2014 Workshop on Optimization for Modern Computation

September 2-4, 2014

Beijing, China

www.bicmr.org/conference/opt-2014

2014 Workshop on
Optimization for Modern Computation

September 2-4, 2014

Beijing, China

www.bicmr.org/conference/opt-2014

Information for Participants

Conference Hotel For Invited Speakers

Hotel:	"Zhong Guan Xin Yuan" Global Village, Building 1 中关新园1号楼
Address:	No. 216 Zhongguancun North Road, Haidian District 北京市海淀区中关村北大街126号
Dates:	By default, the hotel room is reserved for: Sept 1st (check in) and 5th (check-out). Please let us know if you have a different arrival-departure schedule.
Arrival:	By air, please see this link By subway: line 4 to "east gate of Peking University"
Website:	`www.pkugv.com`
Tel:	+86-10-62752288

Conference Venue

Venue:	Lecture Hall, Jia Yi Bing Building 82 Jing Chun Yuan, BICMR 北京大学镜春园82号甲乙丙楼二层报告厅
Map:	PKU campus map

Meals

- Breakfasts will be complementary at the hotel.
- Lunches and dinners are provided by the workshop. Please let us know if you have any dietary restrictions or preferences.

Currency

Chinese currency is RMB. The current rate is about 6.13 RMB for 1 US dollar. The exchange of foreign currency can be done at the airport or the conference hotel. Please keep the receipt of the exchange so that you can change back to your own currency if you have RMB left before you leave China. Please notice that some additional processing fee will be charged if you exchange currency in China.

Parking at PKU Campus

If you plan to drive to PKU, please send us your license plate number; otherwise, your car cannot enter the PKU campus.

Contact Information

If you need any help, please feel free to contact

- Dr. Zaiwen Wen: +86-188-1310-8562 (in Chinese or English)
- Ms. Qiaoming Wu: +86-10-62744133 (in Chinese)

Sponsors

National Natural Science Foundation of China

Beijing International Center For Mathematical Research, Peking University

The Mathematical Center of the Ministry Education, Peking University

The Mathematical Programming Branch of OR Society of China

Committees

Organizing Committee

Naihua Xiu, Beijing Jiaotong University

Yuhong Dai, Chinese Academy of Sciences

Wotao Yin, University of California, Los Angeles

Zaiwen Wen, Peking University

Scientific Committee

Yaxiang Yuan, Chinese Academy of Sciences

Yinyu Ye, Stanford University

Weinan E, Peking University and Princeton University

Pingwen Zhang, Peking University

Conference Schedule

Each talk is 35 minutes + 5 minutes for questions.

September 2, Tuesday

08:30-09:00 Opening Ceremony

- **08:30-08:40 Welcome Address**
- **08:40-09:00 Group Photo**

09:00-10:20 Session T1
Chair: Yaxiang Yuan

- **09:00-09:40 Stephen Wright**, Asynchronous Parallel Methods for Optimization and Linear Algebra
- **09:40-10:20 Qing Ling**, Decentralized Optimization for Multi-Agent Networks

10:20-10:40 Coffee Break

10:40-12:00 Session T2
Chair: Yinyu Ye

- **10:40-11:20 Dachuan Xu**, A Complex Semidefinite Programming Rounding Approximation Algorithm for the Balanced Max-3-Uncut Problem
- **11:20-12:00 Yong Xia**, Solving Quadratic Integer Programs: Small Changes Yield Big Improvements

12:00-13:30 Lunch

13:30-15:50 Session T3
Chair: Lieven Vandenberghe

- **13:30-14:10 Xiaoming Yuan**, Accuracy v.s. Implementability in Algorithmic Design —- An Example of Operator Splitting Methods for Convex Optimization
- **14:10-14:50 Chelin Su**, Estimating Dynamic Discrete-Choice Games of Incomplete Information
- **14:50-15:30 Shiqian Ma**, Recent Developments of Alternating Direction Method of Multipliers with Multi-Block Variables

