

LB Datalog (LogiQL)

- Predicates
 - Represent relations between values:
 - $\text{father}(x,y)$: x is the father of y
 - $\text{mother}(x,y)$: x is the mother of y
 - $\text{grandfather}(x,y)$: x is the grandfather of y
 - $\text{age}(x)=y$: y is the age of x
 - Can be "given" or calculated
- Rules
 - "Make it so" predicates that are derived from other predicates
 - $\text{grandfather}(x,z) \leftarrow \text{father}(x,y), \text{father}(y,z)$
 - $\text{grandfather}(x,z) \leftarrow \text{father}(x,y), \text{mother}(y,z)$
- Constraints
 - "Must always be so" logical statements about what must always hold of a consistent database
 - $\text{grandfather}(x,y) \rightarrow \neg(x = y)$

Mother	Jane	Bill	Father	Joe	Jane	Grandfather	Joe	Bill
	Alice	Joe		Joe	Tim		Joe	Andy
				Tim	Andy			

LB Rules are incrementally maintained by the database

- Insertion: adding a fact that Jane is Martha's mother, results in insertion and update to the "Grandfather" predicate asserting that Joe is Martha's Grandfather
 - Rules are evaluated minimally – only re-computing additions / removals to predicates
 - Atomicity, consistency, isolation and durability of transactions is automatically maintained
- Rules can talk about state/changes (not that important for our talk)
 - $\text{+mother}(m, y), \text{+father}(x, m)$
 - $\text{<- +grandfather}(x, y), \text{!has_parent}(y), m = \text{"motherOf"} + y.$
 - $\text{has_parent}(x) \leftarrow \text{mother}(_, x); \text{father}(_, x).$

Mother	Jane	Bill	Father	Joe	Jane	Grandfather	Joe	Bill
	Alice	Joe		Joe	Tim		Joe	Andy
	Jane	Martha		Tim	Andy		Joe	Martha
	motherOfJill	Jill		Eddie	motherOfJill		Eddie	Jill

LB Deletion

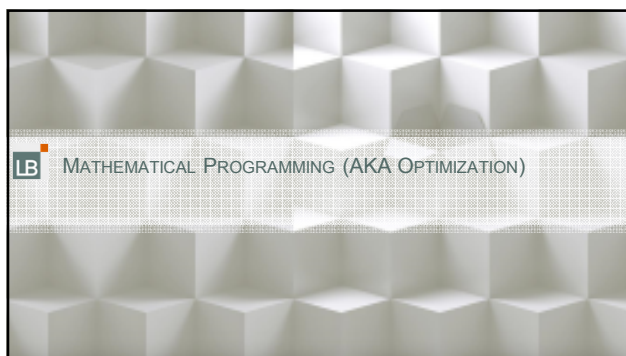
- Removal of a fact describing the relation between Tim and Andy from predicate Father, causes the database to remove the relationship between Joe and Andy in the predicate Grandfather

Mother	Jane	Bill	Father	Joe	Jane	Grandfather	Joe	Bill
	Alice	Joe		Joe	Tim		Joe	Andy
	Jane	Martha		Tim	Andy		Joe	Martha

LB Constraints

- Asserting that both Tim is Joe's father and Joe Tim's violates the constraint that no one can be his own grandfather; the system aborts the transaction and the database is restored to the original state prior to the addition of the violating facts.

Mother	Jane	Bill	Father	Joe	Jane	Grandfather	Joe	Bill
	Alice	Joe		Joe	Tim		Joe	Andy
				Tim	Andy		Joe	Tim



Mathematical (Linear) Programming

- A potter can make either plates or cups with different profits from each (\$1.07 and \$1.05 respectively). Each plate takes 0.24 hours to make and each cup 0.2 hours. The total time available to the potter is 11 hours. How much should she make given that she has already existing orders of 4 plates and 8 cups.
- Maximize value of the *objective function* $o(x,y) = 1.07 * x + 1.05 * y$
- Subject to the conditions that:
 - $0.24 * x + 0.2 * y \leq 11$
 - $x \geq 4$
 - $y \geq 8$
 - x, y integer

