ORSIS CONFERENCE – MAY 2011
ABSTRACTS BOOK

Naor Lecture: Peter Gritzmann

On Geometric Clustering and its Applications

In geometric clustering m (weighted) objects in some $\mathbb{R}^d$ have to be partitioned into k clusters according to certain balancing constraints so as to optimize some objective function. The most prominent example in our context is that of the consolidation of farmland but we will also indicate other applications. While the relevant specifications of the optimization problem are NP-hard in general, we use polyhedral approximations of corresponding clustering bodies to devise surprisingly tight approximations. Also we show that optimal clusterings have strong separation properties. In particular, we introduce gravity polytopes in $\mathbb{R}^{kd}$ whose vertices encode precisely the feasible power diagrams in $\mathbb{R}^d$. (Joint work with Andreas Brieden)

Invited Talk: Assaf Zeevi

Making correct decisions based on incorrect models

Most models only capture at best some key characteristics of the underlying phenomenon or system we are interested in. Performance functions which are optimized to infer operating decisions and control rules are therefore typically predicated on misspecified models. In statistics and econometrics there is a fair bit of literature that deals with model misspecification and its implications on estimation. In contrast, in this talk, we will focus on how said misspecification impacts operational decisions, and illustrate some of the key ideas in the context of a prototypical dynamic pricing problem.

ORSIS Award Talk: Yigal Gerchak* and Eran Hanany

Nash Bargaining Over Allocations in Inventory Pooling Contracts

When facing uncertain demand, several firms may consider pooling their inventories leading to the emergence of two key contractual issues. How much should each produce or purchase for inventory purposes? How should inventory be allocated when shortages occur to some of the firms? Previously, if the allocations issue was considered, it was undertaken through evaluation of the consequences of an arbitrary priority scheme. We consider both these issues within a Nash bargaining solution (NBS) cooperative framework. The firms may not be risk neutral, hence a nontransferable utility bargaining game is defined. Thus the physical pooling mechanism itself must benefit the firms, even without any monetary transfers. The firms may be asymmetric in the sense of having different unit production costs and unit revenues. Our assumption with respect to shortage allocation is that a firm not suffering from a shortfall, will not be affected by any of the other firms shortages. For two risk neutral firms, the NBS is shown to award priority on all inventory produced to the firm with higher ratio of unit revenue to unit production cost. Nevertheless, the arrangement is also beneficial for the other firm contributing to the total production. We provide examples of Uniform and Bernoulli demand distributions, for which the problem can be solved analytically. For firms with constant absolute risk aversion, the agreement may not award priority to any firm. Analytically solvable examples allow additional insights, e.g. that higher risk aversion can, for some problem parameters, cause an increase in the sum of quantities produced, which is not the case in a single newsvendor setting.

(S1A) Optimization in Machine Learning and Data Mining [Abirim 1]

What else can we do with more data?

Shai Shalev Schwartz*

In some applications, data is plentiful. By now, we have a rather clear understanding of how more data can be used to improve the accuracy of learning algorithms. In the talk I'll show how more data can be beneficiary for other purposes. In particular, I'll present algorithms that can leverage large data sets to reduce the required training runtime, prediction runtime, and algorithms that can use the multitude of examples to compensate for lack of full information on each individual example.

Clustering and Approximating High-Dimensional Streaming Data using Coresets

Danny Feldman*

A coreset (or, core-set) for a given problem is a "compressed" representation of its input, in the sense that a solution for the problem with the (small) coreset as input would yield an approximate solution to the problem with the original (large) input. Using traditional techniques, a coreset usually implies provable linear time algorithms for the corresponding optimization problem, which can be computed in parallel, via one pass over the data, and using only polylogarithmic space (i.e, in the streaming model). During the recent years, coresets were suggested for problems such as k-means clustering, facility location, linear regression, PCA, and matrix approximation.

I will give an introduction to this new paradigm, including recent implementations in the context of computer vision. Based on the paper "A Unified Framework for Approximating and Clustering Data" (STOC'11).

Improving Learning by Tracking Uncertainty

Koby Crammer*

I will focus in a class of algorithms that maintain uncertainty information about classifier parameters. Informally, the uncertainty is modeled via a covariance matrix over the weights of a linear hypotheses. Learning in this framework modify parameters by estimating weights and reducing the uncertainty. I will describe few algorithms to learn models in this framework, for which a mistake bound analysis shows that indeed the algorithms perform better under some conditions.

In the second part of the talk I will demonstrate the usefulness of maintaining confidence in few settings: obtaining state-of-the-art performance in binary, multiclass and structured predictions; improving classifier combination after parallel training; performing domain adaption; and estimating the confidence in the structured-prediction.

Based on joint work with Avihai Mejer, Francesco Orabona, Mark Dredze, Alex Kulesza and Fernando Pereira.

(S1B) The Yale-Technion Daniel Rose Center for Homeland Security [Abirim 2]

Separation and Protection Efficiency in Defense Strategy Against Terrorism

Gregory Levitin*, Kjell Hausken

We consider strategic defense and attack of a system which can be separated into independent identical homogeneous parallel elements. The defender distributes its resource between separation of the elements and their protection from terrorist attacks. The terrorist distributes its effort evenly among all attacked elements and has no preferable targerts. The vulnerability of each element is determined by a contest success function between the terrorist and the defender. The defender can choose a subset of the elements to protect. The terrorist does not know which

elements are protected and can choose a number of randomly chosen elements to attack. Separation efficiency conditions are formulated dependent on the defender's and terrorist's budgets, separation costs, contest intensity, and system demand. An algorithm for determining the optimal number of protected elements is suggested for the case when both the defender and the terrorist can choose the number of protected and attacked elements freely. We consider both the cases without and with performance redundancy.

Bi-Linear Inspection Game – Single Inspector vs. Single Manager Who Controls Multiple Inspectees
Yael Deutsch*, Boaz Golany, Uriel Rothblum

This talk addresses "games" between an inspector and a manager who controls multiple sites. The manager has to decide whether or not to cooperate in each of the sites and the inspector has to decide which sites to inspect in order to reveal cheating. The manager and the inspector have, respectively, a finite amount of a resource for executing their actions. Generalizing earlier work which the manager has a single unit of the resource, the current work considers situation where it has two units, where at most one unit can be spent at each site. Two models are considered: one has a consistency between the utility payoffs of the players and the second one does not. Expressions for all possible Nash equilibria are provided. In addition, the inspectors problem in a non-strategic environment is explored.

Civil Defense: A Military-Strategic Perspective
Itzhak Ravid*

There is a tendency to consider civil defense issues (e.g. threats of terror and ballistic missiles) similar to natural and industrial hazards, under the common umbrella of "Emergency Management". Traditional "Risk Analysis" methods are applied to those issues, with the main criterion of minimizing expected damage given a probabilistic array of threats. However, looking at civil defense as an integral part of war management leads to different objectives, alternatives, models and criteria.

(S1C) The Analysis and Optimization of Stochastic Systems [Main]

The Offered Load in Fork-Join Networks: Calculations and Applications to Service Engineering of Emergency Department
Itamar Zeid*

We develop tools to calculate the "Offered-Load" of fork-join networks; we propose mathematical models, analytical approximations and, finally, a simulation methodology for validation and extensions beyond the analytically tractable.

A fork-join network is a queueing network where customers are identified with a collection of tasks that can be processed sequentially and in parallel. The Offered-Load is a function that represents the average amount of work within the system, as a function of time. Fork-Join networks arise in many applications, and the one that motivated the present research are Emergency Departments (EDs): here patients fork into lab-tests, x-ray'ing and physician checkups which, ultimately, must of join in order for the medical process to continue.

We use our calculated Offered-Load of an ED to solve the capacity-sizing problem for various ED resources. Specifically, the Offered-Load, combined with the "square root rule", are shown to reduce congestion of rush hours and stabilize key operation measures over time.

Risk Sensitive Portfolio Optimization in a Semi-Markov Modulated Market
Anindya Goswami*

We address a portfolio optimization problem in a semi-Markov modulated market. We find the optimal portfolio selections by optimizing the risk sensitive criterion for both finite and

infinite time horizon. We use a probabilistic method to establish the existence and uniqueness of classical solution of the HJB equation for finite horizon problem. For infinite horizon problem we obtained the optimal growth rate for Markovian subclass in terms of a maximal eigenvalue of an appropriate matrix, using Perron-Frobenius theorem. A numerical procedure is also developed to compute the optimal expected terminal utility for finite horizon problem.

