Optimization-Based Solutions: Smarter Decisions for a Smarter Planet

Event ID# 267632

Moderator: Welcome to today's IBM webcast, "Optimization-Based Solutions: Smarter Decisions for a Smarter Planet." You may submit your questions at any time during the presentation using the Q and A tab at the bottom of your screen. The presenter will answer your questions during the presentation if possible, otherwise he will answer them at the end.

Our presenter today is Dr. Irv Lustig. Dr. Lustig is the Business Unit Executive, ILOG Optimization Solutions Leader for IBM. Prior to this position he was the Vice President at ILOG Worldwide ILOG Direct Channel, responsible for sales, marketing and presales of ILOG's optimization and visualization products.

He has been a manager of technical services for optimization and visualization, ILOG's optimization evangelist, and the product manager for the ILOG CPLEX and ILOG OPL Studio products. He was also one of the lead developers of ILOG's market leading optimization product, ILOG CPLEX.

Prior to joining ILOG, Dr. Lustig was a professor at Princeton University. He received his Ph.D. in Operations Research from Stanford University and his Bachelors in Applied Mathematics and Computer Science from Brown University.

At this time I will hand the presentation over to Dr. Lustig so we can get started. Irv?

Irv Lustig: Thank you, Wendy. I would like to thank everybody for taking your time out today to listen to this presentation. So what I'm going to be covering is a broad view of the different kinds of applications that we see in mathematical optimization and constrain-based optimization that is used throughout a variety of different industries. The presentation is meant to be introductory, so if you have no familiarity with optimization, then this will give you an introduction. And if you have some, maybe you are somebody who actually really understands all the math that is behind it, hopefully some of the messages I give you hear will make it easier for you to talk to your colleagues about some of the benefits that are seen by applying optimization to real world problems.

So I'm going to start off today by talking about the analytics landscape and mathematical optimization fits into the spectrum of analytics. Then we will say what is optimization all about, how does it actually work. Talk about IBM as the market leader in the space. And then I will go through a potpourri of different applications in a number of different industries. And then we will finish off with some concluding remarks and some time for some Q and A.

So in this slide here, you will see a picture that is based on a diagram that appears in Tom Davenport's book called, "Competing on Analytics." And if you take a look at this diagram, the way to actually read it is from the bottom going up. There are different levels of types of analytics that you can do. If you look on the bottom third you will see things like standard reporting and ad hoc reporting and query drilldown. These are the types of things where you are looking at what has happened or what is happening and this falls into the realm of what we call descriptive analytics.

As you move up you then can do things that are based upon alerts where you look for some patterns that are happening and you might alert for a certain situation. You also can do things like simulation and forecasting and predictive modeling and we have in those areas which fall into the realm of predictive analytics, the IBM SPSS product line. And then sitting above that you have the area of optimization. So as were descriptive analytics looks at what has happened and what is happening, and predictive analytics tries to predict what will occur in the future, prescriptive analytics is represented by optimization and the area of stochastic optimization really tries to apply techniques to say this is what you should do about it based upon what is happening and what has happened, based upon what you think is going to happen in the future, what decisions should you make so that you can achieve the best outcomes.

And if you combine the areas of predictive and prescriptive analytics, we call that advanced analytics and here at IBM we actually have our business analytics and optimization service line in our global business services organization and they have a special practice in the area of advanced analytics and optimization that focuses on predictive and prescriptive analytics.

So here today we are going to be focusing, of course, on optimization but I wanted to make it clear where optimization fits within the analytics landscape.

So we like to think of optimization as sort of the science of better decisions. And what we see today is that optimization is helping businesses and governments create the best possible plans where they can consider different alternatives and understand the tradeoffs and how they can also respond to different changes in business operations.

Now let me illustrate this with a few examples as represented by the pictures on this slide. So starting at the top you will see a picture of an airplane. And the question: How to best allocate aircrafts and crews? So almost all airlines use optimization to decide which airplanes should be put on which flights. You don't want to put a large airplane with a lot of capacity on flights with low demand. You don't want to put smaller airplanes on flights with large demand. But you also have to make these airplanes connect up onto their schedule. So for example, if you take a look at Delta Airline with their hub in Atlanta, you might have an airplane that lands in Atlanta, now it has to be assigned to a flight with sufficient turnaround time to clean the airplane and put the food on, etc, and then move it on to the next flight.

So you construct tours for the flight, for the airplanes and they also have to be maintained. You similarly have to do something to allocate crews. The flight attendants and the pilots have to jump from flight to flight but they are limited in how much time they can work per day. All of these decisions are made using optimization.

Moving over clockwise to the right you will see a picture of some food and the line "inventory cost vs. customer satisfaction." So if you are a retailer, you have to make decisions about how much you should keep of inventory of different products on your store shelf. Now in an ideal world you would like to be able to never have a stock out and always have the right amount of units for you ready for your customers to pick up on your shelves. But you have to trade that off with of course the cost of your working capital and the cost of your inventory.

So how to make those tradeoffs between the products that are in high demand that require restocking frequently versus products that have very variable demand and maybe don't require stocking as frequently so that you balance these issues of cost versus customer satisfaction is something that optimization can take care of.

Moving now to the bottom right, you see "cost vs. carbon emission." So as companies now are trying to make carbon neutral decisions, for example deciding how can you set up a supply chain distribution network in such a way that you are minimizing costs but also minimizing the effects on the environment. In some countries in areas we are now seeing the requests to limit the amount of carbon emissions that are done in terms of the transportation of goods. So optimization can help make those kinds of decisions.

