for $\mathcal{C}_A(c)$ if for all $A' \in \mathcal{C}_A(c)$, we have

$$L_{A'} \preceq dL_B$$

In other words, B spectrally dominates within a d factor all graphs in $\mathcal{C}_A(c)$.

A basic question we will consider in this project is the following:

Question: Are there absolute constants a, b such that for all graphs A there is a graph $B \in \mathcal{C}_A(\log^a n)$ which is a $\log^b n$-spectral maximizer for $\mathcal{C}_A(\log^a n)$?

3.4 Räcke Decompositions, Spectral Maximization and Spectral Distance

A laminar decomposition tree $T = (V_T, E_T, w_T)$ of a graph $A = (V, E, w)$ is defined by the following properties:

- Each node v of T corresponds to a set $S_v \subseteq V$.
- The root of T corresponds to V and the n leaves to the nodes in V.
- If $u_1, \ldots, u_k$ are the children of v in T then $S_v = \bigcup_{j=1}^{k} S_{u_j}$.
- The weight of the edge from $u \in V_T$ to its parent in T is equal to $cap_A(S_u, V - S_u)$.

Räcke [Räc02, BKR03, Räc08] introduced a special kind of laminar (or hierarchical) decompositions with several properties that lend to the design of a nearly optimal oblivious routing strategy. For example, for the square grid graph, Räcke's tree is the natural quadtree, with its edges weighted by the values of the underlying cuts.

Let T be the Räcke tree of A, and L_T be its Laplacian, ordered so that the non-leaf nodes appear in coordinates greater than n. Let b be an n-dimensional vector orthogonal to the constant vector. Then there are vectors x_b and y_b such that

$$L_T \begin{pmatrix} x_b \\ y_b \end{pmatrix} = \begin{pmatrix} b \\ 0 \end{pmatrix}$$

Gremban [Gre96] showed that there is a (dense) graph $B[T]$ such that for all b

$$L_{B[T]} x_b = b.$$

A key to the **graph modification** approach will be the following.

Claim. If T is Räcke's tree for A then $B[T]$ is a $\log^b n$-spectral maximizer for $\mathcal{C}_A(\log^a n)$, for some fixed integers a, b independent from A. Hence the answer to the question posed in section 3.3 is affirmative.

It is important to observe that -because of the slack in the definitions- a spectral maximizer is not unique. In particular the existence of a spectral maximizer for every graph also implies the existence of a *sparse* spectral maximizer [BSS09].

In what follows we will drop the functions $\log^a n$ and $\log^b n$ from the notation. By definition, a maximizer M nearly-dominates all graphs in $\mathcal{C}_A$, including A itself. The eigenvalue $\lambda_{\max}(M,A) = 1/\lambda_{\min}(A,M)$ is a natural measure of the **spectral distance** of A from M. When $\lambda_{\max}(M,A)$ is close to 1 then A is itself a spectral maximizer.

3.5 Generalized Cheeger Inequalities

Let us focus on the conductance $\phi = \phi(A)$ of a graph A. The guarantees for spectral partitioning follow from the Cheeger inequality (see for example [Chu97]), which relates ϕ with $\lambda_{\min}(L_A, L_P)$, where P is the matrix defined in Section 3.2 †:

$$\lambda_{\min}(L_A, L_P) \leq \phi \leq 2\sqrt{\lambda_{\min}(L_A, L_P)}.$$

The proof of the Cheeger inequality actually rounds the generalized eigenvector to a 0-1 vector which achieves a cut of sparsity at most $\phi' \leq 2\sqrt{\lambda_{\min}(L_A, L_P)}$. However the actual conductance may be as small as $\lambda_{\min}(L_A, L_P)$. Hence the cut returned by the rounding is not an $f(n)$-approximation for the sparsest cut, for any non-trivial function of n.

The Cheeger inequality assumes a somewhat different form in the case of the uniform sparsest cut. In particular, in the case the graph A is uniform, and the maximum degree is d, the inequality becomes

$$\lambda_{\min}(L_A, L_K) \leq \phi(A, K) \leq \sqrt{2d\lambda_{\min}(L_A, L_K)}. \tag{2}$$

In this case, the guarantee is weaker for graphs which contain high degree nodes.