15:30-15:50 Coffee Break

15:50-17:50 Session T4
Chair: Xiaoming Yuan

- **15:50-16:30 Naihua Xiu**, Optimality and Support Projection Algorithm for Sparsity Constrained Problem
- **16:30-17:10 Deren Han**, Asymmetric Proximal Point Algorithms with Moving Proximal Centers
- **17:10-17:50 Yuan Yao**, Sparse Recovery via Differential Inclusions

18:30 Dinner

September 3, Wednesday

08:30-09:50 Session W1
Chair: Stephen Wright

- **08:30-09:10 Lieven Vandenberghe**, Decomposition by operator-splitting methods and applications in image deblurring
- **09:10-09:50 Xin Liu**, An Efficient Gauss-Newton Algorithm for Symmetric Low-Rank Product Matrix Approximations

09:50-10:10 Coffee Break

10:10-12:00 Session W2
Chair: Yu-Hong Dai

- **10:10-10:50 Liwei Zhang**, A Smoothing Majorization Method for l_2^2-l_p^p Matrix Minimization
- **10:50-11:30 Yafeng Liu**, An SSQP Framework for a Class of Composite L_q Minimization over Polyhedron
- **11:30-12:10 Dongdong Ge**, An Improved Algorithm for the L2-Lp Minimization Problem

12:10-13:30 Lunch

13:30-15:50 Session W3
Chair: Shiqian Ma

- **13:30-14:10 Jinyan Fan**, A semidefinite algorithm for checking monotonic positivity
- **14:10-14:50 Wenxun Xing**, p-Norm Constrained Quadratic Programming: Conic Approximation Methods
- **14:50-15:30 Yuhong Dai**, A Barzilai-Borwein-Like Gradient Method with Positive Stepsizes and An Extension for Symmetric Linear Systems

15:30-15:50 Coffee Break

15:50-17:50 Session W4
Chair: Liwei Zhang

- **15:50-16:30 Zhiquan Luo**, Efficient optimization algorithms for big data applications
- **16:30-17:10 Wotao Yin**, Tight Rates and Equivalence Results of Operator Splitting Schemes
- **17:10-17:50 Zirui Zhou**, Non-Asymptotic Convergence Analysis of Inexact Gradient Methods for Machine Learning Without Strong Convexity

18:30 Dinner

September 4, Thursday

08:30-09:50 Session V1
Chair: Naihua Xiu

08:30-09:10 **Jong-Shi Pang**, A Unified Distributed Algorithm for Non-cooperative Games with Non-convex and Non-differentiable Objectives

09:10-09:50 **Yinyu Ye**, Optimization with Online and Massive Data

09:50-10:10 Coffee Break

10:10-12:00 Session V2
Chair: Wotao Yin

10:10-10:50 **Tong Zhang**, Recent Progresses in Stochastic Algorithms for Big Data Optimization

10:50-11:30 **Defeng Sun**, SDPNAL+: A MajorizedSemismooth Newton-CG Augmented Lagrangian Method for Semidefinite Programming with Nonnegative Constraints

11:30-12:00 **Zaiwen Wen**, Optimization in Electronic Structure Calculation

11:30-12:30 Lunch

Abstracts

Asynchronous Parallel Methods for Optimization and Linear Algebra

Stephen Wright

We discuss two related asynchronous parallel stochastic methods. The first is based on a randomized Kaczmarz scheme for solving a linear system $Ax = b$, which is identical to a standard stochastic gradient process applied to the least-squares formulation $\min ||Ax - b||^2$. The second is based on a randomized coordinate descent method for composite convex optimization. We analyze the expected convergence behavior of both algorithms, showing that a threshold number of processors can be identified below which near-linear speedup can be expected in the parallel implementation. The analysis of AsySPCD pays particular attention to the issue of "inconsistent reads," in which the partial gradient may be evaluated at an iterate which never actually exists in storage any point in time. This talk covers joint work with Ji Liu and other colleagues at UW-Madison.