On the bottom left you will see a picture of a stock ticker. We have lots of different financial companies that use optimization to determine things like trade settlement as well as trade matching and also portfolio optimization. That is even here at IBM, we have one of our CPLEX customers actually has a product that helps our employees allocate their retirement investments into different portfolios, all driven by optimization.

And finally at the top left, you will see a "what to build, where and when." One of the key applications is in supply chain manufacturing. You have to decide what products to build, where to build them, what plants should build what products in order to meet demand which is now going global, and how to build them, when to build them, etc, when do you order the raw materials. And all of those decisions are driving by optimization.

So the kinds of things optimization is doing is really to understand how to make complex decisions about limited resources. We often see that you are able to find options for how you would make decisions that you wouldn't' be able to make by hand. So you would otherwise having to evaluate millions of choices and optimization is able to make the best choice along all of those millions if not billions of choices.

You also have types of decisions that are automated and streamlined. So you can have decisions that are made by optimization where it is matching buyers and sellers and trading markets or doing things such as determining the routes for crews that would otherwise be difficult to make. And what this allows people to do is then to try different alternatives. So for example when one airline is acquiring another, they look at what would be the impact of adding the second airline to the first airline's route network in terms of planes and crews etc.

And of course with optimization you can explore lots of different alternatives. So you can determine what would happen if your demand should increase or if demand should decrease. And you can gain insight into the different inputs that are driving decisions and view the results in new ways.

So, as I have mentioned, optimization is used to solve resource allocation problems. The types of resources that we see people managing, companies manage, are things like capital. How do you allocate money to different types of investments that are strategic? You have people that need to be scheduled, assigned to tasks, or if you want to grow your operation, how many people should you acquire to cover different things at the reasonable costs. If you have equipment that you have to acquire or locate and manage. Facilities that have to be located and scheduled. Vehicles that have to be routed. Or raw material that has to be acquired. So when we talk to our customers and they have different kinds of resources and we look for these types of key words like buy and sell and schedule and if you are making these types of decisions about resources, then optimization is often most applicable.

So the types of activities that we see that are covered by optimization are things like planning and scheduling activities, so you are deciding a plan of where should you located warehouses. Which warehouses should deliver to which stores? But you have complex operating constraints. There is going to be a limit of capacity of a warehouse. You might have other complex design decisions that are involved. And you typically are trading off of different objectives. It might be to minimize cost, maximize profit, but you also have to consider different service levels. And then you also want to be able to do some what-if analysis. What would happen if things were to change from what we predicted in the future so that you could understand the effect of uncertainties in your demand? And optimization can be used to help make those kinds of decisions.

So we see optimization being used along different time horizons of decision making, from the long-term planning where you are making decisions on an annual, quarterly or occasional basis. You are going to decide whether you want to make a plant bigger, that is a very capital intensive decision. If you are going to close down a warehouse or close down a distribution center or open a distribution center, those are decisions that are not made every day. You make them because they are strategic to you growing business.

You also move then to type of short-term planning decisions where you might be determining a weekly or a monthly schedule of a plan or different people that should be working in order to meet the demand for your products. So, for example, now we are in the holiday time. How many people do you need to cover the manufacturing costs or your extra distribution if you are in the distribution areas?

And then you move down into the detailed scheduling where you are making decisions that can come down even to the hourly. This could be determining what should be our schedule for our operations. What resources can be used? Do we need to do some short term maintenance operations? So you see that optimization gets applied in a variety of different timeframes for making decisions.

Now let's talk about how does optimization work. So any part of analytics is always based upon data and optimization is no different. What you typically have to understand first is to collect some inputs. So on the left here you will see that you have to understand, for example, the demands that might need to be met for your production schedules. You need to understand the resources that are available, what are the limits on those resources, how are they currently committed. What different kinds of yields and recipes cause these resources to interact? If you are going to have a machine making a product, typically there is going to be operators for those machines. Or you have a limited number of machines and operators and demand for product – these things all interact.

You have other operational constraints that can come in and have different preferences for wanting to look at the tradeoffs between cost and service. And then you also have different business goals that come into play. When you combine all of these inputs, what needs to be done then is to write a mathematical model of the decision problem. The decision problem is, how should we allocate these resources in the most efficient way. How do we represent that allocation? How do we measure the different KPIs of that allocation? And that needs to be done in a mathematical model.

And that is sort of the secret sauce is somebody has to write that model. And we will talk a little bit more about that in a moment.

When the model is complete and you combine it with all the data, that is sent into some optimization engines. The optimization engines are very powerful mathematical algorithms that combine the model with the inputs to produce an answer. That answer shown on the right is going to schedule or plan with some different metrics that can measure the efficacy of the schedule or plan. So the things you look at are minimizing the costs or maximizing profit. You had as the output the specific resource assignments or the best possible timing of your activities in a schedule.

So the key aspect of optimization is that you collect a bunch of data, somebody has to write that mathematical model. The engine, which is the key technology that IBM is supplying is then used to produce an answer that is the best allocation of the limited resources.

So now let's talk about how the IBM ILOG Optimization and Supply Chain solutions fit into this offering. So our key offering is shown at the bottom here. It is called CPLEX Optimization Studio. And what that contains is the engine that I mentioned, the mathematical engine that is used to make these decisions. And it also includes tools that make it easier to write models. So I mentioned about the importance of writing a mathematical model, we have a language that is used to express those mathematical models.