The Cheeger inequality leads naturally to the definition of the **Cheeger gap** of a pair (A, B) as the ratio $\phi(A, B)/\lambda_{\min}(L_A, L_B)$ There is no explicitly known bound on the gap for the pair (A, E_{st}), let alone for the generalized sparsest cut (A, B). It will be proven that the Cheeger gap of the pair (A, B) is (within a constant) at most $1/\phi(A)$, **independently from** B. More concretely, the following will be shown.

Claim A. For any pair of graphs A and B, and some fixed constant C, we have

$$C\phi(A, B)\phi(A) \leq \lambda_{\min}(L_A, L_B) \leq \phi(A, B) \tag{3}$$

It can be seen immediately that the inequality recovers the usual Cheeger inequality when applied to the sparsest cut problem. It can also be seen than –when specialized to the uniform problem– inequality 3 is more natural, and in many cases stronger than inequality 2. We will discuss the consequences on the s-t cut problem in Section 3.7.

The conductance can be a very pessimistic upper bound for the Cheeger gap. A central claim of this proposal will be that the quality of spectral partitioning is a function of the spectral distance, quantified by the following inequality.

Claim B. Let A be a graph and M be a spectral maximizer for the set of graphs $\mathcal{C}_A$ as defined in Section 3.3. Then for any graph B, and some constant c, we have

$$\phi(A, B)\frac{\lambda_{\min}(L_A, L_M)}{\log^c n} \leq \lambda_{\min}(L_A, L_B) \leq \phi(A, B).$$

In other words, the Cheeger gap of (A, B) is upper bounded (up to $\log^c n$) by the spectral distance of A from its maximizer, again **independently from** B.

3.6 Spectral Modification

The Cheeger inequality in Claim B of the previous Section suggests a **spectral modification** approach. Given a graph A, find a spectral modifier graph $B \in \mathcal{C}_A$ such that B is **spectrally close** to a spectral maximizer M for $\mathcal{C}_A$. Then apply **spectral partitioning** on B. Observe that since B is in $\mathcal{C}_A$, it has (up to poly-log

† Most proofs of the Cheeger inequality use the pair (L_A, D) instead of (L_A, L_P), where D is the diagonal of L_A. Using (L_A, L_P) doesn't affect the inequality.

factors) the same cuts with A, and at the same time spectral partitioning provides stronger guarantees for B via Claim B. These stronger guarantees can then be translated back to A with only a polylogarithmic loss.

The major goal of this approach is to design a nearly-linear time algorithm that returns a spectral modifier B such that

$$\lambda_{\min}(L_B, L_M) > f(n)$$

where $f(n)$ is a non-trivial function of n only. This will be a major step forward for spectral partitioning and it may eventually lead to a nearly-linear time algorithm with poly-logarithmic approximation ratio.

We plan to explore the following general framework for constructing the spectral modifier B.

1. Compute a graph $A' \in \mathcal{C}_A$ with $O(n \log n)$ edges, using for example [SS08].
2. Identify in A' a number of $\log^c n$ weighted spanning trees T_i
3. For each tree T_i find its spectral maximizer $R_i \in \mathcal{C}_{T_i}$.
4. Let B be the sum of A' and the R_i's in the sense that $L_B = L_{A'} + \sum_j L_{R_j}$.

This general framework guarantees that B will be in $\mathcal{C}_A$ and that $B \succeq A$. In fact a variant of this framework returns polylogarithmic guarantees for the counterexamples in [GM98].