Reference: M. K. Ghosh, A. Goswami and Suresh K. Kumar. Portfolio optimization in a semi-Markov modulated market. Applied Mathematics and Optimization 60 (2009) 275–296.

On many-server limits of the parallel server model
Rami Atar*, Haya Kaspi, Nahum Shimkin

We study the multi-class parallel server model with reneging in the many-server fluid limit. Building on tools developed by Kaspi-Ramanan and Kang-Ramanan, we identify the behavior under fixed priority, in the form of measure-valued fluid equations. In the special case where reneging is exponential (but service time distribution is general), a consequence of our result is that a particular priority policy referred to as $c\mu/\theta$ (Atar-Giat-Shimkin) is asymptotically optimal for a natural cost function, at equilibrium.

(S2A) Production Systems [Main]

The Relationship between Yield and Flow Time in a Production System under Inspection
Israel Tirkel*, Gad Rabinowitz

This work studies the relationship between Yield and Flow Time (FT) in a production system monitored by in-line inspection. It originates in the known semiconductors Yield vs. FT trade-off premise, but can be adapted to other industries. We challenge the common premise, and suggest alternate analytical model to demonstrate this relationship. The model relies on a simplified production system that represents a repetitive segment in a production line. It illustrates that rising inspection rate increases both Yield and FT while exhibiting a trade-off. However with further growing inspection rate the Yield reaches a maximum and then starts to decline, while FT continues to increase. The Yield decline is explained by longer delay of inspection results which trigger the repair of an out-of-control machine. Clearly, lower Yield performance and higher FT are undesired. Our work defines this relationship with the analytical model and validates it with simulation. The model can be embedded in a decision support tool to pre-determine the inspection policy, while simultaneously considering Yield and FT.

Equipment Maintenance and Inspection Policies impact on Flow Time and Yield
Arava Ben Horin*, Israel Tirkel, Gad Rabinowitz

In order to maximize profitability of wafer fabrication, semiconductors industry constantly drives for low Flow-Time (FT) and high Yield.

Corrective maintenance is mostly triggered by in-line inspection which affects FT and Yield, and thus calls for studying both impacts simultaneously. Each can cause a positive effect on one measure, while causing a negative effect on the other. We suggest a model that provides an integrative evaluation in order to achieve the defined goal.

The inspection policy applied is Fixed Measure Rate (FMR), a common static policy. It is defined by sampling one lot every fixed number of lots that departed the first production step. The inspection time distribution is exponential. Corrective maintenance characteristics can be managed (e.g. preparation activities, technicians training), and can thus impact repair time duration. Applying repairs reduces availability and extends the machine's queue, which directly increases FT. They can also affect the metrology queue and extend the feedback time, which directly influences the Yield. We suggest a comprehensive simulation model that presents

the effects of changing in-line inspection and corrective maintenance characteristics combined, in order to drive the best operation decisions.

Simulation results show that applying repairs reduces machine availability and directly extends the FT; the longer queues formed due to maintenance and the growing variability as a result, reduces Yield. Statistical analysis of results shows that the measure cycle and the repair time length has an influence on both FT and Yield, while repair-time variance effects FT, and has no effect on Yield.

Cycle Time and Yield of a Multi-feeding Production Cell Under in-line Inspection

Elad Gati*, Israel Tirkel, Gad Rabinowitz

This research examines the relationship between yield and Cycle Time (CT) in a production system. It is based on past work done on this topic, including the last research that investigates the factors which affect both CT and Yield, by Tirkel and Rabinowitz (2010), and which focuses on the impact of inspection policies on these measures. In their work, a Production Cell model is established. The model is a short segment of a production line. It includes two production steps: Step i is named Feeder, and Step $i+1$ is named Bleeder. Following the Feeder and prior to the Bleeder, there is an inspection station, named Metro.

This study investigation is based on a Production Cell model with two Feeders. The goal here is to find the best inspection policy for two Feeders monitored by one Metro. The best inspection policy is the one that finds the optimal combination of maximizing the yield and minimizing the CT. The inspection policy is FMR (Fixed Measure Rate), which means that every n_i lots processed in Feeder i, a lot is sent to the Metro. This is an analytical-based research, validated by simulation. The research copes with the arrival distribution to the inspection station. It is assumed that each machine (Feeder) sends to the metro items every ni lots. The composed arrival process to the Metro is not a renewal process (unless $n_1 = n_2 = 1$). For any (n_1, n_2) the approach applied defines the process as CTMC (Continuous Time Markov Chain) and employs QBD (Quasi Birth and Death) process and MG (Matrix Geometric) to solve it. One challenge in this research is to find the sojourn time of lots in the metro, based on which we find the feedback delay (FD) distribution – the distribution of lots entering service of the Feeder.

(S2B) Algebraic and Geometric Methods for Discrete Optimization [Abirim 1]

Multiple via Single Criteria Optimization and Projections of Integer Hulls

Michal Rozenblit*

We consider multiple criteria optimization problems of the form $\max\{f(w_1x, \ldots, w_dx) : x \in S\}$, over sets $S \subset \mathbb{Z}^n$ of integer points, where d linear criteria are traded off by a convex function f. Using a combination of geometric and algebraic methods, we show that under some assumptions on the criteria w_i, such problems can be reduced to the solution of polynomially-many or even constantly-many single criteria counterparts over S. Applications include network flows, matroids, multiway tables and more. For instance, with binary w_i, for S the set of spanning trees in a graph, 96 greedily-solvable single criteria counterparts suffic to solve any tri-criteria problem regardless of the size of the graph, and for $3 \times 3 \times n$ tables with margin constraints, 7200 Graver-solvable single criteria counterparts suffice to solve any bi-criteria problem regardless of n and the margins.

Hilbert's Nullstellensatz and Linear Algebra: An Algorithm for Determining Combinatorial Infeasibility

Susan Margulies*

Unlike linear models, systems of multivariate polynomial equations over the complex numbers or finite fields can be compactly used to model combinatorial problems. In this way, a problem

is feasible (e.g. a graph is 3-colorable, Hamiltonian, etc.) if and only if a given system of polynomial equations has a solution. In the work of M. Laurent, J. Lasserre and P. Parrilo, Y. Nesterov, and others, continuous optimization problems which are modeled by zero-dimensional radical ideals have been shown to have a finite sequence of semidefinite programs that converge to an optimal solution. For yes/no combinatorial decision problems (e.g., is a graph 3-colorable?), we observed that Hilbert's Nullstellensatz gives a sequence of linear algebra problems that eventually determines feasibility. This has advantages as linear algebra is quite stable on computation and sparsity is well-understood.

In this talk, we present theoretical and experimental results on these sequences of large-scale, sparse, linear algebra relaxations to the combinatorial optimization problem. We show that the size of the smallest Nullstellensatz linear algebra system certifying that there is no stable set of size larger than the stability number of the graph grows as the stability number of the graph. We additionally describe ideas for optimizing the method, such as utilizing alternative forms of the Nullstellensatz, adding carefully-constructed polynomials to the system, branching and exploiting symmetry. Finally, in the case of 3-colorability, we use this method to successfully solve graph problem instances having thousands of nodes and tens of thousands of edges.

N-Fold Integer Programming in Cubic Time (Abraham Mehrez Award)

Raymond Hemmecke, Shmuel Onn, Lyubov Romanchuk*

N-fold integer programming is a fundamental problem with a variety of natural applications in operations research and statistics, including multiindex and multicommodity transportation problems, privacy and disclosure control in statistical databases, and stochastic integer programming. Moreover, it is universal and provides a new, variable-dimension, parametrization of all of integer programming. The fastest algorithm for n-fold integer programming predating our present work, runs in time $O(n^{g(A)}L)$ with L the binary length of the numerical part of the input and $g(A)$ the so-called Graver complexity of the bimatrix A defining the system. In this talk we provide a drastic improvement and establish an algorithm which runs in time $O(n^3L)$ having cubic dependency on n regardless of the bimatrix A. Our algorithm can be extended to separable convex piecewise affine objectives as well, and also to systems defined by bimatrices with variable entries. Moreover, it can be used to define a hierarchy of approximations for any integer programming problem.