The people who typically purchase CPLEX Optimization Studio are in an analytics department or an operations research department. And the most important things that they care about are the performance of the products. These are people who generally know how to write these math models. And they are going to apply it to their specific business problems.

Then we also end up talking to people who are in IT departments or line of business and we have products there called ODM Enterprise. ODM stands for Optimization Decision Manager. And it is a solution platform that makes it easier to develop an application on top of the base engine that is provided in CPLEX Optimization Studio. So it includes things to make it easier to connect with data. It includes a scalable deployment platform based IBM's WebSphere Application Server. It also has the ability to be extended with visualization so that you can see the answers and work with the data in a real application.

And then on top of ODM Enterprise we have vertical assets that we are developing that are pre-built models that can be used for specific kinds of applications.

And then finally we have at IBM our supply chain applications. And you can think of these applications as being based upon the same optimization engine, but these are for specific decisions that are required in the supply chain. So we have a product called LogicNet+ that is used for sourcing decisions as well as strategic network design. We have a product called Inventory and Product Flow Analyst that is used to do inventory optimization as I referred to earlier. And there are a couple of other products as well.

So the difference between how these products are used is that when you use CPLEX Optimization Studio, you typically need to be a person who understands how to build those mathematical models. When you use ODM Enterprise, you generally, you could be one of those people, but we also offer and we sell as services people who can build those models and then the platform is used to quickly deliver an allocation.

And then finally in the she Supply Chain Optimization Applications, these have pre-built models that are package inside as an application that is for a very specific problem domain.

Now, our optimization customers are far and wide around the world, over 160 of the Global 500 build custom applications using our engines and tools, most of which is in manufacturing, transportation and finance. We have over 1,000 commercial customers under maintenance. We also have a big program where software vendors embed our engine inside their offerings. So SAP, Oracle, JDA, Manhattan Associates, etc, all of these companies use our optimization engine inside offerings that are for very specific problems in specific industries.

There is a set of other tools beyond our own tools that can be used to build the mathematical models. We have partners that resell CPLEX with their modeling languages. We are also the leaders in university research. So almost any university that is doing research in optimization, applying optimization will be using CPLEX in their research activities. And then finally as I mentioned earlier, our Supply Chain Applications leverage this technology. So our technology is not only directly seen by customers in touching our products but is also seen in very indirect ways through our software vendor partners as well embedded inside of our own offerings here at IBM.

So now what I would like to do is turn over and start talking about some of the applications of optimization and in a variety of different industries.

So on this next chart here what you will see is a different set of industries going across the top from manufacturing, transportation logistics, etc, and there are a number of different applications that are listed within each one of those industries. And I'm going to be covering some specific examples in a moment. But let me point out a few of them.

So in manufacturing we talked earlier about inventory optimization. I have alluded to supply chain network design where you are deciding which distribution center should supply which stores. Which plants supply which distribution centers. Where should you locate you distribution centers. These are all strategic supply chain decisions. You have applications such as in maintenance scheduling, determining when should you maintain your equipment. And transportation and logistics. We mentioned fleet assignment already in the area of the airlines. Different types of vehicle container loading. In rail we see things related to the management of yards and crews and locomotive engineers. In financial services, portfolio optimization is used and almost all of the financial firms are leveraging mathematical optimization to optimize portfolios. We see things like balancing portfolios or creating portfolios that track an index fund. These are all done with optimization.

In the utilities area you do things like determine the schedule of when you generate power. Another problem is called unit commitment problem, is when should you turn on a generator, how do you turn off, because it takes time and energy to turn things on and off.

Mining is another big area where optimization is used to determine how to dig into the ground efficiently to pick out the material that you want to use for mining.

In telecommunications one of the older areas where the capacities of networks are planned using optimization. Also they have to route and schedule the people who come out and fix the wires. All of that is done with optimization.

And then you see some examples in the last column across the different industries. One of the newer ones we have seen recently is in marketing campaign optimization where you have different potential marketing campaigns that you wish to allocate to different customer segments and how do you do that to maximize your return on investment.

So here, again, is a sampling of different applications across a wide variety of industries and what is really interesting here is that the same fundamental mathematical model is used or class of mathematical models are used to solve these different problems which is why our engines can go out and solve these problems because all of these problems can be reduced to a very similar mathematical expression.

I also want to talk about some well documented use cases and really where you can see the return on investment that our customers get by applying optimization.

So on this slide here is a list of a number of different customers that have used the optimization. The first group that has an asterisk next to them come from what is called the Edelman Competition. I think the footnote got lost on our slide here. This is a competition that is run by Informs on an annual basis to compete – five or six companies will compete for the best project that was done in the year, is based on operations research which optimization is a subfield.

About something like 65% to 70% of them involve optimization and the concept has been going on for 30+ years and our product has been around for about 20 of them. And I would say easily over half of the ones that use optimization are based using our products.

So the ones with asterisks here, what's interesting, these have been vetted by a third party, by the judges to actually verify that these returns were seen. So let me point out a couple of the interesting ones. At UPS they did Air Network Design. So this was the fleet, the network for the overnight air operations. And they had to decide how many planes would they need in order to meet their demand. Now at the time they were unsure of whether they would have a little bit of demand, medium sized demand or high demand, so they ran different scenarios which allowed them to determine the fact that they could reduce the number of planes by 10% and also by allocating the planes more efficiently reduce their overall operational costs.