Profit	Plate	1.07
	Cup	1.05
Time	Plate	0.24
	Cup	0.2
AvailableTime		11
Order	Plate	4
	Cup	8
Make	Plate	x
	Cups	y

Make	Plate	5
	Cups	49
TotalProfit		56.8
TotalHours		11

Mathematical (Linear) Programming

- You can pretty quickly get more elaborate
- Maximize value of the *objective function* $o(x,y) = 1.07 * x + 1.05 * y$
- Subject to the conditions that:
 - $0.24 * x + 0.2 * y \leq 40$
 - x, y integer
 - $x1 + x2 + x3 + x4 + x5 = x$
 - $y1 + y2 + y3 + y4 + y5 = y$
 - $0.24 * x2 + 0.2 * y2 \leq 4$ (dentist appointment)

Profit	Plate	1.07	
	Cup	1.05	
Time	Plate	0.24	
	Cup	0.2	
AvailableTime		11	
Order	Plate	4	
	Cup	8	
Make	(Product)	(Day)	(Amount)
	Plate	1	x1
	Cups	1	y1

Make	Plate		
	Cups		
TotalProfit			
TotalHours			40

Mathematical Programming

- All about maximizing (minimizing) value some function subject to constraints
 - Discrete vs continuous (integer, binary, continuous)
 - Forms of functions: linear, quadratic, non-linear and so on
 - Different classes have different algorithmic difficulty
- Constraints are expressed as inequalities featuring variables ranging over numeric quantities (e.g., x and y above)
- Objective function is an expression involving some subset of the variables
- Efficient and effective solvers exist for solving mixed integer (linear constraints, linear objective, and some variables integers)
 - Solve problems with (many many) millions of constraints and variables
 - Usually highly specialized (commercial) software
- Interfaced through **modeling languages**
 - How do we write families of related mathematical programming problems
 - For example, can we generalize to multiple products (plates, cups, etc.) and write single problem that depends on input values such as profits per product etc.

THE MAIN IDEA

LB Databases can store variables, not just values

Same Problem in LogicBlox

Rules:

```

objective[] += profit[p] * make[p].
total_time[] += time[p] * make[p].

```

Constraints

```

make[p] = v -> v >= order[p].
total_time[] = v
-> v <= available_time[].

```

And we declare the predicate **make** as a special kind of predicate

- The system assumes that any value that exists in the second column of make is an unknown (variable) which must obey the constraints
- The system will "fill in" the optimal values for the variables such that the objective function is maximized and the constraints satisfied

Profit	Plate	1.07
	Cup	1.05
Time	Plate	0.24
	Cup	0.2
AvailableTime		11
Order	Plate	4
	Cup	8
Make	Plate	x=5
	Cups	y=49