(S2C) Stochastic Models [Abirim 2]

Universal control of search swarms

Iddo Eliazar*, Joseph Klafter

A general target-search model is considered: independent agents are stochastically emitted into a space with general structure and topology; the agents' paths in space are governed by a general random transport mechanism; the agents are detected only when in a target zone which is an arbitrary subset of space. We explore emission intensities that render the agents' "target-zone statistics" universal - i.e., invariant with respect to the underlying spatial topology and transport mechanism. The results established determine how to universally control the target-zone statistics via the emission intensity alone. The emergent "universal statistics" are shown to be governed by intrinsic power-laws, and include: the Beta, Pareto, Weibull and Fréchet extreme-value laws, the Central Limit Theorem's one-sided Lévy law, the linguistic laws of Zipf and Heaps, anomalous diffusion, and $1/f$ noise.

An insensitivity of reflection mappings on an orthant

Offer Kella*, S. Ramasubramanian

A reflection map on the nonnegative orthant (with constant routing matrix) is applied to two multidimensional processes differing by a constant vector. A question that has been open for over

15 years is under what conditions the difference between the two resulting processes converges to zero for any choice of the constant difference as time diverges. This in turn implies that if one imposes enough stochastic structure that ensures that the reflected process when the initial condition is zero converges in distribution then it will also converge in distribution when the initial condition is any almost surely finite random vector that may even depend on the original process (in contrast to Markov processes). We obtain a useful equivalent characterization of this property. As a result we are able to identify a natural sufficient condition which also gives a simple proof of stability when the driving process is (multi-dimensional) Levy or even just has stationary ergodic increments. A similar necessary condition is also indicated. Extensions of the sufficient condition are then developed for reflection maps with drift and routing coefficients that may be time and state dependent.

Tandem Queues with Batch Service
Iddo Eliazar, Shlomi Reuveni, Uri Yechiali*

Inspired by the biological process of gene translation and its modeling as a flow of particles through a tandem array of limited-capacity 'servers' - known in biophysics as the Totally Asymmetric Exclusion Process (TASEP), we devise and analyze a new model: the Totally Asymmetric Inclusion Process (TASIP). The TASIP is a system of tandem Markovian queues with batch service. That is, when service is completed at queue k all particles present at that queue move simultaneously to queue $k+1$. Equations for the time-dependent Probability Generating Function (PGF) of the system-state (occupancy) vector at arbitrary time instants, as well as at 'event epochs' (arrival or service completion), are derived. We show that, at steady state, the two PGFs coincide - implying a PASTA-type phenomenon. We explicitly compute the mean of the system-state vector at steady state, and establish a Little-type law. Furthermore, we explicitly compute the PGF of the TASIP's total load in steady state, and obtain a stochastic decomposition result for the total load. Finally, we show that TASIP models with identical service rates are, under various perspectives, optimal.

(S3A) Topics in Operations Research 1 [Abirim 1]

Stochastic Enumeration Method for Counting
Reuven Rubinstein*, Radislav Vaisman

We present a new generic algorithm, called stochastic enumeration (SE) for counting $\#P$ problems. A particular emphasis is placed on self-avoiding walks (SAW's) and satisfiability (SAT). We show that

1. An importance ingredient while counting with SE is the classic DPLL solver.
2. SE can be viewed as a sequential importance sampling (SIS) algorithm.
3. SE presents a natural generalization of a specially designed splitting algorithm, which is presented in the paper as well. Our simulation studies show its superiority to splitting.

Location Reliability Models: New Directions and Insights
Oded Berman*, Dmitry Krass

Over the last few years there has been a growing interest in developing location reliability models in order to understand the impact of imperfect reliability on the optimal location patterns and the resulting costs. In this talk we will review some of the recent results for this class of models. We will focus on models where failures are due to some stochastic events, rather than actions of an intelligent adversary (the latter are the subject of location interdiction models – a distinct stream of literature).

We distinguish between two types of models with regard to the information available. In one type it is assumed that the state of each facility is known to the customer in advance, hence it is called the Median (or Center) Problem with Unreliable Facilities and Complete Information (in the median problem we minimizing the expected cost of the system while in the center problem we minimize the maximum cost traveled). In the second type it is assumed that customers do not have advance knowledge of the state of the facilities and thus have to search for an operating facility, hence called the Median (or Center) Problem with Unreliable Facilities and Complete Information.

Optimization of Real-Time Systems (Simulation and Analytic Approach)
Joseph Kreimer*, Edward Ianovsky

We consider a multiserver Real-Time system (RTS) with identical servers (e.g. unmanned aerial vehicles, machine controllers, etc.) that provide service for requests of real-time jobs arriving via several different channels (e.g. surveillance regions, assembly lines, etc.) required to be under nonstop surveillance (working under maximum load regime).

We compute analytically (for *exponentially* distributed service/operation times and *generally* distributed maintenance times) and via simulation using Cross-Entropy method (for *generally* distributed service and maintenance times) optimal routing/assignment probabilities which maximize system availability and minimize the damage. We show that a proper choice of routing probabilities improves system's availability more than arbitrary increasing of the number of servers.

The Value of Accurate Inventory Information in Supply Chain Management
Avrahami Assaf, Avinoam Tzimerman*, Yale T. Herer, Avraham Shtub

Uncertainty in a supply chain is often manifested as inventory discrepancies. In particular, we are concerned with discrepancies known as inventory errors, which are the result of events that can be classified as one of the following three types: *Shrinkage*, *Misplacement* and *Wrong-Scanning*. The outcome of these events is discrepancies between *IT inventory* (inventory as it appears in a computerized system) and *available inventory* for sale. The literature describes both the extent to which these types of discrepancies are commonplace and the degree to which they cause increased lost sales and lower profits. Nevertheless, these discrepancies and/or their associated negative affects can be eliminated if more information is available.

Our work explores three different models that operationally take these errors into consideration in three different ways: the *ignorant case*, the *informed case* and the *full information case*. The objective of our work is to explore and quantify the value of the additional information in the *full information case* in comparison to the *ignorant case* and the *informed case*. We develop an analytical model for each case, where the objective in all three is to minimize the average expected cost per period. We find the optimal order-up-to level for each period. In addition, we find the optimal frequency for performing inventory counts for the *ignorant case* and the *informed case*. There is, of course, no need to conduct an inventory count under *full information* case.

Our findings indicate that additional information greatly influences the optimal policies and it can be exploited wisely to save a great deal of money. Moreover, we find that certain problem parameters do not affect the *value of information* (e.g., the expected value of the demand), whereas other parameters influence it tremendously (e.g., the expected value of the mean shrinkage).

(S3B) Game Theory [Abirim 2]

A Stochastic Competitive R&D Race where "Winner Takes All": Explicit Computation of a Unique Nash Equilibrium

Boaz Golany, Pelin Gulsah Canbolat*, Inbal Mund, Uriel G. Rothblum

The paper considers an environment in which multiple firms compete over the development of a product. The first firm to complete the project gains a reward whereas the other firms gain nothing. Each firm is to decide on the level of investment in developing the project; the completion time of the project by a firm is random having exponential distribution whose rate depends linearly on the investment level. The paper provides a method for explicitly computing a unique Nash equilibrium, parametrically in the interest rate. The structure of the solution yields insights about the behavior of the participants. Further, an explicit expression for a unique globally optimal solution is obtained and it is compared with the unique Nash equilibrium.

Explicit Pricing Strategies in Competitive Market with Reference Effect
Eyal Dayan, Arieh Gavious*, Oded Lowengart

Pricing decisions are regular activities of managers in most firms. These tasks are somewhat complex as they are require to handle the supply side of the firm (i.e., cost), the demand side (i.e., consumers), and the type of competitive environment that the firm operates within. Obtaining dynamical optimal pricing decisions, therefore, is even more complex as adjustment over time to these three variables is needed. This is especially valid where some decisions the firm is making will affect future outcomes of that variable. For example, changing the price of the product at one pint of time might create certain expectations among consumers in future purchase occasions.