Another example is at Continental Airlines. So there, this project was done shortly after 9/11 in 2001 and the issue here was Crew Re-scheduling. So as I mentioned earlier the airlines use optimization to determine the scheduling of their crews but for airlines they have disruptive events such as 9/11 or for Continental it would be a blizzard in the Northeast or a hurricane in Houston and you can imagine other types of events or other different types of airlines. When these events occur, flights get cancelled, delayed, and it affects these crew schedules. So what you need to do is now adjust the schedule so that the crews can be rescheduled and get back towards the original schedule. Well there is some new work that was done, Continental along with a partner of theirs and a partner of ours, they used optimization to drive better decisions about how they would do this and in fact Continental was the first airline that was back up and running after 9/11.

So those are a couple of examples. On the bottom four you will see some examples from our supply chain applications, here companies like Pepsi, which I will talk about later, have used our applications to recognize $6 million in inventory reduction. Fonterra is a New Zealand Dairy Cooperative that ships powdered milk around the world. And in fact I think it is the powdered milk inside of Kraft cheese is one of the things I learned possibly.

In any case they had a big supply chain problem that they had to face and they used our supply chain applications to significantly reduce their annual costs in supply chain. So I want to point these out because what you can see is that the impact of applying optimization to more efficiently manage your resources is really across industries, these are real measured impacts that save companies millions and millions of dollars.

So now let me talk about some more specific examples in some more detail and we will start with some applications in finance. So first I would like to talk about insurance underwriting at a US Health Benefits Company. So the, at this company there was a, they had to do health benefits and they had to modify their underwriting system to make it easier to configure potential policies. So here what you want to do is determine what should be the price and plan and options that best correspond to the customer's requirements. So by doing it, applying this to optimization, they were able to improve their customer service and their experience by automatically customizing these plans from lots and lots of different possibilities. So the idea here is instead of having a menu of plans to pick from, the customer can come in and say I would like to have these types of things in my plan and the plan can be dynamically customized to meet their needs and to also meet the underwriting requirement. And so now the underwriters are much more productive.

In another example here we have a portfolio management at a major investment firm. So here in this firm they actually have three different applications and manage lots of assets for investment banking, big and large investments that they have to manage. And there are three different applications that they use. So the first is in fund rebalancing. So they have to have funds that actually track indices. And you want to be able to rebalance those funds while reducing your overall transaction costs. So you use optimization to figure out how to do that. Another example is in trade crossing. So here what they would look at is different, they have buy and sell orders that are happening and there are similar assets in different classes and they want to be able to match them up in an efficient way, so again to reduce overall

transaction costs. And then the third application was what is called in-kinding and here the idea was that they had different funds and instead of taking a security from one fund and putting it out in the market and then having another fund buy it, the idea was to match that up and do an inside trade to, again, avoid the transaction costs. By using optimization they have chosen some significant competitive differentiation, being able to meet their different guidelines and requirements that are there. And actually save, their estimates, $500 million in transaction costs for their clients.

The next example was S.D. Indeval and the Central Bank of Mexico. And this is a recent winner of the Edelman Competition that I mentioned earlier. In this application, they applied optimization in a novel way for trade settlement. So trade settlement is typically done on a T+3 basis, it take three days to settle a trade. They wanted to do something that would allow them to settle trades in real time. So what does it mean to settle a trade? Well if you think about having a specific security you could have brokerage A selling to brokerage B who is also selling to Brokerage C. Well, what would have to happen there is for brokerage B to be the middle man, he has to have cash to buy from A and then he gets the cash when he sells to C.

Well, if you can match A and C together, you could avoid B as the middleman and this is called netting. And the idea here was you could reduce the liquidity risk that these different traders have to have a lot of cash to be able to move around, but instead of settling trades one by one, you put them all in a pool, look in at all of the buys and sells, matching them together, and making it so the securities and the cash moves a lot more efficiently. So by doing this, and they do it every two minutes using an optimization, they have reduced their intraday financing costs by $130 million per year and they have also reduced the daily liquidity requirements by $150 million per year which means the amount of cash that these different brokerage and bankers would need to have in order to be in the market.

So this is a major project. It was a revolutionary way of changing the way that trade settlement happens and it was all driven by optimization.

So now let turn our attention to the energy and utilities area. So the first application I would like to take about is a unit commitment at Red Electrica de Espana. And this was actually a problem that they had been solving for many years. And the unit commitment problem is a problem of determining when should you turn on different generators based on different types of power – nuclear, coal, water, etc, and recognizing that there is cost for turning on things and cost for shutting them down. And obviously you want to make sure that things are on when there is high demand and you need to have things that are off within low demand times. And there are limits to how often you can turn things on and off.

So they used to do things in a very heuristic way with an interactive methodology that they would do things in a very trial and error type of way, methodology. By using optimization based on our CPLEX and OPL products and also ODM, they were able to make it easier for the engineers and the people to make decisions and to really see and understand the impact of different decisions by letting optimization figure out the best possible schedule.

They have actually recognized per day their overall production costs dropping between $50,000 and $100,000. And it turned out it also had a side benefit of reducing carbon dioxide emissions. So on an annual basis they are estimating that they have reduced their carbon impact on the environment by 100,000 tons on an annual basis. And that is pretty significant and that wasn't even a goal of the project. It was as side benefit by the fact that they were doing better use of their limited resources for generating electricity.

