Declarative Networking

Boon Thau Loo, Tyson Condie, Minos Garofalakis, David E. Gay, Joseph M. Hellerstein, Petros Maniatis, Raghu Ramakrishnan, Timothy Roscoe, and Ion Stoica. "Declarative Networking." *CACM* 2009

http://www.albinoblacksheep.com/text/copyright

A different kind of language …a different kind of class.

JEOPARDY!

- Choose a category.
- Given the answer, provide the correct question.
- Jun pays the winning team.

Some slide contents from Boon Thau Loo's presentation Declarative Networking: Extensible Networks with Declarative Queries

Jeopardy Template from http://teach.fcps.net/trt10/PowerPoint.htm

Datalog	Big Picture	Not So Social Networking	Making It Happen	Pot Pouri	Review Points
10	10	10	10	10	10
20	20	20	20	20	20
30	30	30	30	30	30
40	40	40	40	40	60
50	50	50	50	50	

Datalog is a restriction of this paradigm for which Prolog is considered the standard.

What is Logic Programming?

Also called predicates, these can appear in both the body and the head of Datalog Rules.

What are relation symbols?

Datalog rule syntax:

<result> ← <condition1>, <condition2>, … , <conditionN>.

Head **Body**

- Types of conditions in body:
  - Input tables: *link(src,dst)* predicate
  - Arithmetic and list operations
- Head is an output table
  - Recursive rules: result of head in rule body

The q in this Datalog rule.

T(x):- q(x,y)

What is a conjunction of relational atoms?

Datalog fixpoint semantics rules define the immediate consequence operator, which can be expressed in relational algebra, and does this.

What is mapping one idb (predicate at head of a rule) to another?

Applies rules to existing facts.

$F(T) = G \bowtie T \cup G$

Naïve evaluation, the simplest bottom-up evaluation strategy based on fixpoint semantics, does this.

What is apply the immediate consequence operator on all previous steps until a step yields nothing new?

Declarative networking has the potential to reduce this by orders of magnitude relative to traditional approaches.

What is program size?

Network protocols deal at their core with computing and maintaining this basic declarative networking observation.

What is distributed state?

This language approach opens opportunities for automatic protocol optimization and hybridization, program checking, and debugging.

What is declarative?	**Ndlog is explicitly missing these, often expressed as “send” or “receive”.**
What are communication primitives?	**NDlog builds on Datalog with location specifiers, which do this.**
What is provide explicit control over tuple storage location?	**NDlog must be able to enforce these constraints, locally.**

What are routing policies?

This distributed network protocol computes the shortest paths to all other nodes for each node.

What is path vector?

Nodes advertise entire path to a destination.
Node receives ad, adds self, and forwards to neighbors.

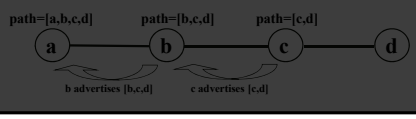

This network protocol of which BGP is an example, is a class of distance vector.

What is path vector?

```
sp1 path(@Src,Dest,Path,Cost) :- link(@Src,Dest,Cost),
    Path=f_init(Src,Dest).
sp2 path(@Src,Dest,Path,Cost) :- link(@Src,Nxt,Cost1),
    path(@Nxt,Dest,Path2,Cost2), Cost=Cost1+Cost2,
    Path=f_concatPath(Src,Path2).
sp3 spCost(@Src,Dest,min<Cost>) :- path(@Src,Dest,Path,Cost).
sp4 shortestPath(@Src,Dest,Path,Cost) :-
    spCost(@Src, Dest,Cost), path(@Src,Dest,Path,Cost).
Query shortestPath(@Src,Dest,Path,Cost).
```

NDlog example program implementing path-vector protocol.

Dynamic source routing, used in wireless mesh networks, requires these in its packets.

What are the addresses of each device in the source-chosen path?

Link state protocol implies that each router builds one of these.

What is a network map?

Each iteration updates and propagates these of increasing hop count in the below depiction of a shortest-path NDlog program.

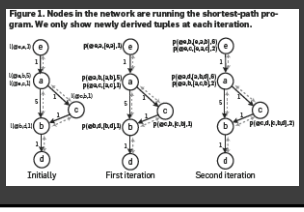

What are paths?

Semi-naïve fixpoint Datalog evaluation uses only these tuples from previous recursive iterations.

What are new tuples?

Centralized plan generation uses semi-naïve fixpoint evaluation in order to prevent this.

What is redundant evaluation?

Distributed plan generation requires a rule localization rewrite step because of this situation.

What happens when nonlocal rules have location specifiers that require joining tuples from different nodes?

Pipelined semi-naïve iteration addresses this problem in distributed plan generation.

What happens when completion of an iteration is difficult to determine in a distributed environment?

Overlog, a extension of Ndlog, includes a materialized keyword that allows TTL definition useful for this network approach to storing distributed state.

What is soft-state?

Link-restricted rules limit NDlog in order to capture these.

What are physical network constraints?

```
p(@Dest,...) :- link(@Src,Dest...),p1(@Src,...),
                p2(@Src,...),..., pn(@Src,...).
```

Body predicates all have same location specifier (@Src)

Declarative networking uses an optimization technique called Magic Sets, which does this.

What is limiting computation to relevant paths?

Overlog was able to describe a practical implementation of this overlay network based on a distributed hash table in two orders of magnitude less code.

What is Chord DHT?

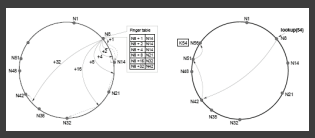

This plan generation mode supports continuous rule updates for a changing network topology and strives toward *eventual consistency*.

What is incremental maintenance?

The authors made it clear that NDlog can reduce code size, but at least six reviews mentioned that the paper was missing this.

What is a performance comparison to a traditional implementation?

Two reviewers thought that while NDlog saves development time, users will pay at least some of that time back here.

What is the compiler?

NDlog support for continuous rule execution assumes a network eventually quiesces for a time, which implies that some nodes will be in this state at some point.

What is inconsistent / unsychronized / invalid?

Magic sets rewriting and predicate reordering can do this to implementations of distance-vector and dynamic source routing.

What is make them indistinguishable?

FINAL JEOPARDY!

http://boemo.blogspot.com/2008/04/then-norm-showed-upthe-best-norm.html

Make your wager

The paper mentions many times that declarative networking (NDlog) is a natural fit for this method of query evaluation.

What is recursive?