Comment: Mądry [Mąd10] recently proposed a framework that generates partitioning algorithms whose running time is (roughly) $O(m^{1+1/\varepsilon})$ at the expense of an $O((\log n)^{1/\varepsilon})$ factor in the approximation ratio. At a high level, his algorithm generates a *decomposition* of the given graph into a set of graphs (called j-trees), solves the partitioning problem on a small randomly selected subset of the j-trees, and extracts an approximation for the problem on the input graph. The computation of the trees T_i in step 3 of our framework is similar only superficially to Mądry's graph decomposition. Our framework doesn't approach the problem by reducing in to smaller or 'easier' graphs, as this imposes very stringent requirements on the decomposition. Rather, it is based on the plausible conjecture that one can modify the spectrum of the graph by operating on about $O(\log n)$ trees since this number is enough to get a cut-preserving sparsifier [FHHP11]. In particular, the 're-wiring' of the graph that effectively takes place in step 4 doesn't have an analogue in Mądry's approach and it may be inherently more powerful.

3.7 The descriptive power of spectral distance

The proposal will show that **spectral distance** is a measure of fundamental nature in understanding the performance of known algorithms on families of graphs. It will also examine the design of algorithms whose running time is **parameterized** by the spectral distance. Particular instantiations of this reasoning include the following.

• **Spectral partitioning:** The discussion of Section 3.5 already implies that spectral partitioning works better than previously thought for certain classes of graphs. For example, we conjecture that it works better for *dense graphs*:
Conjecture. Let A be a uniform graph with average degree d and B any graph. We have

$$\phi(A,B)\frac{n}{d \log^c n} \leq \lambda_{\min}(L_A, L_B) \leq \phi(A,B).$$

• **Approximate max-flow and *s*-*t* min cuts:** The recent approximate max-flow algorithm [CKM+11], is based on computing electrical flows which are used as (crude) approximation of combinatorial flows. The quality in the approximation determines the number of electrical flows that have to be computed and controls directly the running time of the algorithm. An electrical flow can also be viewed as the eigenvector of the generalized system (A, E_{st}). The inequalities in Section 3.5 imply that the approximate min *s*-*t* cuts computed via a rounding of the eigenvector can be good approximations of the exact min *s*-*t* cut if the spectral distance is small. This suggests the possibility that the quality of the electrical flow as an approximate combinatorial is a function of the spectral distance which may lead to an improved running time on graphs with smaller spectral distance. Such improved bounds already exist for other approaches [LR99].

• **Faster Computation of Räcke Decompositions.** Given the relationship between maximizers and Räcke Decompositions one can consider the following **problem:** Design an algorithm for computing a Räcke decomposition for A, provided that A is a near-maximizer for the set $\mathcal{C}_A$. The algorithm should be faster than known algorithms, at the expense of some poly-log factors for the parameters capturing the quality of the decomposition.
It is clear that if there is a positive answer to the promise problem, then combining it with the spectral modification approach will yield a faster algorithm for computing good Räcke decompositions for all graphs.

• **Information-passing in networks.** In the longer term the project will investigate the descriptive power of spectral distance in contexts including (i) the congestion-efficiency of simple routing protocols, (ii) rumour spreading mechanisms [CLP10, CHHKM11], and (iii) network navigability [Kle00, CFL08].

3.8 The ground work: Practical SDD Solvers

The proposal aims to design algorithms that are not only asymptotically nearly-optimal but also elegant and **practical**. Practical SDD linear system solvers are necessary to attain this objective. The currently fastest SDD solver [KMP11] is believed by many to be asymptotically optimal, up to $\log\log n$ factors. However, it appears possible that SDD solvers based on newer ideas can come closer to practice. From a theoretical point of view this requires new ideas that will solve certain carefully posed problems:

• **Practical spectral sparsification.** The solvers for Laplacian systems, or more generally for Symmetric Diagonally Dominant (SDD) systems are based on **spectral sparsification**, which given a graph A seeks to 'modify' it to a sparser graph B, such that the condition number $\kappa(L_A, L_B)$ is small. In particular, the $\tilde{O}(m \log n)$ time algorithm of [KMP11] is based on the notion of **incremental sparsification** which given the graph A returns a graph B with $n + m/\log^2 n$ edges such that $\kappa(L_A, L_B) < \log^4 n$. Even for mildly dense graphs, incremental sparsification is quite weak comparing to the **full sparsification** of Spielman and Srivastava [SS08] who proved that for every graph A there is a graph B with $O(n \log n)$ edges such that $\kappa(L_A, L_B) < 2$. Furthermore, B can be computed with $O(\log n)$ calls to a Laplacian solver.