The next example is at a Midwestern utility that is doing pumped storage optimization. And this is a pretty interesting problem. So what the pumped storage facility is that you have imagine having a lake or pond of water at the top of a mountain or a hill. And the ability to release water from that pond, and it would go though some turban generators that can then generate electricity and then the water ends up at a reservoir at the bottom. Well, now you have the ability to turn on a pump, reverse the turbans, and make the water go back up to the top of the hill so that it can be used again. So what you want to do is you are limited to how often you can turn on the pump per day to pump the water back, but you also want to push water through the turbans where you can sell the electricity because it is worth more and you want to pump it back up when electricity costs less because you have to pay for the electricity that turns on the pump.

So there are a number of different factors that you have to consider and these decisions are typically made by hand. Well, now by using optimization you get a standardized business process that will give you a valid schedule that you know is mathematically valid. And you can find opportunities that may not otherwise be obvious. It also helps the operators to understand what is the value of the water in the pond and make a decision as to understand because that water can turn into electricity, how do they move it. So they get better utilization of the plant, and as they are in the process of implementing this they anticipate they will get an improvement opportunity of about $8 million annually and they have an initial goal of only hitting 10% of that. Because they have to manage the people to get them to start following these new plans that are driven by the optimization engine.

Here is another different type of application. So in the first few applications we talked about application in energy which were about managing sort of the energy itself, the generation of electricity, for Red Electrica de Espana or (inaudible) water which is used to generate electricity.

Here in the National Electricity Market of Singapore there is a market for trading electricity between the different utilities that are actually selling the electricity. So you have to actually make the market and meet the buyers and the sellers while also making sure that nobody is buying all of the electricity in the market because they would supply their customers leaving other customer in the dark.

So here they used CPLEX to do this and they solved the problem within thirty seconds, every half hour, to determine what is the best way to operate this electricity trading market.

By using CPLEX in this application they have been able to make sure that all of the constraints that are regulatory are being followed and they also know that the customers are now getting the lowest generation costs of electricity for everything in the wholesale market. And the market is now more efficient. It has also make it easier for them to maintain the trading system. So if a new regulation comes along, it is easy for them to introduce that regulation into the mathematical model.

And so they are very efficient because our tools are easy to use from the development standpoint in making those types of changes to the systems.

The next success story is looking at natural gas and E.ON Ruhrgas which is in Europe is looking to understand when should they purchase contracts for natural gas production or use gas that they might have in storage. So they developed an optimization solution that looks at the margins of the quantities of these contracts and performed some different sensitivity analysis. And so then they used these results to do a better job of managing their contracts in existing storage facilities so that they could understand their pipeline capacities and their negotiated purchase costs.

So the benefits here again are that they have better plans, better decision making. It is allowing them to react quickly to market changes. And they can look at a lot of different potential scenarios. So they can say what happens if, hey, we see a big storm coming in and we think it is going to get very cold and demand is going to go up for natural gas – could they hedge against that and make a better decision on a purchase contract.

And so it is very important for them to have a reliable implementation of the underlying mathematical engine so that they can respond to these changes quickly within the gas market.

The next area I would like to turn to is applications that are in transportation. So I have already talked about some of the sample applications in the airlines from crew scheduling to making sure that you are moving the crews through the network. The fleet assignment to assign the right planes to the right flights. There are also applications in ground staff scheduling. So we have airports in the world where you are deciding how to you schedule the staff to make the planes come into the gates and take them out of the gates or unload and load cargo and baggage, all of these types of things have to be allocated. You have to allocate people with skills, those are applications that are using ground staff planning in airports. We have airports where optimization is used to allocate gates towards incoming flights. So in airports where the gates are shared between airlines, how do you efficiently allocate that according to the schedule.

Another major area is in revenue optimization in yield management. So here the airlines are deciding how many different seats to offer at the different price levels. And if you take an airline that has a hub, when flights come into a hug and when flights depart the hub, you have people who are paying different

fares for the different connections. And how do you allocate that in an efficient manner in order to maximize revenue.

And then I mentioned earlier the success of Continental Airlines to do irregular operations. At disruptive events, how do you handle that in order to get your crews and your fleets back online.

We have a large tire manufacturer who has $150 million transportation budget but is using optimization to better manage their hubs, their driver assignments and to decide which kinds of fleets they should use in terms of for-hire or private fleet decisions. They have a lot of dealers and a lot of different trucks to manage. And they actually, these results are used by a lot of different planners on a daily basis. They saved a lot of – several percent – off of their transportation budget by using optimization. Made it easier to max their supply and demand match. They are able to plan for a more finer level in terms of not just doing daily planning but actually planning for every 15 minutes. And they are able to manage more orders by using optimization.

Our next example at transportation is Netherlands Railways, I realize there are a lot of words on this slide, this is also another Edelman winner. And here they were using optimization to basically redo the entire way they were doing scheduling and allocation of resources to fore their network.

So one of the things that they needed to do was allocate the rolling stock which is the actual cars to the different train schedules in order to meet demand. So once of course you do this, you then need to connect the trains up with the cars to the next train it would be on and meet the vans so all of this was done using optimization. And the types of benefits that they recognized was getting a €40 million annual cost reduction by applying optimization to better operate the railway.

At the United States Postal Service, there was a project that was actually done by IBM Global Business Services leveraging CPLEX to determine how to best manage the different classes of mail within a corridor. And so the idea here is to make sure that you – instead of having a truck per different type of freight, could you consolidate some of the freight, priority mail, first class mail, the standard mail into one truck without sacrificing service levels. They estimated that there would be a $5 million in annual savings in one area of the country and they are now looking into rolling this out into different areas so that we can keep our postal service afloat and hopefully keep our prices of stamps down.