Change some values stored in predicates

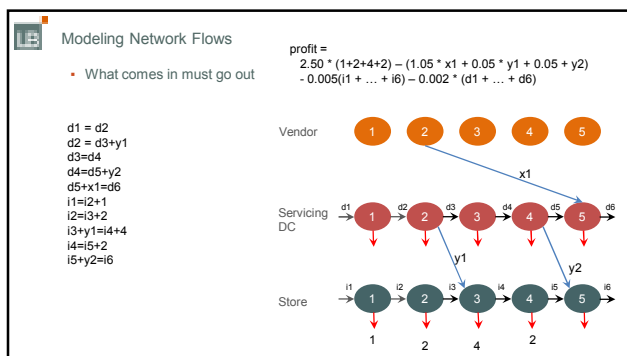
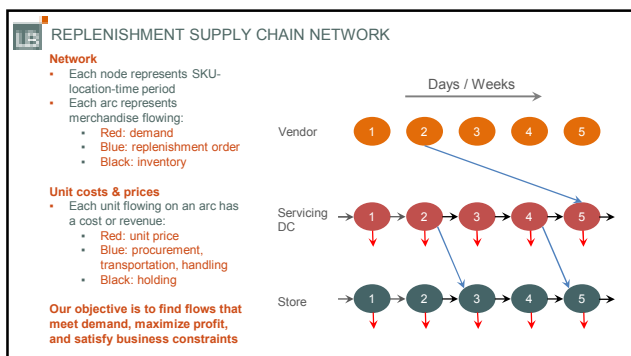
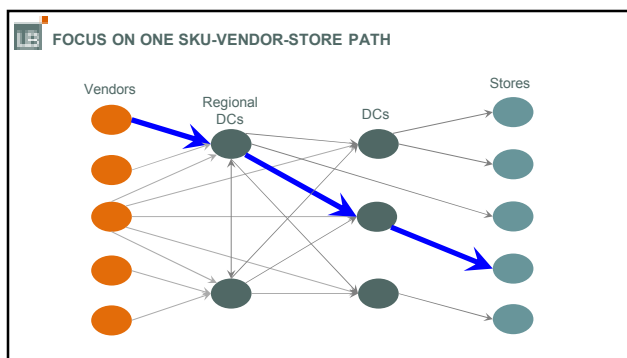
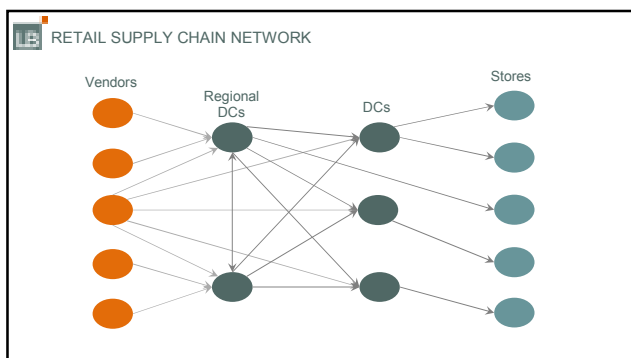
- The system will
 - Automatically re-synthesize a new system of inequalities
 - Invoke the solver to obtain new solution values
 - Update the predicate **Make**
- So that
 - At the end of the transaction that changed the values in red
 - The original semantics is maintained: the predicate **Make** still contains the optimal solution to the constraints that maximizes the objective function
- Benefits
 - Interoperation with solver 100% hidden from programmer
 - Support for multiple solver back ends
 - Built-in optimization of the instance part of the platform
 - Maintenance ensures that all rules, values etc. that depend on the predicate **Make** get recalculated to obtain consistency

Profit	Plate	0.3
	Cup	1.05
Time	Plate	0.24
	Cup	0.2
AvailableTime		42
Order	Plate	4
	Cup	8
Make	Plate	x=4
	Cups	y=164

And automatically get an updated solution

- We can even add more facts. Suppose we the potter can also make vases. We simply insert the relevant facts into the database
- Invariant
 - At any given time (i.e., at the end of any transaction) the database is guaranteed to contain a set of optimal values for the second column of the predicate Make or, if a solution cannot be found, the transaction is aborted, and the database restored to its previous state

Profit	Plate	0.3
	Cup	1.05
	Vase	6.2
Time	Plate	0.24
	Cup	0.2
	Vase	1.0
AvailableTime		42
Order	Plate	4
	Cup	8
	Vase	2
Make	Plate	x=4
	Cups	y=8
	Vase	z=39

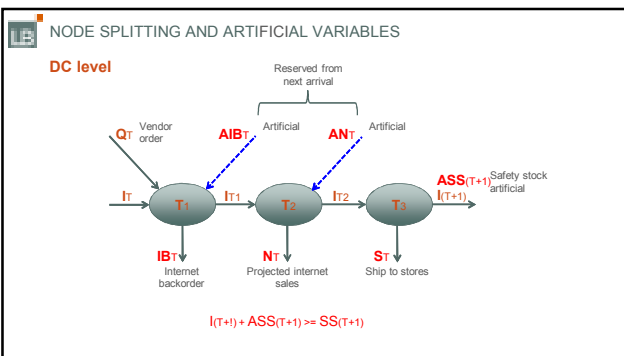


OBJECTIVE FUNCTION

- Maximize profit
- Revenue
 - Units sold (forecasting, seasonality, promotions)
 - Unit selling price (typically varies)
 - Lost sales = (shortages – backorders) * unit selling price
 - Transforming a portion of shortages into backorders
 - Loss of goodwill due to shortages
- Cost
 - Procurement cost = quantities bought * unit purchasing price (incorporate quantity discounts and special vendor agreements)
 - Transportation cost = shipped quantities * unit shipping cost (can incorporate shipping in full containers or trailers)
 - DC and store inventory holding cost
 - Warehouse handling costs
 - Capital cost
 - Penalty costs
 - Many more