Our experience with implementations of SDD solvers suggests that a full sparsification routine would be very useful in practice. It remains an open problem whether a full sparsifier with $O(n \log n)$ edges can be computed via a **combinatorial sparsification** algorithm without resorting to a Laplacian solver. We propose then to study the following problem.

Problem. Find an algorithm that runs in time $O(m + n \log^c n)$ which given a graph A returns a graph B with $O(n \log n)$ edges, such that $\kappa(L_A, L_B) < O(\log^2 n)$.

A candidate approach could be based on an interesting recent development [FHHP11], summarized in the following proposition.

Proposition. There is an algorithm that runs in time $O(m + n \log^c n)$ which given a graph A returns B with $O(n \log n)$ edges such that $1/2 < \phi(A, B)\phi(B, A) < 2$.

The above proposition claims the existence of a **cut-preserving sparsifier** which has the running time properties we seek, but it only guarantees a small **discrete** condition number. We plan to address the problem by proving a variant of the following conjecture.

Conjecture. There is an algorithm that runs in time $O(m + n \log^4 n)$ which given a graph A, it decomposes the vertices V of a graph A into vertex-disjoint clusters V_i such that: (i) Only a constant fraction of the edges of A connect different clusters, (ii) For all i, running [FHHP11] on the subgraph $A[V_i]$ induced by V_i returns a graph B_i such that $\kappa(L_{A[V_i]}, L_{B_i}) < O(\log^2 n)$.

• **Bypassing low-stretch tree computations.** A major obstacle in the practicality of the [KMP11] solver is the complicate nature of the algorithm for computing a low-stretch tree [AKPW95, EEST05, ABN08]. For the special case of spectral partitioning, the SDD system doesn't have to be solved for a fixed a graph A, but for **some graph** in the cut structure $\mathcal{C}_A$. We will show the following.

Claim. There is a graph B in $\mathcal{C}_A(4)$ such that, in effect, a spanning tree of B with average stretch $\log n$ (no hidden constants) can be found in linear time.

• **Parallel SDD solvers.** The recent parallel algorithm in [BGK$^+$11] runs in time nearly $O(n^{1/3})$. Perhaps the most significant open question concerning SDD solvers is the design of a parallel algorithm which does $O(m \log n)$ of $O(m \log^2 n)$ work, and runs in **polylogarithmic parallel time**. This question seems to be outside the grasp of current methods and previous work shows that it will most probably require a deeper understanding of generalized eigenvalues and eigenvectors [Kou07].

3.9 Broader consequences in Spectral Graph Theory

This proposal will bring us closer to major advances in algorithms, such as the design of nearly-linear time algorithms for approximate max-flow. But beyond the realm of very efficient algorithms, the novel notions of this proposal will be fertile ground for research that will introduce new techniques and may uncover beautiful theorems. Similar to the notion of spectral sparsification that led to the discovery of twice-Ramanujan graphs [BSS09], the notion of spectral modification may very well lead to breakthroughs of similar magnitude. A wealth of interesting questions can be posed. There is however one that stands out and has the potential for surprising consequences:

What are the theoretical limits of spectral maximization?

In more technical terms, what is the smallest function $\theta(n)$ such that every cut structure $\mathcal{C}_A(c)$ contains a $\theta(n)$-maximizer, for some constant c? And how fast can a $u(n)$-maximizer can be computed as $u(n)$ varies?

We are indeed optimistic that the proposed research activities will open a new cycle of innovations in the foundations of algorithms and in theoretical computer science more generally.

4 Broader Impacts in Applied Sciences and Industrial Applications

4.1 Neural Imaging and Connectomics

The automated or semi-automated segmentation of individual neurons in electron microscopic (EM) images is a crucial step in the acquisition and analysis of **connectomes**, i.e. maps of the neural connections [STK05, LS08, JST10]. The acquisition of connectomes is a computational problem of colossal size, and algorithms for it are in their infancy. The PI will collaborate closely with Professor *************, and his biological imaging group at ***********. A related **letter of collaboration** is included.