Decentralized Optimization for Multi-Agent Networks

Qing Ling

In multi-agent optimization, a group of networked agents collaboratively minimize the summation of their local cost functions with respect to a common optimization variable. Its applications include coordination of an aircraft or vehicle network, data processing of a wireless sensor network, and spectrum sensing of a cognitive radio network, in which transmitting distributed data to a fusion center is prohibitive and the network relies on collaboration of neighboring agents to fulfill the optimization task. A favorable multi-agent network optimization algorithm requires fast convergence which implies low communication cost over the network and simple implementation which translates to low computation cost for the agents.

This talk discusses (nearly) linear convergence of two multi-agent network optimization algorithms, the decentralized gradient method (DGM) and the decentralized exact first-order algorithm (EXTRA). We prove that if the local cost functions are strongly convex and have Lipschitz continuous gradients, DGM converges linearly to a neighborhood of the optimal solution. To eliminate the inaccuracy, we propose EXTRA, whose update has a similar per-iteration complexity as that of DGM. We prove that through a simple modification to DGM, EXTRA is able to converge linearly to the exact optimal solution. The convergence properties of the two algorithms are determined by the condition numbers of the cost functions, the algorithm parameters, as well as the network topology.

A Complex Semidefinite Programming Rounding Approximation Algorithm for the Balanced Max-3-Uncut Problem

Dachuan Xu

In this paper, we consider the balanced Max-3-Uncut problem which has several applications in the design of VLSI circuits. We propose a complex discrete linear program for the balanced Max-3-Uncut problem. Applying the complex semidefinite programming rounding technique, we present a 0:3456-approximation algorithm by further integrating a greedy swapping process after the rounding step. One ingredient in our analysis diffierent from previous work for the traditional Max-3-Cut is the introduction and analysis of a bivariate function rather than a univariate function. (Co-authored with Chenchen Wu, Donglei Du, and Wen-qing Xu)

Solving Quadratic Integer Programs: Small Changes Yield Big Improvements

Yong Xia

In this talk, we solve two classes of quadratic integer programs: the box-constrained nonconvex quadratic integer program and the probabilistically constrained quadratic program. We show by computational results that small changes in the reformulations when using CPLEX could highly outperform the state-of-the-art solvers.

Accuracy v.s. Implementability in Algorithmic Design —- An Example of Operator Splitting Methods for Convex Optimization

Xiaoming Yuan

Accuracy and implementability are two common yet usually conflicted criteria for developing an efficient algorithm. In this talk, I will focus on the context of convex optimization models with separable structures to show how to make a trade-off between these two criteria for some operator splitting methods originated from the PDE literature (e.g., the Douglas-Rachford and Peaceman-Rachford schemes). The resulting algorithms could be applicable to large-scale dataset; and they could find efficient applications in many areas such as statistical learning and image processing. Some theoretical results such as the convergence rates of these algorithms will also be mentioned briefly.

Augmented Lagrangian method for total variation related problems over triangulated surfaces

Chunlin Wu

Total variation regularization has been proven very useful in image processing and computer graphics applications. Recently many efforts have been contributed to efficiently solve this type of problems which are non-differentiable. Augmented Lagrangian method is one of the most efficient methods.In this talk, we will discuss this method for total variantion related problems over triangulated surfaces, including image denoising and segmentation on surfaces, as well as surface denoising.

Recent Developments of Alternating Direction Method of Multipliers with Multi-Block Variables

Shiqian Ma

The alternating direction method of multipliers (ADMM) has been widely used for solving structured convex optimization problems. In particular, the ADMM can solve convex programs that minimize the sum of N convex functions with N-block variables linked by some linear constraints. While the convergence of the ADMM for N=2 was well established in the literature, it remained an open problem for a long time whether or not the ADMM for $N \geq 3$ is still convergent until very recently. In this talk, we discuss the recent developments on the convergence of multi-block ADMM for $N \geq 3$ and its variants.