It is intuitive to describe the idea of reference price, therefore, as a price (p) to which consumers are comparing their expected price on the next purchase occasion that serves as a reference point (RP). Competition between firms can primarily be analyzed based on two dimensions. The first is competition that is based on the quantity in the market between the different firms and the second is based on price. Extensive research gas been done in exploring competitive environments under these two approaches. In the recent past dynamic models that capture cost, reference price, and competition has been introduced in the literature.

The purpose of this paper, therefore, is to fill the void in addressing this issue of price competition by developing a differential game model for such competition using a Bertrand mechanism that capture this process. Our results indicate that a dynamic optimal pricing strategy under price competition complexity can be obtained and that we account for the asymmetric effects of reference price. We also analyzed a case of a myopic firm (i.e., instantaneous optimization) and show that optimal pricing strategy under this assumption can be obtained for this type of competition and with reference price effects in the demand formation process. Furthermore, we show that this alternative pricing strategy will result in lower profits than the dynamical optimal one.

Sequential All-Pay Auctions with Head Starts
Ella Segev*, Aner Sela

We study a sequential all-pay auction where heterogeneous contestants are privately informed about a parameter (ability) that affects their cost of effort. In the case of two contestants, contestant 1 (the first mover) makes an effort in the first period, while contestant 2 (the second mover) observes the effort of contestant 1 and then makes an effort in the second period. Contestant 2 wins the contest if his effort is larger than or equal to the effort of contestant 1; otherwise, contestant 1 wins. This model is then generalized to any number of contestants where in each period of the contest, $1 \leq j \leq n$, a new contestant joins and chooses an effort. Contestant j observes the efforts of all contestants in the previous periods and then makes an effort in period j: He wins if his effort is larger than or equal to the efforts of all the contestants in the previous periods and strictly larger than the efforts of all the contestants in the following

periods. We characterize the unique sub-game perfect equilibrium of these sequential all-pay auctions and analyze the use of head starts to improve the contestants performances.

On the balancedness of homogeneous of degree one games with applications to queues with outsourcing
Shoshana Anily*, Moshe Haviv

A cooperative game with transferable utility is defined as homogeneous of degree one if m times cloning all players at any given coalition leads to m times the original value of the coalition for any integer $m \geq 1$. We show that this property coupled with the game being sub-additive guarantee the non-emptyness of the core. We then present a few examples for such games which naturally emerge when servers in queueing systems cooperate in order to outsource some of the activities, while the rest of the activities is optimally assigned to the servers in the system.

(S3C) Transportation [Abirim 3]

The Locomotive Fleet Fueling Problem
Tal Raviv*, Mor Kaspi

Operators of large fleets face the problem of deciding upon the location and schedule of fueling operations. This operation is associated with three types of costs. Namely, the direct cost of fuel, the cost of the fueling procedure and medium-term contracting costs paid to the suppliers of fuel at the various locations. These costs may vary significantly among locations due to the distance to the nearest refinery, local taxes and other factors. Clearly, the operators seek a contracting plan and fueling schedule that minimize the total costs.

In this study, the problem is modeled as a mixed integer program. The formulation is strengthened using several classes of valid inequalities and domination rules. The enhanced formulation is capable of delivering near optimal solutions of large-scale instances using a commercial solver.

An instance of the problem that is based on the schedule of CSX (a major USA railway company) was introduced as a challenge for the 2010 Problem Solving Competition of INFORMS Railway Application Section (RAS). See http://www.informs.org/Community/RAS/Problem-Solving-Competition. The solution submitted by the authors won the competition.

Evaluating the efficiency of local municipalities' traffic safety programs via Data Envelopment Analysis
Doron Alper*, Zilla Sinuany-Stern, David Shinar

The purpose of this study was to estimate the relative efficiency of 197 local municipalities in traffic safety in Israel during 2004-2009, using Data Envelopment Analysis (DEA). DEA efficiency is based on multiple inputs and multiple outputs, when their weights are un-known. We used here 2 inputs, 6 main outputs, and 8 intermediate variables (which were used sometimes as inputs and sometimes as outputs). The inputs reflect the resources allocated to the local municipalities by Israel's National Road Safety Authority (NRSA) (such as funding). The outputs include measures that reflect reductions in accidents (such as accidents per population), and intermediate variables are safety performance indicators (SPI): measures that are theoretically linked to crash reductions (such as use of safety belts). Using DEA, the local municipalities were ranked from the most efficient to the least efficient. Benchmark analysis was performed, where the best local municipalities were identified, and required improvements for the other inefficient municipalities were calculated. We found that most of the improvements were required in two intermediate variables related to traffic tickets. Further analyses utilizing multiple regression and factor analysis were preformed to verify the effect of various environmental parameters on the efficiency of the municipality. The environmental parameters, tested for each local municipality were related to the size, age, and socio-economic level of the population. These three

groups of variables were statistically significant with the DEA efficiency score of the main model in our research. As far as we could determine, this is the first time that DEA is used to measure the efficiency of local municipalities in improving traffic safety.

Benchmarking airports from a managerial perspective
Nicole Adler, Vanessa Liebert, Ekaterina Yazhemsky*

Benchmarking airports is currently popular both in the academic literature and in practice but has proved rather problematic due to the heterogeneity inherent in any reasonably sized dataset. Most studies either treat the airport production technology as a black box or separate terminal and airside activities, assessing them individually. In this paper we analyze airports as a single unit due to the direct complementarities, avoiding the artificial separation of inputs between the terminal and airside but opening the black box through network data envelopment analysis (DEA). To further improve the benchmarking process, we identify appropriate peers for 43 European airports over 10 years through a dynamic clustering mechanism according to pre-defined characteristics and restrict the integer linear program with respect to potential reductions in capital inputs. Compared to basic DEA models, the results of the network DEA structure provide more meaningful benchmarks with comparable peer units and target values that are achievable in the medium term. By identifying each airports individual reference set, unique airport outliers influence relative efficiency less severely than occurs under basic DEA. In addition, the formulation is shown to be suitable in assessing different strategies with respect to aeronautical and commercial activities not only separately but also in combination.

New Industry Formation for Sustainable Transportation
Michael Naor*, Cheryl Druehl, Ednilson Bernardes, Yoram Shiftan, Shalom Hakkert

The environmental catastrophe in the Gulf of Mexico caused by the leakage from the oil well explosion of the Deepwater Horizon rig underlined the necessity to cut dependency on oil and protect the environment. Electrification of the automobile industry has the potential to accomplish these goals by diminishing both oil consumption and air pollution. Thus, this research proposes to investigate, using case study methodology, a new formation of sustainable transportation in the automobile industry. The subject of this case study is an Israeli company specialized in building infrastructure for electric vehicle (EV). The significance of the study stems from the unique opportunity to conduct an experiment on a country wide scale in Israel specifically. Israel has been globally recognized as a start-up nation and has the technological capabilities and engineering knowledge to build such an infrastructure. Israel Ministry of Environmental Protection has classified automotive pollution as a severe problem. Furthermore, Israel has strategic and political reasons to cut its dependency on oil making it a national priority. The scarcity of natural resources at Israel such as water and oil play a central role in driving Israeli companies to investment in R&D seeking for alternative sources of energy. The location is certainly a critical factor in this case. According to the International Road Federation, World Road Statistics (2007), Israel has 122 vehicles per KM of paved road. This ranks Israel very highly in road congestion, alongside with only a few other countries in the world such as Honk Kong and Monaco. Israel's small size, the concentration of its population in only a few districts, and its highly congested roads provide a rare arena for investigation. Another important factor making the investigation of the case study for adoption unique specifically in Israel is the high taxation on vehicle propelled by combustion engine in Israel of up to 95%, compared to a current 10% tax purchase for EVs. Another innovative aspect examined is the notion of separating the vehicle from the battery to make this green solution economically feasible by reducing the vehicle's price and extending the driving range. Finally, the study describes the opportunities emerging from this and it sketches new avenues for future research.