So this was really an important project for the United States Postal Service.

Now let's look at some of the applications in consumer packaged goods. So we have a food manufacturer that is doing supply chain network design. And they had a number of different food products, a lot of different production facilities all around the globe. Different, multi-tiered production

manufacturing process where you make sort of an intermediate product at one plant and have to ship that to another plant to make the final goods. And they were forecasting some very much increased demand. So they did a strategic project to answer some questions. Where should they expand in their network considering that they think demand is going to be increased? Do they need some new plants, if so, where should they be? Should they do more outsourcing of production? How do they best utilize their existing facilities? And where should they do different products in different regions?

And so the whole idea of the project was to really make it so that they were doing a much fact based decisions based upon real data and understanding what the different [tradeoff] are because they have to bring in things such as labor and real estate, freight and duties and depreciation and taxes. So in order to be able to facilitate the decision making, they use optimization to make those better decisions.

So the results of this analysis were that they should do some regional sourcing in Southeast Asia and Russia. They should do some co-manufacturing in the short-term in order to meet the increased demand until the new plants can come online. They also evaluated closing certain plants that would meet their demand. And they were able to have a much better capital plan for how they should expand their business. So the tool was used to make these very strategic capital intensive types of decisions.

Pepsi Bottling Group, this was documented in a public case study, they were seeing that some new patterns in their demand was outstripping their bottle lines availability to actually deliver the product. Everything was at capacity and in fact their peak demand was outstripping their production capacity.

So what they did is they created a process which would make it have a better strategy for doing production sourcing to get the raw materials and making sure that they were getting better customer service. So now they like to say that they have a 21st Century Supply Chain as you can see on the quote on the bottom, and they really got some nice results by being able to reduce the number of out-of-stock and reduce the number of raw material and supply inventory. And so they reduced their inventory on an annual basis by $6 million of this raw material.

And then because even though demand was growing and revenue was growing, they were able to have a more efficient network for delivering products so they reduced their transportation costs. So again this comes from applying our supply chain tools to these kinds of problems.

The next example we have an apparel manufacturer. They were very, had low inventory turns throughout the network. So that meant that they were very rarely restocking, etc. They were generally overstocking. They did not have a good system for how they would set the target inventory levels across all of these different types of brands that they had and for all to use. So they used our Inventory Analyst product to actually compute the optimal inventory levels for each of the products based upon its forecasting of what they believed demand would be.

And so now you have a centralized methodology for determining the inventory levels and everything is tied into their ERP systems. So they were able to reduce their inventory levels and the formalization of the process means that it is much easier to train somebody new to actually look at it, do the review of the suggested inventory levels that come out of the tool.

So now let's talk about some applications that are in manufacturing. So we have a diversified mining company. I mentioned earlier these applications in mining where you have to dig in the ground. And you want to determine where and when to dig so the pit doesn't collapse on top of itself. But you also want to be able to dig for the good stuff that is actually worth a lot more.

You also have to meet your demands and you have to handle all different things like setup costs, what are your royalty obligations, etc. So by applying optimization this particular mining company saw a cost savings of 5% which was more than $35 million. Instead of having taken months to create a dig plan, they can create dig plans in days and they can understand what different types of things would occur over the long term for them in terms of their overall plans to mine the pits.

At Nissan, this was an interesting example that our team worked on with a partner. They have a plant in the United Kingdom that was actually a very efficient plant and was building two models of a car. Typically when you build cars you will have one production line per car. So the first car will be – so one car will be on production line one. And the second car will be on production line 2. Well, they wanted to start building a third car model which typically would involve a large capital expense to create that third production line.

Well what they discovered was that there was capacity that was available in the first and the second line and that if they changed the rule and allowed a car to bounce back and forth between the two production lines, they could more efficiently use their existing capacity but that they could essentially build a third car on a virtual production line that would be shifting back and forth.

So they actually, this was actually built by our professional services team. They were able to actually increase their overall production capacity but the other important aspect was that they were able to create schedules at the beginning of the day that at the end of the day would adhere to, substantive things like breakdowns, etc. So they went from a 3% schedule adherence, the schedule they thought they would have, to a 90% schedule adherence which made the plant much more efficient and they also were able to avoid the capital costs of building the third production line.

We have another example here in the petrochemical area. So this was a European company that was trying to decide if they could consolidate some of their distribution centers. So they had five production facilities and 14 DCs and wanted to know what would be the optimal number of DCs. Now what you can

see here on the left is that the current design had 14 distribution centers with a total cost and an average service level of around 304 kilometers. If they increased the service level they could reduce their cost because they would have less DCs. And there are different drivers that drive this kind of decision. So they have to make a very strategic decision if they are going to sacrifice service but reduce costs, both operationally as well in their overall distribution center costs.

So in this particular example here, this was a strategic study to determine that they could better operate their overall operations by using optimization to figure out which DCs were really actually needed.

So now let me conclude with some final remarks and then we can open the floor up for questions. So I would like to say hopefully you have seen with these different set of examples and different industries that optimization is pretty much everywhere. No matter where you go, we used to have a marketing person who used to say our products touch you every day. If you buy a product in a food store, most likely our optimization product was used at some point along the supply chain to bring that product to your store shelf.