Professor ******************** and his group will not only provide data to the project but also develop **customized** EM imaging techniques based on interactions with the PI's group. In particular, the project will consider **genuinely 3D** (three-dimensional) algorithmic approaches on which little work has been done, despite their clear advantages. This is due to the misguided (given the progress in solvers) intuition that approaches which use contextual information from distant parts of the image should be computationally infeasible [JST10]. The initial joint work [KGLFZ$^+$] shows that dendrites of the same neuron that appear disconnected or ambiguously connected in 2D frames, are ultimately connected through the soma in the serial 3D image. In addition, the expected noise artifacts and the natural fluctuations in the brightness of pixels have an adversary effect on the connectivity properties of 2D affinity graphs, but they are insignificant in a 3D context.

4.2 Industrial Applications in Computer Vision

Symmetric Diagonally Dominant systems are ubiquitous in Computer Vision applications [KMT11]. Several segmentation algorithms are based on *s-t* min cuts computations [BFL06]. The PI collaborates with ObjectVideo on deploying to the real-world solutions based on SDD solvers and approximate *s-t* min-cut algorithms stemming from this proposal. A related **letter of collaboration** is included.

5 Broader Impacts in Education

The *************** is the largest Hispanic-serving institution in the United States. The PI is the newest member of the Computer Science faculty which currently consists of 9 members. The department enrolls a total of about 110 students. Between 2005 and 2010, an average of 7 students have graduated per year. The enrollment and completion numbers are similar for the Mathematics department, which however has a significantly larger faculty. The two departments have a joint graduate program in Applied Mathematics.

An understanding of the **challenges and opportunities** at *********** is the first step to the development of appropriate educational activities. The PI has shaped his views around his numerous interactions with faculty and students and his personal experience with more than one educational system; the PI received his undergraduate degree in ********* and is aware of the persistent impact (positive or negative) that **social and cultural traits** can have on all levels of education.

5.1 Challenges and Etiology

It is well documented that the Hispanic demographic group lags in educational attainment and is underrepresented in science and engineering [fES03]. Secretary Arne Duncan recently called the improved academic performance of Hispanic children a **"national priority".**

Puerto Rican students who live in the island have somewhat **different characteristics** comparing to Hispanics in the main US. They don't face neither social inclusion and assimilation difficulties nor a significant language barrier as a result of the fact that the English language is mandatory by law throughout schooling. An interesting additional fact is that University students in Puerto Rico receive **significant financial support** that dramatically reduces the financial burden of their studies. In fact, the availability of financial assistance motivates several 'stateside' Puerto Ricans to return to their homeland for their undergraduate studies. However, the statistics clearly indicate that –especially in Computer Science– an overwhelming majority of Puerto Rican students **do not fully realize their potential**; from a total population of 4 million, excluding certain private non-accredited institutions, the total number of CS graduates per year is not more that 20!

The problem is obviously multifaceted, and with deep social and cultural roots. There is however a number of causes that are relatively isolated and addressable within the University:

1. Many students enter the department with **excitement** about the more illustrious aspects of Computer Science such as video games, graphics, machine learning, or smartphone applications. However, a non-negligible percentage of the students, being frequently **first generation students**, they also enter with a **conflicting view** of the University as some type of vocational school, where the will quickly learn skills in order to get a job. On top of that, they start with **wrong expectations** about the amount of work required to achieve that. In some (relative) sense, many students lack the psychological mechanism or preparation for a **delayed professional gratification**, that will get them through the many stages of sometimes mundane preparation required to reach what they imagine. At the same time, the lack of 'light of the end of the tunnel'-course offerings in specialized topics disappoints the students who quickly become indifferent, move on to other majors or even drop out under financial pressure.