Optimality and Support Projection Algorithm for Sparsity Constrained Problem

Lili Pan, Naihua Xiu and Shenglong Zhou

Recently, the research on sparse recovery for nonlinear measurements has a growing tendency. In this talk, we mainly consider the nonlinear minimization with sparse and nonnegative constraints. By discussing tangent cone and normal cone of sparse constraint, we give the first necessary optimality conditions, α-Stability, T-Stability and N-Stability, and the second necessary and sufficient optimality conditions for the nonlinear problem. By adopting Armijo-type stepsize rule, we present a gradient support projection algorithmic framework for the problem and establish its full convergence and computational complexity under mild conditions. By doing some numerical experiments, we show the excellent performance of the new algorithm for the least squares and nonlinear problem without and with noise.

Asymmetric Proximal Point Algorithms with Moving Proximal Centers

Deren Han

We discuss the classical proximal point algorithm (PPA) with a metric proximal parameter in the variational inequality context. The metric proximal parameter is usually required to be positive definite and symmetric in the PPA literature, because it plays the role of the measurement matrix of a norm in the convergence proof. Our main goal is to show that the metric proximal parameter can be asymmetric if the proximal center is shifted appropriately. The resulting asymmetric PPA with moving proximal centers maintains the same implementation difficulty and convergence properties as the original PPA; while the asymmetry of the metric proximal parameter allows us to design highly customized algorithms that can effectively take advantage of the structures of the model under consideration. In particular, some efficient structure-exploiting splitting algorithms can be easily developed for some special cases of the variational inequality. We illustrate these algorithmic benefits by a saddle point problem and a convex minimization model with a generic separable objective function, both of which have wide applications in various fields. We present both the exact and inexact versions of the asymmetric PPA with moving proximal centers; and analyze their convergence including the estimate of their worst-case convergence rates measured by the iteration complexity under mild assumptions and their asymptotically linear convergence rates under stronger assumptions.

A Unified Distributed Algorithm for Non-cooperative Games with Non-convex and Non-differentiable Objectives

Jong-Shi Pang

We present a unified framework for the design and analysis of distributed algorithm for computing first-order stationary solutions of non-cooperative games with non-differentiable player objective functions. These games are closely associated with multi-agent optimization wherein a large number of selfish players compete non-cooperatively to optimize their individual objectives under various constraints. Unlike centralized algorithms that require a certain system mechanism to coordinate the players' actions, distributed algorithms have the advantage that the players, either individually or in subgroups, can each make their best responses without full information of their rivals' actions. These distributed algorithms by nature are particularly suited for solving huge size games where the large number of players in the game makes the coordination of the players almost impossible. The distributed algorithms are distinguished by several features: parallel versus sequential implementations, scheduled versus randomized player selections, synchronized versus asynchronous transfer of information, and individual versus multiple player updates. Covering many variations of distributed algorithms, the unified algorithm employs convex surrogate functions to handle non-smooth non-convex functions and a (possibly multi-valued) choice function to dictate the players' turns to update their strategies. There are two general approaches to establish the convergence of such algorithms: contraction versus potential based, each requiring different properties of the players' objective functions. We present the details of the convergence analysis based on these two approaches and discuss randomized extensions of the algorithms that require less coordination and hence are more suitable for big data problems.

This is a joint work with Meisam Razaviyayn, USC and Univ. of Minnesota.

Decomposition by operator-splitting methods and applications in image deblurring

Lieven Vandenberghe

We discuss decomposition methods for convex optimization, based on the Douglas-Rachford splitting algorithm applied to the primal-dual optimality conditions. The methods handle the classical structure in primal and dual decomposition (separability after removal of coupling variables and constraints) as well as extensions that involve more general sums of structured matrices such as sparse or block-circulant matrices. The primal-dual approach is often more straightforward than dual decomposition methods, which may require the introduction of a large number of splitting variables and constraints. The results will be illustrated with applications in image deblurring, including problems with general boundary conditions, domain decomposition formulations, and deblurring problems with spatially varying blurring operators. This is joint work with Daniel O'Connor.