There are applications that I mentioned in multiple different kinds of industries. You get real ROIs, and mathematical models that are doing, have built into them the objective of minimizing cost or maximizing profit and the kind of return on investment incur in months, sometimes even weeks because there are such dramatic returns on investment. So you can avoid capital expenses, you can reduce your operating expenses, and you get better revenue, revenue mix and margin improvements.

You get better customer satisfaction. So you can build better service into things, making sure that the right products are on your store shelves for managing inventory. That you have better schedules that you know that you can adhere to which then turns into better service for your customers.

You have better employee satisfaction. So the employees who are getting scheduled feel better because they know this is a schedule that can be done, but you also have the people who are generating the plans and the schedules happier because they have a tool that has really make their job a lot more easier, and now they can use that tool to explore lots of different scenarios. You get better decisions by applying optimization. You get faster decisions because you can now use the tools to automate and make decisions more quickly. And we would like to say that you get smarter decisions for a smarter planet. We think that optimization is a key aspect of IBM's movement into creating a smarter planet and we hope to have a lot more success stories to come now that we are all part of IBM.

So I would now like to turn it over to the Q and A session. I will remind you, and I think Wendy will as well, the questions can only be asked via the Q and A tab that is in your presentation manager there. So feel free to submit a question and then we can get around to answering them.

So taking a look at the first question here which says the following. It says, IBM has ILOG LogicNet+ for supply chain, I mentioned that earlier, and his question was: Should you go for the package tailored specific for supply chain or should you build everything from scratch?

So I think it is actually better to go for the packaged application as long as it meets your business needs. And we tell this not only for the packages that we sell but even for if you are comparing building a custom solution using our tools yourself versus buying a package that comes from one of our partners that embeds optimization.

And what we tell our customers to do and our prospects to do is to look at your business problem and make sure that all of the requirements, the constraints, the different preferences that you have can really be made, met, by the application tool at hand. We often debate this ourselves; we have prospects that have some interesting things and we think that it is going to fit for LogicNet+ and we get down very deep into some very deep detail and we find out that LogicNet+ is not able to handle a certain kind of business constraint. It is pretty rare. I know I have had that in my experience. So my answer to there is that if you know that you have a supply chain that would design a problem, a sourcing problem that you really should evaluate the capabilities of LogicNet+ first and then say is it worth it for, is it going to meet those needs and that it is easier to use that application out-of-the-box then it is to build something yourself by hand because we have built in connections to databases and cut and paste from Excel and better graphics and visualizations to getting all of these types of things all built on top of the engine.

If you build something on top of CPLEX yourself, that is all you are doing. You only have the engine. You have to build the visualizations, the database, the Excel connectivity yourself which can be quite laborious from the implementation standpoint.

So at this point I don't see any other questions that were addressed to me. Let me hit the refresh button – there are some more. Okay. So let me read this next question and see how I will handle that. Okay, so somebody is asking a question on when we solve mixed integer programs which is one of the types of problems that we solve with CPLEX, what kind of sizes have we actually solved.

And I would say this is not like some of those ads you might get in your spam that say that "size is everything." Size doesn't matter. And the reason I say that is we have mixed integer programs that have just a few hundred variables that are extremely difficult to solve the optimality. We have others that have millions of variables and constrains that we were able to solve the optimality.

So problem complexity has a real big issue. One of the things that is true is that we have been pushing the envelope more and our capabilities and we are now able to – those problems that we were not able to solve just a few years ago in any reasonable amount of time that now we can solve in minutes.

So what we have seen from our customers is they create bigger problems by using more granularity in their decision making. So in supply chain, you might actually make decisions now at an SKU level as opposed to groups of SKUs. Or you might want to make a decision on a daily level or even hourly level as compared to weekly or monthly types of planning.

So the size allow you to get to bigger granularity and then you have to ask the question of, well is it that important for me to get a provably optimal solution. So, hopefully that addresses that question.

Let me go to the next one. The next question says, do we use a simulation package to see what will happen if a company implements the optimal product plan generated by ILOG CPLEX.

We have seen our customers do that type of thing. As it is today, IBM nor ILOG has a simulation package. I know that this is actually done within the business analytics and optimization services team. They are big users of simulation. So that is certainly a technique that is used to evaluate the quality of a planning decision that comes out of the optimizer.

Let's see if there are any more. There are a couple more that are assigned here.

The next question says how do you include model intangibles such as customer perception in your modeling? That is a really good question. The whole idea here is to be able to model those types of things using some quantitative methodology. So for example, if you take customer perception, I'm going to move it over to sort of customer preferences.

I mentioned earlier the application of marketing campaign optimization. So in a marketing campaign optimization you are trying to decide which offer to make to which customers. And you have the different type of offers and you have to have some way to measure what is the probability of a response from a customer and what is the likelihood that they will purchase and what do you think the purchase amount will be?

Well the techniques for all doing that come from predictive analytics. So we look to our predictive analytics folks from SPSS to produce the predictive models to measure things like customer perception or customer preferences. And then we use that as input data into our software so that we can then optimally allocate the resources. So we are very dependent on optimization on measures of goodness and measures of resource availability and measures of what we think a return would be on a certain decision. And so we are dependent on other methodologies to produce that.

The next question says, can I solve the stochastic in your programs using ILOG CPLEX. So we do not supply a direct stochastic linear programming solver. The reason being is that in our experience we see a lot of different ways that people want to apply uncertainty modeling and stochastic programming to different types of problems. So on the other hand we have people who build stochastic programming solutions using CPLEX. So based upon their uncertainty model, they might build their own algorithms in order to represent the problem and pass it to CPLEX in an algorithmic way to take advantage of doing decisions under uncertainty.