(M1A) Military OR [Abirim 1]

Mathematical Models of Insurgencies

Moshe Kress*

Most recent military conflicts (e.g., Iraq, Afghanistan, Libya) differ from traditional force-on-force engagements, taking the form of an asymmetrical counterinsurgency situation where the civilian population plays a major active role. Legacy warfare and combat models (e.g., Lanchester models) do not capture this "social" aspect of military conflicts and therefore must be modified. In this talk I will review some recent research in this area relating to dynamic and rational choice models.

Multi-Criteria/Multi-Domain Optimization and Target Cascading for System Level Design

David V. Strimling*

Product development is a major undertaking involving numerous decisions. These decisions cannot be made in isolation, but rather must be made in the context of the total system being developed. Multi-Criteria/Multi-Domain Optimization and Target Cascading provides a conceptual framework within which to formulate and solve complex design problems. The multiple criteria include performance attributes and burdens. The multiple domains are the disparate subsystems of the total system. The problem must be solved in the context of three types of constraints: requirements constraints, program constraints, and combinatorial constraints. Requirements constraints are those related to desired and/or required levels of the performance and integration criteria. Program constraints are those related to cost, risk, etc criteria. Combinatorial constraints address physical and technical issues regarding feasible combinations of subsystem options. Two different procedures for solving this problem are a preference oriented method that effectively collapses the multiple criteria into one aggregated criterion, and a generating method that generates a Pareto set of solutions. The multi-criteria/multi-domain optimization methodology was implemented at General Dynamics Land Systems (GDLS) in their proprietary Armored Combat System Optimization Modeler (ACSOM). ACSOM is designed to function on a standard desktop PC. The vision of ACSOM is to quickly span the entire design space considering the full spectrum of subsystem options and generate a Pareto set of balanced design concepts to aid design engineers and decision makers in selecting the preferred balance among performance and integration requirements, and program burdens. This preferred balanced solution focuses the design engineer's attention on optimizing the design of the selected subsystem option. Results of this optimization can be sent back to ACSOM for further balancing of the system level performance and burdens. This sequential process of balance and result feedback, called target cascading, can be continued until the system level design is deemed best by the customer. The multi-criteria/multi-domain optimization and target cascading process for system level design can be used throughout the product development lifecycle.

Traveling Salesman Problem in Dynamic Environment

Eran Simhon*, Liron Yedidsion, Shraga Shoval

In this research we are studying dynamic versions of traveling salesman problem (TSP) in particular we are interested in two dynamic elements, the first one is different arriving time, each city (targets in our case) has it own arrival time r_j and the second element is that the targets are not stationary and can move at fixed velocity.

Algorithms for Dynamic TSP can be used in real applications, mainly for optimization of defense missions and robotic tasks. For example, airplane that needs to supply goods to moving ground forces, Maritime border control ship that need to encounter with other ships, automated shut gun that need to hit multi targets or robot in production floor that need to encounter parts

in motion. We believe that as the use of automated and autonomous systems in battle field as in production floors is increasing, the need for such algorithms will increase respectively.

Since TSP is NP-hard and therefore any generalization of it is NP-hard as well, we will look at special cases of Dynamic TSP, in which both pursuer and targets are confided to one dimensional plane. In this paper we define two special cases and suggest polynomial algorithm for finding optimal solutions for those problems. The problems are stated as follow: Given a set $S = \{s_1, \ldots, s_n\}$ of targets, each s_i show up in time r_i at position p_i on one-dimensional plane, and an agent starting at the origin (P_a) and having maximum speed $v_a = 1$, find the fastest tour starting at the origin, which intercepts all objects, in the first problem targets speed is zero and in the second problem target speed is $0 < v < 1$ (all targets have the same speed). To simplify the problems we are making the following assumptions: there are no collisions between objects, there are no constraints on agent motion (it can change its direction and speed instantaneously).

(M1B) Stochastic Queueing/Inventory Systems [Main]

State Dependent Priorities for Service Systems with Exchangeable Items

O.J. Boxma, David Perry, Rachel Ravid*

We consider a repair facility which consist of one server that arrive from two sources. The service times are independent, exponential random variables with equal rates. The failed repairable items arrive according to independent Poisson processes. The items are admitted to a unique line but the customers wait and are marked according to their sources. We assume that the system is in steady state. A backorder is created for each arrival. At a completion of service the item is delivered to a customer from the base which currently has the most backorders. Ties are broken by random assignment.

Our aim is to compute certain steady state performance measures of the system. Underlying these performance measures is the joint steady state distribution of the backorders.

A Fluid EOQ Model of Perishable Items with Intermittent High and Low Demand Rates

Onno Boxma, David Perry, Shelley Zacks*

We consider a stochastic fluid EOQ-type model with demand rates operating in a two-state random environment. This environment alternates between exponentially distributed periods of high demand and generally distributed periods of low demand. The inventory level starts at some level q, and decreases linearly at rate β_H during the periods of high demand, and at rate $\beta_L < \beta_H$ at periods of low demand. The inventory level is refilled to level q when level 0 is hit or when an expiration date is reached, whichever comes first. We determine the steady-state distribution of the inventory level, as well as other quantities of interest like the distribution of the time between successive refills. Finally, for a given cost/revenue structure, we determine the long-run average profit, and we consider the problem of choosing q such that the profit is optimized.

Accounting for Exact Temporal Behavior of Costs in Stochastic Inventory Systems

Tal Avinadav*, Mordecai I. Henig

Most periodic review inventory models do not attribute holding and shortage costs to their actual timing and therefore may miss optimality. This work shows how it is possible to overcome this deficiency in a general periodic review model. Exact formulas are presented to compute the optimal order quantity when inventory costs accrue in discrete or continuous times. The computational effort, not much different than that of the common periodic review models, involves some extra arithmetic. We introduce formulas to calculate the optimal policy for the Lévy process and especially for the Poisson, Wiener and gamma processes. The one for the

Poisson process can be solved analytically while those for the Wiener and gamma processes can be approximated to the desired accuracy.

(M1C) Discrete Optimization [Abirim 2]

Graph Coalition Structure Generation

Thomas Voice, Maria Polukarov*, and Nicholas R. Jennings

We give the first analysis of the computational complexity of coalition structure generation over graphs. Given an undirected graph $G = (N, E)$ and a valuation function $v : \mathcal{P}(N) \to \mathbb{R}$ over the subsets of nodes, the problem is to find a partition of N into connected subsets, that maximises the sum of the components values. This problem is generally NP-complete; in particular, it is hard for a defined class of valuation functions which are *independent of disconnected members*that is, two nodes have no effect on each others marginal contribution to their vertex separator. Nonetheless, for all such functions we provide bounds on the complexity of coalition structure generation over general and minor free graphs. Our proof is constructive and yields algorithms for solving corresponding instances of the problem. Furthermore, we derive polynomial time bounds for graphs of bounded treewidth. However, as we show, the problem remains NP-complete for planar graphs, and hence, for any K_k minor free graphs where $k \geq 5$. Moreover, a 3-SAT problem with m clauses can be represented by a coalition structure generation problem over a planar graph with $O(m^2)$ nodes. Importantly, our hardness result holds for a particular subclass of valuation functions, termed *edge sum*, where the value of each subset of nodes is simply determined by the sum of given weights of the edges in the induced subgraph.

General Techniques for Accelerating FPTAS for the routing and knapsack problems

Amir Elalouf, Eugene Levner*

A fundamental problem in quality-of-service (QoS) routing in communication networks is the constrained path problem (CPP) where a path connecting a source node to a destination node is to satisfy QoS constraints, such as cost, delay and/or reliability. The network is modeled by a directed graph where the vertices represent computers or routers and edges represent links. To model QoS parameters, each edge is associated with an edge cost (or benefit) and a weight representing delay or reliability of the edge. In some cases the objective is to minimize a criterion (say, total costs) while in other cases to maximize some function (e.g. total benefit). Regarding the problem constraint, in some cases the considered constraints are of type "less than or equal" (for example, "path's transition time is to be less than or equal to a given value"); whereas another constraint type may be "larger than or equal" (for example, "the information gathered by a mobile agent during its travel through web sites is to be larger than or equal to a given threshold"). We study ϵ-approximation algorithms (FPTAS) for solving the considered problems. We suggest a general three-stage technique that follows and extends an earlier computational scheme suggested by Gens and Levner (DAM, 1981) for the min-cost knapsack problem. Its basic building blocks - testing and scaling - are the same as in Hassin's algorithm (1992), but they are used differently. The new techniques permit to essentially improve several known algorithms (Gens and Levner (DAM, 1981); Hassin (MOR, 1992); Camponagaro and Shima (J.UCS 2010)), and to obtain several new FPTAS.