2. In an environment where **student disappointment** abounds, the teaching work becomes even more difficult. The faculty can't work at the rate they would like to work, **faculty disappointment** sets in and reflects in the classroom, essentially creating a vicious cycle. While it is certainly very difficult to accelerate and –within a short period– increase the demands to the students, it is often the case that faculty will go the opposite way and **reduce their demands**, in order to alleviate student disappointment, so that the weaker students *"at least learn the basic"*. At the same time, the faculty who want to increase the intensity of their classes will often be regarded by the students as very strict or even unreasonable, as a result of an exaggerated and ill-posed sense of **social order**. These professors will often find themselves alienated from the students, until they eventually (perhaps partially) succumb to the pressure of the environment. Breaking with this pattern requires a **collective and gradual approach** over the span of a few years. Otherwise, the **culture of reduced effort** will persist and perpetuate.

3. The department and more generally ******* do not lack bright and intellectually curious, **star students**. However these students are scattered among different departments, they don't take classes simultane-ously. They meet socially, but **they don't meet intellectually**. As a result, they usually don't benefit from intellectual kinship, or a feeling of adventure and competition. At the same they are conscious of the fact that they live in a culture of reduced effort; the PI is aware of students who compare ******* courses to similar courses in top institutions by watching lecture videos that are available online. This **disappointment of the bright** leads them to pass on opportunities as they feel that the environment drags them behind.

   Overall, the result is **underachievement** at **all levels of ability**. However, with appropriate stimulation the **excitement** can be **transcended to dedication** and hard work. **Collective faculty efforts** can fight off the culture of reduced effort. **Carefully designed** courses and research projects can restore the **optimism of the bright**.

5.2 Opportunities

The department of Computer Science at *********** offers a number of unique opportunities.

1. The small faculty to student ratio helps bridging the gap between the faculty and the students. The friendliness of the environment enables a closer monitoring of each student *individually*. This instills a higher **sense of responsibility** in the students.

2. The department plans to hire 5 new faculty members in the next 5 years. Recruiting faculty with expertise **complementary** to that of the current faculty will improve in several ways the quality of the program, and increase the availability of **course offerings** that will motivate the students to increase their **effort level**. The new hires will also facilitate undertaking a collective department-wise effort to gradually increase the **intensity** of the program.

3. The department has an established mentality for **undergraduate research**. Most students participate small research/development projects for credit, and many participate in NSF-sponsored REUs (Research Experiences for Undergraduates). In practice this individualizes the learning process and offers an attractive alternative to the traditional lecture/homework/exam approach.

4. The department is gradually assuming a more **active role** in the Math-CS graduate program and is examining its participation in at least one additional **interdisciplinary** MSc program. These are the first steps to increase the popularization of the value of computational thinking within the educational system of Puerto Rico.

5.3 Educational plan component I: Departmental activities

During the academic year 2011-2012, the PI will begin serving the department as a **coordinator of undergraduate research** and also as a member of the **graduate CS-Math committee**. The PI will aim for:

1. Synchronization of the undergraduate research with **current and future trends** of Computer Science research, increasing the **educational effectiveness** of the program and improving the chances of the best students for admission in strong **graduate programs**. In order to achieve this goal the PI will consult with experts in various areas of Computer Science, using his professional and social network.

2. A **tighter integration** of the undergraduate and graduate programs. In particular the PI will advocate the introduction of graduate courses with a strong **algorithmic** component, increasing the appreciation of faculty and students for **Theoretical Computer Science**. This goal will be supported by inviting selected researchers in Theory to visit the department; their presence will lend credibility to the PI's effort.

The PI also expects to play a leading role in the recruitment process, which will be **crucial** for the **near and mid-term future** of Computer Science in Puerto Rico. Besides the usual advertisements the faculty search will include **direct invitations** for interviews.

The project will support the departmental activities of the PI with academic year salary. It will also support invitations for visits of distinguished researchers. **Most importantly**, the prestige of a CAREER award, will increase the credibility and visibility of the department and will exercise a **strongly positive influence** in the outcome of the faculty search.

5.4 Educational plan component II: Curriculum development activities

The PI will build his curriculum development activities around three observations: **(a)** The current curriculum has a limited number of course offerings whose subject can capture the imagination of **current and prospective** students. **(b)** While many students want to acquire knowledge of methods and techniques, they don't appreciate analogously the **intellectual value** of proofs, and the process of discovery. **(c)** The current structure of the curriculum doesn't serve very well the purpose of **bringing the best minds together**.