Our next question, how widespread is optimization for blending in the food industry, for example wine, juices, coffee, tea and tobacco?

I believe I have heard of those types of applications. Blending is actually a traditional application that has been done in the refinery and oil refineries where you are trying to blend different types of fuels to come up with the different stuff you see in the tank. I'm not so sure I have heard about it necessarily being done within the food industries. One possibility is that there is a lot of regulation involved. On the other hand than may just be my own personal familiarity as to what has occurred in that area. Certainly blending problems are the types of things that we have seen where you are trying to come up with the right kinds of recipes in terms of the manufacturing operations. So that is certainly something worth exploring if you have some certain ideas in your business we could discuss further to see if there are some nice applications there.

Next question is how are constraints represented in the IBM system and are there any limitations on applying constraints? Well, it really comes down to the mathematics. So our fundamental products handle linear constraints, quadratic constraints and also what are called logical constraints. So we have two areas, mathematical programming which uses linear programming and nurture programming. And then we also have constraint programming which is used more in the detail scheduling area where you might have a constraint that says something like either this or that or you might have a constraint that says, I'm trying to think of a good example, in scheduling you might have to say this activity precedes that activity. You are making a decision on the basis of time. So these are represented. We have a language that helps represent these types of mathematical constraints and you then write out the language. You can actually create very big problems with that language because we separate modeling data, make it easier to actually represent these kinds of large problems.

And let's see if there are any more. Boy they are really coming in. Next question says, in my view are there any new areas of industries where optimization will be successful in the coming years or has not been successful or used before. That is a good question.

I have to say that I really often get surprised at the new applications that people think of. So I will take as an example the one I mentioned earlier from Indeval where they were doing trade settlement.

This was a brand new application of optimization. Once you understand it is kind of simple. But the issues on actually making it work were less about the optimization but about changing the infrastructure around it and changing the rules for how the different players would accept an optimization-based solution.

So I think one of the aspects in terms of new areas that we see is in terms of acceptance of the automation of the decision. I think one of the areas that we will start seeing more applications are in healthcare because that is certainly a place where we have limited resources that we could do better allocations of resources. I think the marketing campaign stuff is starting to take off. I think there will be other – sometimes things are driven by how technology is changing. You can start seeing optimization now being applied in areas where you are trying to allocate IT resources.

So, I think the sky is the limit. It takes really the creativity of people recognizing that they have limited resources to manage and that you have the ability of actually modeling those types of resources. So, I'm sure if you ask me that same question a year from now I will be able to tell a story of some client that I learned about that came up with some new novel way of applying our technology. Let's see what else we have. Wow, they keep coming.

So the next question is do we offer more detail training and do we have online training and do we do certification? So we have actually online training courses for how to use our products as well as course that you can take in person. I believe there are links to that information off of our website. In fact I think I saw it the other day. We don't, I believe the certification that occurs is more on the level of, the fact that you have taken the course. Separately we are in the process of developing a certification program for resellers. That is going to take a few months to roll out and that will be the type of certification that we would do.

The next question says, how do you retrieve interim solutions along optimization. So in our tools we have the ability of actually as the algorithm is progressing there are methods and techniques for actually getting the intermediate solutions that the optimizer is producing and being able to examine them and it has helped them to use.

I'm not going to get into the specifics here of how to use the tools, but we provide that availability within the actual engines for you to sort of inspect intermediate solutions as they are produced.

The next question. What kinds of probability distributions can CPLEX handle in stochastic optimization? The answer is none. So as I mentioned earlier in stochastic programming we do not assume an uncertainty model and that is exactly why we have not done anything yet in stochastic programming is because we see lots of different demands for lots of different probability distributions. So, we don't make any assumptions with respect to stochastic optimization. This is a very active area of research in IBM, not just within our own development teams but also IBM research is spending a lot of time looking at different

ways and novel ways to do stochastic optimization, really being driven by client demand as opposed to being driven by advancing the mathematics.

So we are doing a lot of active discussions about what direction we can take in that area.

And the next question is do we have a database of sample solutions in addition to the installed examples. Well, I think the best way to answer that question is we have a lot of different sample models. We actually have some internally developed what we call as assets that act as starting points for certain solution areas. There are IP issues that our lawyers are talking to us about, about what we are allowed to release. So anything that is included in our examples has been approved for release through our development processes, etc. And as we wanted to release more advanced examples we are getting questions about whether we should. On the other hand that doesn't mean we can't leverage them in certain ways, we just can't hand them out willy-nilly. We have to have a process for handing those out.

So if there are specific things that you are looking for, I encourage you to have a discussion with us after this presentation, can put you in touch with members of our team and we can explore further how we can meet your needs in that area.

And that concludes all of our questions. A lot of great questions. I want to thank everybody for all of the great questions and attending today. And certainly if you have more you can contact us and we will be able to follow up with more one-on-one discussions. So, Wendy, do you have any concluding remarks?

Moderator: Well, that is our time. And thank you Irv and thank you to all the attendees who stayed on 8.5 minutes past the hour. We know that you all have busy schedules and we appreciate you spending the time with us. This presentation has been recorded and you will receive a URL for on-demand viewing in the next day or two, and I believe also from the screen you see now you have a button where you can download the slides and when you do download them as a PDF, they will be much larger than you were able to see them today in the live presentation.

So thank you. That concludes our presentation.