A Sequential Competitive Resource Allocation Game

Rafael Hassin, Barak Y. Reif*

We define a competitive resource allocation game in which there are two players. An attacker and a defender. The game is played sequentially in turns. In each turn the attacker commits one resource to some territory, slot, and in turn the defender responds by committing one of its

own resource units to that slot. The resource units differ from each other by the subset of slots they can be placed at. If the defender has no available resource to responds to the attacker, he looses. If the defender can match resources to each of the attackers, he wins the game. We wish to find winning strategies and the necessary and sufficient conditions for either players win. We also observe some of the games properties.

Later we expand the above model, instead of an asymmetrical game of an attacker and defender we look into a game where both players attack simultaneously and defend simultaneously. The game is still played sequentially, and in turn both players commit resources and both players respond to their respective opponents play. We inquire as to which properties carry on from the original game to the simultaneous one, and observe some probabilistic attributes resulting from this expansion.

(M2A) Game Theory [Abirim 1]

Contests with a Random Number of Players and Private Values

Moshe Haviv, Liron Ravner*

The widely used Tullock contest model has been researched under assumptions of both complete information and incomplete information regarding one of the model's variables, the number of participants or the private values of each of the participants. In this presentation we consider contests with incomplete information regarding both the number of players participating in the contest and their valuation of the prize. We seek the symmetric Bayesian-Nash equilibrium points and the centralized socially optimal strategies for the various possible model assumptions. In addition we examine the properties of the equilibrium strategies and test their sensitivity to the model's parameters. Furthermore, we discuss the incomplete information model for the exponential race, where the winner is the first participant to complete a task which requires an exponentially distributed amount of time. The exponential race is similar both in concept and in analysis to the Tullock contest.

Reverse Game Theory in Case Evaluation with Differential Information

Ram Orzach*, Stephen J. Spurr

This paper provides an example showing the bene.t of mechanism design in a nonbinding arbitration procedure called case evaluation that is widely employed in U.S courts. Under the current system, a party who rejects the mediation award is penalized, unless the trial verdict is more favorable to her than the mediation award. This penalty is designed to minimize the frequency of trial, by inducing both parties to accept the award. We provide procedures that motivate the parties to disclose their private information to the mediator. In the example, under the proposed new rules of game the mediation award is likely to be more accurate, and the parties are more likely to accept it, thereby reducing the frequency of trial, while providing an ex ante gain for both parties.

(M2B) Workforce Management [Main]

Applying Constraint Programming for the Identification and Assignment of Service Professionals

Sigal Assaf, Yael Ben Haim, Daniel P. Connors, Haggai Eran, Donna L. Gresh, Oded Margalit*, Michael J. Mcinnis, Julio Ortega, Yossi Richter

We solve the challenge of assigning highly-skilled practitioners (people) to high-end positions (jobs) while taking into account numerous constraints. Unlike existing solutions, which focus primarily on scheduling many similar workers on shift assignments, our case is complicated by the fact that every professional and every role are unique.

We model the situation as a constraint satisfaction problem (CSP). The rich expressive language supported by CSP provides a convenient mechanism for changing and adding new matching and preference constraints. We use the CSP solver designed in IBM Research - Haifa to provide a solution. The solver was enhanced with the innovative SomeDifferent propagator to handle this problem.

The new tool was successfully piloted in 2005, and has since been essential in helping resource managers and deployment managers implement better assignments. Today it is widely deployed by IBMs Global Business Services (GBS) in all geographies.

Manpower allocation for service jobs: The Balayla' Model
Isaac Balayla*

Time standards are required for efficient manpower allocation. The process of determining standards for service jobs, create inherent difficulties because of: a) the wide variation in TBA and in service performance time (particularly in face-to-face jobs); b) the numerous ancillary tasks & interruptions.

The Balayla' model overcomes the above-mentioned difficulties by deploying a series of indicators for a correlation & BEP between output (of worker) and waiting time (of customer). The indicator values are affected by service level of urgency and the initial number of service workers.

The model was applied successfully on various types of service systems. such as tellers in supermarkets and medical doctors.

(M2C) Water Management [Abirim 2]

Assessment of the intake and the pretreatment design in existing SWRO (SeaWater Reverse Osmosis) plants by Hasse Diagram Technique (HDT)
Amos Bick*, Rainer Bruggemann, Gideon Oron

Seawater desalination is a rising technology to overcome water scarcity and is increasingly being used for the production of drinking water. A good quality pretreatment process is prominent to the successful operation of a seawater reverse osmosis (SWRO) plant. The tasks of the intake and the pretreatment modules are to prevent fouling and to extend the lifetime of the Reverse Osmosis (RO) membranes.

To practice the method of assessment design, 19 seawater plants with 11 different design strategies have been evaluated. The assessed strategies consist of: (i) Intake treatment, (ii) Pretreatment aid (Coagulation, Flocculation, Clarification), and (iii) Pretreatment unit. Every possible combination of the design elements has been evaluated by a set of 15 attributes. The comparative evaluation has been done by Hasse Diagram Technique (HDT) using PyHasse software.

The result of the HDT application includes 7 favorable incomparable design strategies. In order to solve the incomparability, the Condorcet-Kemeny-Young-Levenglick (C-K-Y-L) ranking procedure was additionally used and outputted the optimal design that is based on subsurface intake and membrane (Ultrafiltration) pretreatment facilities. According to the results, both mathematical methods, i.e. the HDT in combination with the C-K-Y-L, Procedure could be powerful and helpful tools to assist decisions concerning design strategies.

Optimal multi-year management of a water supply system under uncertainty: robust counterpart approach
Mashor Housh*, Avi Ostfeld, Uri Shamir

In this paper the Robust Optimization (RO) methodology [Ben-Tal et al., 2009] is applied to optimize management of a Water Supply System (WSS) fed from aquifers and desalination

plants. The water is conveyed through a network to meet desired consumptions, where the aquifers recharges are uncertain. The objective is to minimize the total cost of multiyear operation, satisfying operational and physical constraints. The RO methodology uses a min-max approach to solve the problem assuming that the uncertain parameters can only reside within a user-defined uncertainty set. This static version of RO is called Robust Counterpart (RC), in which the original problem is converted into a deterministic equivalent problem. The robust policy obtained by the RC approach is compared with polices obtained by other decision making approaches.

(M3A) Optimization and Analytics for Water Management [Main]

Automated Leakage Identification in Water Distribution Networks using Optimization

Ofer M. Shir*, Segev Wasserkrug, Yossi Shiloach, Lena Granovsky

With growing populations, increasing industrial demands, and changes in climate patterns, water loss throughout distribution networks constitutes a significant issue worldwide with prominent environmental as well as financial consequences. In most cases, water loss is primarily attributed to leakage from the water distribution networks' pipelines. Treating these leaks firstly pose the leakage identification challenge - that is, detecting the occurrence of leaks and finding their exact locations - a problem which is known to be hard. At present times, solving these problems primarily relies on either customer complaints, or disruptive procedures such as service shut-down to parts of the network.

The advent of modern technology, and especially the availability of remotely connected and reliable sensing instrumentation, makes field measurements accessible at high rates and accuracy. Consequently, in many of today's water networks, sensors are widely available. These sensors can provide real-time hydraulic information, e.g., water pressure and the water flow-rates at various points in the network. In this work, we introduce a novel framework that considers such measurement data and incorporates it into an optimization scheme in order to solve the leakage identification problem within a water distribution network in a non-invasive, cost-effective, and proactive manner. The proposed scheme compares the values provided by the sensor networks to values estimated by a hydraulic model, and utilizes an optimization engine, which searches for the location of the leaking sites that best explain the discrepancy between the measured and estimated values. Our presentation will describe the algorithmic approach, the underlying details of solving a hydraulic network, and outline experimental results, based on real-world water distribution scenarios, corroborating our proposed scheme. Our contribution therefore comprises a novel leakage detection technique, utilizing optimization, as well as results demonstrating its successful application.