An Efficient Gauss-Newton Algorithm for Symmetric Low-Rank Product Matrix Approximations

Xin Liu

We derive and study a Gauss-Newton method for computing the symmetric low-rank product (SLRP) XX', where X is a n by k matrix for k¡n, that is the closest approximation to a given symmetric n by n matrix A in Frobenius norm. When A=B'B (or BB'), this problem essentially reduces to finding a truncated singular value decomposition of B. Our Gauss-Newton method, which has a particularly simple form, shares the same order of iteration-complexity as a gradient method when k ¡¡ n, but can be significantly faster on a wide range of problems. In this paper, we prove global convergence and a Q-linear convergence rate for this algorithm, and perform numerical experiments on various test problems, including those from recently active areas of matrix completion and robust principal component nalysis. Numerical results show that the proposed algorithm is capable of providing considerable speed advantages over Krylov subspace methods on suitable application problems. Moreover, the algorithm possesses a higher degree of concurrency than Krylov subspace methods, thus offering better scalability on modern multi/many-core computers.

A Smoothing Majorization Method for l_2^2-l_p^p Matrix Minimization

Liwei Zhang

We consider the l_2^2-l_p^p (with $p \in (0,1)$) matrix minimization for recovering the low-rank matrices. A smoothing approach for solving this non-smooth, non-Lipschitz and non-convex l_2^2-l_p^p optimization problem is developed, in which the smoothing parameter is treated as a decision variable and a majorization method is adopted to solve the smoothing problem. The convergence theorem shows that any accumulation point of the sequence generated by the proposed approach satisfies the first-order necessary optimality condition of the l_2^2-l_p^p problem. As an application, we use the proposed smoothing majorization method to solve the famous matrix completion problems. Numerical results indicate that our algorithm can solve the test problems efficiently.

This is a joint work with Yue Lu and Jia Wu from School of Mathematical Sciences at Dalian University of Technology.

An SSQP Framework for a Class of Composite L_q Minimization over Polyhedron

Yafeng Liu

The composite L_q ($0 < q < 1$) minimization problem over a general polyhedron has received various applications in machine learning, wireless communications, image restoration, signal reconstruction, etc. In this talk, we shall present a theoretical study on this problem. Firstly, we show that for any fixed $0 < q < 1$, finding the global minimizer of the problem, even its unconstrained counterpart, is strongly NP-hard. Secondly, we derive Karush-Kuhn-Tucker (KKT) optimality conditions for local minimizers of the problem. Thirdly, we propose a smoothing sequential quadratic programming framework for solving this problem. The framework requires a (approximate) solution of a convex quadratic program at each iteration. Finally, we analyze the worst-case iteration complexity of the framework for returning an ϵ-KKT point; i.e., a feasible point that satisfies a perturbed version of the derived KKT optimality conditions. To the best of our knowledge, the proposed framework is the first one with a worst-case iteration complexity guarantee for solving composite L_q minimization over a general polyhedron.

An Improved Algorithm for the L2-Lp Minimization Problem

Dongdong Ge

We consider a class of non-Lipschitz and non-convex minimization problems which generalize the L2-Lp minimization problem. We propose an iterative algorithm that decides the next iteration based on the local convexity/concavity/sparsity of its current position. We show that our algorithm nds an ϵ-KKT point within $O(log(1/\epsilon))$ iterations. The same result is also applied to the problem with general linear constraints under mild conditions. This is a joint work with Simai He and Rongchuan He.