At least until the department expansion materializes, the PI will make a considerable effort to teach classes outside his immediate expertise. For example, in the Fall 2011 semester he will introduce a new **'Undergraduate Machine Learning'** course. The current –much higher than average– enrollment for the class validates the PI's hypothesis that students are excited in the promise of **knowledge of novel type**.

The PI has introduced a new advanced undergraduate class titled **'Linear Algebra for Computer Scientists'**. He plans to develop it further and offer a graduate version. The class reviews and teaches concepts in Linear Algebra via a graph theoretic point of view, motivated and highlighted by visual and social networking applications. In addition, algorithms for basic problems in linear algebra are taught with an emphasis on efficiency. The students are exposed **in class** to the vast difference between, for instance, quadratic and linear time algorithms, in order to appreciate algorithms as **economical and nature-friendly** ways of solving problems. In addition, the class presents and motivates notions through the power of examples and then attempts to walk the students through the full **discovery and proof** of theorems and algorithms. Similar approaches will be used in an undergraduate algorithms class that the PI will teach in Spring 2012. Lecture notes from the classes will be distributed online.

The PI will leverage his position as the undergraduate research coordinator, to advocate the introduction of a double-credit **intense course** comprised mostly by challenging mathematical and algorithmic problems and related reviews of the fundamental nature of **computational complexity**. The course will be available only to students with high GPA, and students who have excelled in research projects.

The PI also plans to experiment with **novel** classroom and grading methods. In particular, he will judiciously trade in in-classroom problem solving with recorded/written lectures as homework, as he feels that students can benefit significantly by 'hands-on' problem-solving sessions, while they can choose their lecture-reading pace at home. In addition, in an effort to increase the intellectual and academic **communication** of the students he will experiment with peer-review assessment and grading. The PI will consult with Dr. ********** whose **letter of collaboration** is included.

5.5 Educational plan component III: Undergraduate and graduate research training

The PI has established a group of students, senior **********, and MSc student ************. The group has been supported mostly through the PI's seed budget. This project will support further research activity. To maximize the effectiveness of the training and the **research output**, the PI plans to:

1. Introduce students to spectral graph theory via its most **applied aspects**, by assigning to them small research/development projects with a good chance for an immediate **broader impact** in other fields. This approach has so far yielded a presentation in a top international conference for Neuroscience. As students familiarize themselves with theoretical concepts, they will be gradually moved to more theoretical problems. Indeed, the current students are gaining a good grasp of generalized eigenvalues and Cheeger inequalities and stand a good chance for more theoretical contributions.

2. Instill in the students a sense of **broader academic participation**, by including them in the joint interdisciplinary research with colleagues in ***********. This joint research will benefit equally students in other departments. This goal will also be accomplished by their participation to meetings with distinguished visitors to the department, invited by the PI and supported by this proposal.

3. Offer to the students **summer research experiences** beyond the NSF REUs. In particular, the PI plans to use his affiliation with ********** University and invite the students to work in collaboration with Professor ********* and his group of students, following the example of a successful trial in Summer 2011. The advantage over REUs is that the students will already be prepared on related material, and via the presence of the PI will have to have a shorter social adjustment time. Of course, the **story-telling** of the students when they return to Puerto Rico from a world-class environment will have a very positive impact to other students and the department as a whole. In addition it will help the PI in the recruitment of new good students. A related **letter of collaboration** by ******** is included.

4. Encourage students to participate in **local conferences** such as the annual PRISM and SIDIM conferences, and **national conferences devoted to minorities**, such as the SACNAS (Society for Advancement of Chicanos and Native Americans in Science) annual conference, so that their activity impacts Hispanic populations beyond Puerto Rico. This project will partially support related expenses not covered by the conference organizers.

5. Attempt to obtain additional financial support by utilizing the extensive network of support for Latinos, for example SACNAS and the NSF-funded Puerto Rico Annex of the Louis Stokes Alliance for Minority Participation program (PR-LSAMP).

In summary, we are optimistic that the proposal will have a long lasting positive effect in the educational system of Puerto Rico.