Minimizing the Production Costs of Providing Water

Yulia Tseytlin*, Segev Wasserkrug, Lena Granovsky, Mila Keren

Providing water from multiple water sources, while attempting to minimize overall production costs is one of the major pain points of water providers worldwide. This is a complex problem, as the water sources may have differing qualities and differing production costs. In addition, in many cases, the main production cost is the cost of electricity required to pump the water. As water can be stored in high elevation storage areas such as reservoirs, it is possible to reduce the electricity cost by pumping water at hours with low electricity tariff, and then provide the water from the reservoirs at virtually no additional electricity costs.

In this talk, we shall focus on an optimization model that minimizes the production costs, and which is based on the needs of an actual water provider. We shall detail the relevant objectives, control points, and constraints, and present the formal definition of the optimization model.

Water Analytics
Einat Kermany*, Yehuda Naveh, Pnina Vortman, Dan Pelleg

As part of the smarter cities initiative, we are developing analytical tools to support management of water resources. The current emphasis is on consumption monitoring at the municipal level. We are using statistics and machine learning methods to detect irregular consumption (low or high consumption) in order to find leaks, water theft or faulty meters. In the talk I will show some examples of information that could be derived from water consumption data and will discuss our methods for irregular consumption detection.

(M3B) Scheduling [Abirim 1]

Online Scheduling of Two Job Types on a Set of Multipurpose Machines with Unit Processing Times
Dvir Shabtay, Shlomo Karhi*

We study a problem of scheduling a set of n jobs with unit processing times on a set of m multipurpose machines in which the objective is to minimize the makespan. It is assumed that there are two different job types, where each job type can be processed on a unique subset of machines. We provide an optimal offline algorithm to solve the problem in constant time and an online algorithm with a competitive ratio that equals the lower bound. We show that the worst competitive ratio is obtained for an inclusive job-machine structure in which the first job type can be processed on any of the m machines while the second job type can be processed only on a subset of $m/2$ machines. Moreover, we show that our online algorithm is 1-competitive if the machines are not flexible, i.e., each machine can process only a single job type.

Two-Machine Flow-Shop Scheduling with Rejection
Dvir Shabtay*, Nufar Gasper

We study a scheduling problem with rejection on a set of two machines in a flow-shop scheduling system. We evaluate the quality of a solution by two criteria: the first is the makespan and the second is the total rejection cost. We present four different problem variations to deal with these two criteria (P1-P4). In problem variation P1, the objective is to find a schedule which minimizes the sum of the makespan and total rejection cost. In problem variation P2, the objective is to find a schedule which minimizes the makespan subject to an upper bound on the total rejection cost. In problem variation P3, the objective is to minimize the total rejection cost subject to an upper bound on the value of the makespan. Finally, the objective in problem variation P4 is to find the entire set of Pareto optimal points. We prove that the P1 problem variation is NP-hard and for its solution we provide two dierent approximation algorithms, a pseudo-polynomial time optimization algorithm and a fully polynomial time approximation scheme (FPTAS). We also provide a pseudo-polynomial time optimization algorithm to solve P2-P4 problem variations (those problem variations on a set of two machines in a flow-shop system are NP-hard due to the NP-hardness of the same problem variations on a single machine (Zhang (2009))).

A two-stage flow-shop scheduling with a critical machine and batch availability
Enrique Gerstl, Gur Mosheiov*

We study a two-stage flow-shop, where each job is processed on the first (critical) machine, and then continues to one of two second-stage (dedicated) machines. Jobs are processed on the critical machine in batches, and a setup time is required when starting a new batch. The setting assumes batch-availability, i.e., jobs become available for the second stage only when their entire batch is completed on the critical machine. We consider three objective functions: minimum makespan, minimum total load, and minimum weighted flow-time. Polynomial time

dynamic programming algorithms are introduced, which are numerically shown to be able to solve problems of medium size in reasonable time. A heuristic for makespan minimization is presented and shown numerically to be both accurate and efficient.

Scheduling problems with two competing agents to minimize minmax and minsum earliness measures
Baruch Mor*, Gur Mosheiov

A relatively new class of scheduling problems consists of multiple agents who compete on the use of a common processor. We focus in this paper on a two-agent setting. Each of the agents has a set of jobs to be processed on the same processor, and each of the agents wants to minimize a measure which depends on the completion times of its own jobs. The goal is to schedule the jobs such that the combined schedule performs well with respect to the measures of both agents. We consider measures of minmax and minsum earliness. Specifically, we focus on minimizing maximum earliness cost or total (weighted) earliness cost of one agent, subject to an upper bound on the maximum earliness cost of the other agent. We introduce a polynomial time solution for the minmax problem, and prove NP-hardness for the weighted minsum case. The unweighted minsum problem is shown to have a polynomial time solution.

(M3C) Decision Making in Queues [Abirim 2]

Dynamic Control in the Single Server Markovian Queue
Gail Gilboa-Freedman*, Refael Hassin

We explore a queuing system with a non-homogeneous population of customers in the sense of optimizing the social welfare. In particular we modify Naor's observable $M = M = 1$ queue model, which is the most seminal queueing model involving strategic customers. We assume that some customers obey the central planner, who aim to maximize the social welfare, the others don't. We consider two levels of control:

1. Static control, in which disciplined customers join or balk as instructed by the central planner.

2. Dynamic control, in which disciplined customers also abandon the queue when instructed.

We quantify and compare the benefit from each level of control. In particular, we investigate the additional benefit from having dynamic control instead of static control.

On the equilibrium in a two dimensional strategic game
Refael Hassin, Ricky Roet-Green*

In most models of congestion one tries to avoid others, and hence responds inversely to their actions. For example, if more individuals join a queue then this tends to discourage the individual from joining. This behavior is called avoid-the-crowd (ATC). The opposite behavior of follow-the-crowd (FTC) is also common. For example, the more customers buy priority in a queueing system, the more is an individual inclined to follow them and buy priority for himself. FTC behavior typically results in multiple equilibria, whereas ATC provides a unique equilibria.

This discussion is limited to the case where the customer decision is of one dimension. For example, Naor (1969) considered an observable queue, where customers have to make a decision if to join it or not. The decision is based on the length of the queue, and the solution is a threshold strategy. Adding another action as a possibility for the deciding customer increases the model's complexity, and requires a new definition of the ATC and FTC concepts. Our model suggests a case of customer's two dimensional decision problem, and offers a proof of existence and uniqueness of a symmetric Nash equilibrium.

When to arrive to a queue with tardiness costs?
Moshe Haviv*

A random number of customers seek service during some time interval which starts at time zero and can be bounded or not. They face the question of when to arrive so as to minimize their waiting and/or tardiness costs. We look for a Nash equilibrium time of arrival strategy as well as for the the social optimal arrival strategy. Non linear differential equations, which call for numerical solutions, are first suggested. Then, exact solutions for the corresponding fluid approximation models are derived.

A queueing approach to a multi class $M/G/1$ make-to-stock with backlog
Opher Baron, Yoav Kerner*

A single machine produces an item according to a renewal process. Customers of each class arrive according to independent Poisson processes and differ from each other only in their waiting costs, which are assumed to be linear. We assume a finite number of classes. We consider a linear (both in time and inventory level) holding cost. One would like to find a policy that minimizes the overall cost. An Inventory Rationing (IR) policy is defined by a base stock level and a finite sequence of thresholds, such that a class i customer receive an item if and only if the inventory level exceeds the threshold R_i. The IR policy can be presented as an $M/G/1$ priority system with state dependent arrival rate. In general, the IR policy is not optimal, even when the information given to the decision maker is only the queue lengths of various customers types. In this talk we will discuss cases in which the IR policy is optimal, and we will present an improvement when it is not. We will show that the structure of the improved policy depends on the hazard function of the production time distribution. The analysis is based on busy period analysis of $M/G/1$ priority system and on obtaining the conditional residual service time in such a system.