CONSTRAINTS AND BUSINESS RULES

- Network infrastructure
 - Vendors, DCs, cross-docks, stores
 - Order placement dates
 - Lead times, truck schedules
- Shipping multiples (pack size)
- Flow conservation equations
- Service level, safety stock
- Display stock
 - Sellable, non-sellable
 - Primary, secondary
- Inventory should not fall below
 - Safety stock
 - Primary and secondary display
 - Soft constraints
- Prioritization in case of constrained supply
 - Backorders
 - Store and internet sales
 - Display stock
 - Safety stock



NUMBER OF VARIABLES

- Number of inventory variables alone for a typical medium size retail replenishment problem is huge:
 - Active SKUs: 50,000
 - Number of stores: 500
 - Number of periods: 100 (assuming modeling the first 6 weeks at the daily level and the following 420 days at the weekly level)
 - Number of inventory variables:

50,000 * 500 * 100 = 2.5 billion variables

In fact a lot larger (cca 20+ variables per node)

- Use of the cloud, parallelization, and powerful dis-aggregation and aggregation techniques enable us to formulate and solve very large problems

Pragmatics: the devil, the details

- So far, all of the above can be modeled pretty easily with modern MIP solvers
- We represent the structure of the network with predicates and rules
 - Turns out to be a very convenient formalism for specifying very large set of complicated constraints
- Creating problem instances is a challenge
 - Problem instances may be too large to fit into memory
 - Luckily, we happen to have a database that can do out-of-core computation:
 - Represent the instance itself as a set of predicates
 - Generate lots of LogiQL code (invisible to the user) to construct it and marshal it to the solver
 - Generate rules that rewrite the instance (decomposition, aggregation, variable elimination) to obtain smaller instances to solve
 - Incremental update of instances
 - If only a single order change update only a small portion of the instances and solver data structures
- Complex business rules are surprisingly easy to add
 - Often just extra LogiQL constraints:
 - Sum of all the fish ordered from the vendor during months which are marked as no fishing season should be 0
 - the total cubic footage ordered must be some multiple of the volume of a shipping container
 - you should never have more than 20 tons of fish at any distribution center
- Computation time becomes very expensive
 - Amazon EC2 instance with 80G of memory a lot more expensive than one with 24G and we need hundreds or thousands of those daily

SUPPLY CHAIN NETWORK OPTIMIZATION IN PRODUCTION

Input data

SKUs, DCs, stores, vendors, historical demand, DC calendar, store calendar, ordering DCs, initial conditions

Input Parameters


Demand forecast, price, procurement cost, transportation cost, holding cost, cost of capital, fill rate, safety stock

Solver

Nightly

Outputs

Orders, Reports, Charts



Conclusions

- Who would have thought that something as so boring sounding as "enterprise software" could be so much fun to work on?!
- Fancy Booklearnin': Do I use anything I learned along the way?
 - Mr Kunic, my elementary school teacher: $2x=4 \rightarrow x=2$
 - Undergraduate Evergreen [Quite a lot]
 - Judy's Database Class: I learned about datalog [In fact, Judy's advisor invented it]
 - C&C I learned about 1st Order Logic, models thereof, mathematical logic and Prolog [Al Leisenring]
 - OGI: Programming Languages, Functional Programming, Type Theory
 - More about Logic [basis of Datalog]
 - Much of what I do to make the Mathematical Programming work at LB is translation from one language to another
 - Rice: Postdoc
 - Started working on implementing dynamic programming and linear algebra numerical algorithms – came in handy with optimization
 - Theorem proving (Coq): still a dream – Can we use a proof assistant to write our constraints, and actually prove safety, coherence and convergence properties of our MP models a priori.