A semidefinite algorithm for checking monotonic positivity

Jinyan Fan

A matrix A is monotonically positive (MP) if there exists a columnwise monotonic matrix U such that $A = UU^T$. MP-matrices have many applications in order statistics. In this talk, we present a semidefinite algorithm for checking whether a matrix is MP. If it is not MP, a certificate for this can be obtained; if it is MP, a MP-decomposition of the matrix can be obtained.

p-Norm Constrained Quadratic Programming: Conic Approximation Methods

Wenxun Xing

The 2-norm and 1-norm functions are usually used in the data fitting and handling sparse problems. This talk considers a quadratic programming problem with p-norm constraints (p is no less than 1). We shall show that it is NP-hard for general case and it can be reformulated as an equivalent quadratically constrained quadratic programming problem. And then it can be reformulated to a linear conic programming problem. The approximate methods for linear conic programming are thus applied to the problem.

A Barzilai-Borwein-Like Gradient Method with Positive Stepsizes and An Extension for Symmetric Linear Systems

Yu-Hong Dai

The Barzilai and Borwein (BB) gradient method has achieved a lot of attentions since it performs much more better than the classical steepest descent method. In this work, we analyze a positive BB-like gradient stepsize and discuss its possible uses. Specifically, we present an analysis of the positive stepsize for two-dimensional strictly convex quadratic functions and prove the R-superlinear convergence under some assumption. Meanwhile, we extend BB-like methods for solving symmetric linear systems and find that a variant of the positive stepsize is very useful in the context. Some useful discussions on the positive stepsize are also given. This is a joint work with Mehiddin Al-Baali and Xiaoqi Yang.

Efficient optimization algorithms for big data applications

Zhiquan Luo

For many contemporary engineering applications involving big data, traditional general purpose optimization tools are ineffective due to their large memory requirement and high per-iteration complexity. In this talk, we will present an effective algorithmic framework for largescale optimization problems. This primal-dual framework utilizes both random and parallel updates of the coordinates and is able to solve problems with millions or billions of variables efficiently. We will discuss the convergence issues of these methods and describe several recent applications in speech enhancement, wireless communication and smart grids.

Recent Progresses in Stochastic Algorithms for Big Data Optimization

Tong Zhang

In modern big-data machine learning applications, one may encounter optimization problems so large that traditional gradient based numerical methods are difficult to handle. A modern solution to this problem is to compute the gradient approximately using a random sample. This idea leads to a number of new optimization methods with convergence performance superior to the traditional algorithms. We will review some recent progresses on this topic as well as challenges.

Non-Asymptotic Convergence Analysis of Inexact Gradient Methods for Machine Learning Without Strong Convexity

Anthony So

Many recent applications in machine learning and data fitting call for the algorithmic solution of structured smooth convex optimization problems. Although the gradient descent method is a natural choice for this task, it requires exact gradient computations and hence can be inefficient when the problem size is large or the gradient is difficult to evaluate. Therefore, there has been much interest in inexact gradient methods (IGMs), in which an efficiently computable approximate gradient is used to perform the update in each iteration. Currently, non-asymptotic linear convergence results for IGMs are typically established under the assumption that the objective function is strongly convex, which is not satisfied in many applications of interest; while linear convergence results that do not require the strong convexity assumption are usually asymptotic in nature. In this paper, we combine the best of these two types of results and establish—under the standard assumption that the gradient approximation errors decrease linearly to zero—the non-asymptotic linear convergence of IGMs when applied to a class of structured convex optimization problems. Such a class covers settings where the objective function is not necessarily strongly convex and includes the least squares and logistic regression problems. We believe that our techniques will find further applications in the non-asymptotic convergence analysis of other first-order methods.

Tight Rates and Equivalence Results of Operator Splitting Schemes

Wotao Yin

Splitting schemes are the backbone of many first-order algorithms, ranging from alternating projection, to various splitting algorithms, and to alternating direction method of multipliers (ADMM). They are simple, versatile, and nearly state-of-the-art in solving problems with multiple monotone operators, convex constraint sets, and/or convex objective functions.