(M4A) Continuous Optimization [Abirim 1]

A First Order Method for Solving Two-Stage Lexicographic Optimization Problems
Amir Beck, Shoham Sabach*

We consider a general class of two-stage lexicographic convex optimization problems in which one seeks to minimize a convex function over a closed and convex set which is by itself an optimal set of another convex problem. We introduce a gradient-based method, called *the lexicographic gradient method*, for solving this class of problems, and establish the convergence of the sequence generated by the sequence as well as an $O(1/\sqrt{k})$ rate of convergence of the sequence of function values. We illustrate our results via a portfolio optimization problem.

A Sequential Ascending Parameter Method for Solving Constrained Minimization Problems
Amir Beck, Aharon Ben-Tal, Luba Tetruashvili*

In this talk, a method for solving constrained convex optimization problems is introduced. The problem is cast equivalently as a parametric unconstrained one, the (single) parameter being the optimal value of the original problem. At each stage of the algorithm the parameter is updated and the resulting subproblem is only approximately solved. A linear rate of convergence of the parameter sequence is established. Using an optimal gradient method due to Nesterov to solve the arising subproblems, it is proved that the resulting gradient-based algorithm requires an overall of $O(\frac{\log(1/\epsilon)}{\sqrt{\epsilon}})$ inner iterations to obtain ϵ-optimal and feasible solution. An image deblurring problem is solved, demonstrating the capability of the algorithm to solve large-scale problems within reasonable accuracy.

On the Solution of the GPS Localization and Circle Fitting Problems
Amir Beck*, Dror Pan

We consider the problem of locating a user's position from a set of noisy pseudoranges to a group of satellites. Two different formulations are studied: the nonlinear least squares formulation in which the objective function is nonconvex and nonsmooth, and the nonlinear squared least squares variant in which the objective function is smooth, but still nonconvex. We show that the squared least squares problem can be solved efficiently, despite is nonconvexity. Conditions for attainment of the optimal solutions of both problems are derived. The problem is shown to have tight connections to the well known circle fitting and orthogonal regression problems. Finally, a fixed point method for the nonlinear least squares problems is derived and analyzed.

(M4B) Queueing Theory [Abirim 2]

Control of Fork-Join Networks in Heavy-Traffic
Rami Atar, Avishai Mandelbaum, Asaf Zviran*

This work addresses the problem of analysis and control of fork-join networks in the conventional Heavy-Traffic diffusion regime. Standard fork-join networks are feed-forward, which are relatively easy to control. Motivated by healthcare systems, we allow probabilistic feedback, which turns the problem into a challenging one.

In our models, activities are associated uniquely with customers. They are hence non-exchangeable in the sense that one can not combine/join activities associated with different customers - this is the case in healthcare (e.g. emergency departments) and multi-project environments (in contrast to assembly networks).

We introduce a natural concept of optimality for our model, and then solve for the optimal control, asymptotically in heavy-traffic. The central ingredient in the proof is the establishment of asymptotic equivalence between non-exchangeable and exchangeable dynamics.

Finite Two-Queue Systems where Customers of Each Queue Are the Servers of the Other Queue
Efrat Perel, Uri Yechiali*

We consider systems comprised of two interlacing finite queues where customers of each queue are the servers of the other queue. Examples of such systems can be found in file sharing programs, SETI@Home project, and other related applications. Denoting by L_i the number of customers in queue i, $i = 1, 2$, we study three models, distinguished by the way in which service is rendered to each queue. In Model 1, queue 1 (Q_1) operates as a finite-buffer multi-server Markovian queue with Poisson arrival rate λ_1 and exponential service time with mean $1/\mu_1$, for each costumer, such that, at any moment, the L_2 customers currently present at queue 2 are the potential servers of the costumers of Q_1. The overall capacity of Q_1 is N customers. We denote such a queue as $M(\lambda_1)/M(\mu_1)/L_2/N$. Queue 2 ($Q_2$) operates as a finite-buffer single-server Markovian queue with Poisson arrival rate λ_2 and exponential service time, as follows: The L_1 customers in Q_1 join hands together to form a single server with combined rate $\mu_2 L_1$. The overall capacity of Q_2 is K. We denote such a queue as $M(\lambda_2)/M(\mu_2 L_1)/1/K$. In Model 2, Q_1 operates as in Model 1, but Q_2 operates as a finite-buffer multi-server $M(\lambda_2)/M(\mu_2)/1/K$ system. In Model 3, Q_1 is a finite-buffer single-server $M(\lambda_1)/M(\mu_1 L_2)/1/N$ queue, while Q_2 is also a finite-buffer single-server $M(\lambda_1)/M(\mu_2 L_1)/1/K$ queue. The resulting systems are formulated as finite nonhomogeneous quasi birth-and-death (QBD) processes.

For each model we derive its steady state probability distribution function and its related first moments. This is achieved by utilizing the special three-diagonal structure of the QBD generator matrix. We further show that in Model 1, the carried load by Q_1 is larger than its counterpart in Q_2, independent of the queues capacities, while in Model 2 the carried loads are

equal. In Model 3 we show that the effective arrival rate, (due to blocking) in Q_1 (Q_2) is smaller than its mean service rate, being $\mu_1 E[L_2]$ ($\mu_2 E[L_1]$). Numerical examples are presented.

On the optimal service capacity allocation
Rafi Hassin, Yair Shaki*

This paper considers a loss system with a fixed budget for servers. The decision of the system's owner is the number and quality of the servers in order to maximize its profits. We deal with identical and different service rates, preemptive and non-preemptive policies. We present optimal solutions for all ρ larger than zero.

Stability and Performance of queueing networks with infinite supply of work
Yongjiang Guo, Erjen Lefeber, Yoni Nazarathy, Gideon Weiss*, Hanqin Zhang

We generalize the standard multi-class queueing network model by allowing both standard queues and infinite virtual queues which have infinite supply of work. We pose the general problem of finding networks and policies which allow some of the nodes of the network to work with full utilization, and keep all the standard queues stable. Towards this end we show that re-entrant lines, systems of two servers with two re-entrant lines and rings of servers can be stabilized with priority policies under certain parameter restrictions. We further establish simple diffusion limits for the departure and work allocation processes. A third contribution of the paper is with respect to properties of the Markov processes associated with the models. Towards this end, we prove some technical results regarding pettiteness and smallness of compact sets in specific cases. The analysis throughout the paper, depends on model and policy and illustrates the difficulty in solving the general problem.

(M4C) Topics in Operations Research 2 [Abirim 3]

The Split Variational Inequality Problem
Aviv Gibali*

We propose a prototypical Split Inverse Problem (SIP) and a new variational problem, called the Split Variational Inequality Problem (SVIP), which is a SIP. It entails .nding a solution of one inverse prob- lem (e.g., a Variational Inequality Problem (VIP)), the image of which under a given bounded linear transformation is a solution of another inverse problem such as a VIP. We construct iterative algorithms that solve such problems, under reasonable conditions, in Hilbert space and then discuss special cases, some of which are new even in Euclidean space.

Cheap Iterative Schemes for L0-constrained Principal Component Analysis
Ronny Luss*, Marc Teboulle

We present a simple iterative scheme for solving the L0-constrainted Principal Component Analysis (PCA) problem, which can be viewed as a generalization of the power method for finding the first principal component of a matrix. The proposed scheme builds on the so-called conditional gradient method and its simplicity allows for solving large scale PCA problems where usual convex relaxation techniques fail. Furthermore, we show that a variety of recent and novel sparse PCA algorithms which have been derived from various disparate approaches can all be viewed as special instances of our approach. Convergence properties and numerical results are presented.

The Value of Information in a Retailer-Based Distribution Network
Assaf Avrahami*, Yale T. Herer, Retsef Levi

The fact that information has value in the management of supply chains is well accepted. Our work is focused on distribution systems that are based on a network of retailers. The goal is

to explore and quantify the value of additional information in these systems. In particular, we explore the value of the ability to review the state of the system more frequently. In our study we develop a natural formulation for the problem and later on a working formulation for our model. We prove convexity of our model and find optimal solution for the problem. We developed an algorithm to numerically solve the problem and we preformed a numerical study of the problem. Later on we performed a large scale field study. We will report on the savings that the additional information enabled (i.e. the value of the additional information) and discuss in detail what we learned both about the original system and the information rich system.