This talk presents our new findings on Douglas-Rachford, Peaceman-Rachford, and ADMM algorithms such as their primal-dual equivalence and their tight convergence rates under the basic convexity assumption, as well as under stronger regularity assumptions. We show that they automatically adapt to the regularity of the problem at hand and achieve improved convergence rates. This explains why they have good performance in a wide variety of problems. In particular, we show that ADMM is "self-dual", namely, the iterates of ADMM can be recovered from those of the ADMM applied to the dual problem. In fact, ADMM is equivalent to a primal-dual algorithm applied to the saddle-point formulation, which hides one variable and thus may save memory and run faster in some situations.

The convergence analysis is joint with Damek Davies (arXiv:1406.4834 and 1407.5210). The equivalence results are joint with Ming Yan (arXiv:1407.7400).

Optimization with Online and Massive Data

Yinyu Ye

We present several analytic models and computational algorithms dealing with online/dynamic, structured and/or massively distributed data. Specifically, we discuss:

- Near-Optimal Online Linear Programming Algorithms, where the matrix data is revealed column by column along with the objective function and a decision has to be made as soon as a variable arrives;
- Least-squares with Non-convex Regularization, where we give sparse and structure characterizations for every KKT stationary solution of the problem;
- Alternating Direction Method of Multipliers (ADMM) for large-scale data, where we give an example to show that the direct extension of ADMM for three-block convex minimization problems is not necessarily convergent, and propose possible convergent variants.

Optimization in Electronic Structure Calculation

Zaiwen Wen

Minimizing the Kohn-Sham total energy functional with respect to electron wave functions, or orthogonality constraints, is a fundamental nonlinear eigenvalue problem in electronic structure calculation. These problems are very challenging due to non-convex orthogonality constraints. This talk presents a few recent advance on analyzing the convergence of the widely used self-consistent field iteration methods. A few efficient optimization algorithms will also be introduced.

A Majorized Semismooth Newton-CG Augmented Lagrangian Method for Semidefinite Programming with Nonnegative Constraints

Defeng Sun

We present a majorizedsemismooth Newton-CG augmented Lagrangian method, called SDPNAL+, for semidefinite programming (SDP) with partial or full nonnegative constraints on the matrix variable. SDPNAL+ is a much enhanced version of SDPNAL introduced by Zhao, Sun and Toh [SIAM Journal on Optimization, 20 (2010), pp. 1737–1765] for solving generic SDPs. SDPNAL works very efficiently for nondegenerate SDPs but may encounter numerical difficulty for degenerate ones. Here we tackle this numerical difficulty by employing a majorizedsemismooth Newton-CG augmented Lagrangian method coupled with a convergent 3-block alternating direction method of multipliers introduced recently by Sun, Toh and Yang [arXiv preprint arXiv:1404.5378, (2014)]. Numerical results for various large scale SDPs with or without nonnegative constraints show that the proposed method is not only fast but also robust in obtaining accurate solutions. It outperforms, by a significant margin, two other competitive publicly available first order methods based codes: (1) an alternating direction method of multipliers based solver called SDPAD by Wen, Goldfarb and Yin [Mathematical Programming Computation, 2 (2010), pp. 203–230] and (2) a two-easy-block-decomposition hybrid proximal extragradient method called 2EBD-HPE by Monteiro, Ortiz and Svaiter [Mathematical Programming Computation, (2013), pp. 1–48]. In contrast to these two codes, we are able to solve all the 95 difficult SDP problems arising from the relaxations of quadratic assignment problems tested in SDPNAL to an accuracy of 10^{-6} efficiently, while SDPAD and 2EBD-HPE successfully solve 30 and 16 problems, respectively. In addition, SDPNAL+ appears to be the only viable method currently available to solve large scale SDPs arising from rank-1 tensor approximation problems constructed by Nie and Wang [arXiv preprint arXiv:1308.6562, (2013)]. The largest rank-1 tensor approximation problem we solved (in about 14.5 hours) is nonsym(21,4), in which its resulting SDP problem has matrix dimension $n = 9,261$ and the number of equality constraints $m = 12,326,390$. [This is a joint work with Liuqin Yang and Kim-Chuan Toh at National University of Singapore].

The organizing committee wishes you a pleasant stay in BICMR!