

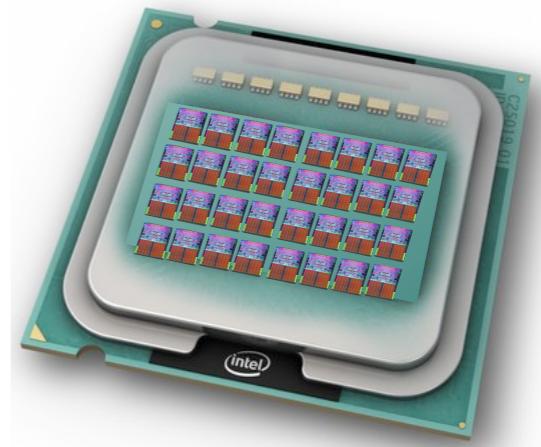
# **CHALLENGE 12**

**DECLARATIVE LANGUAGES  
PLAY THE FIELD**

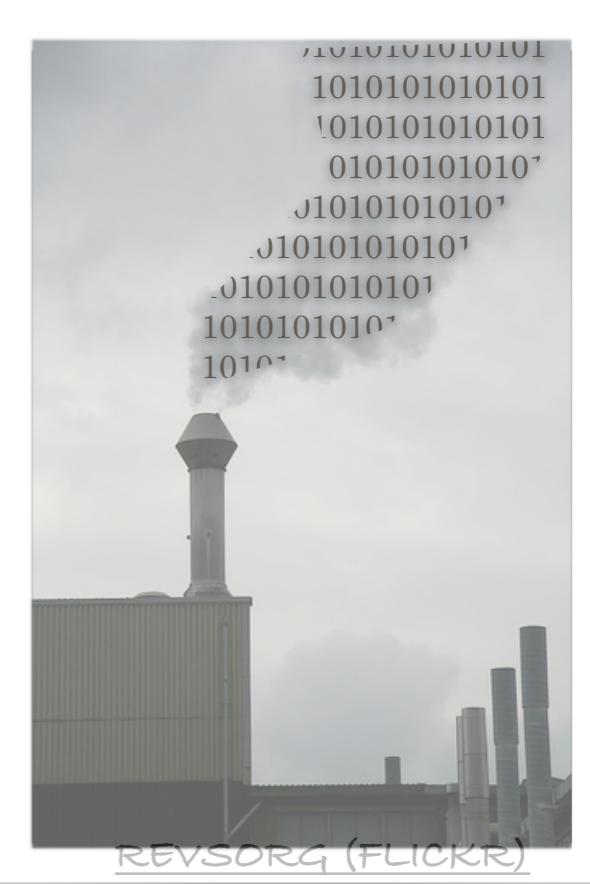
**JOSEPH M HELLERSTEIN**



# CHANGE AFOOT



- ✿ datacenters/virtualization
- ✿ manycore
- ✿ big data (industrial revolution!)
- ✿ => ready to revisit the whole SW stack



# JIM GRAY CHALLENGE #12

# JIM GRAY CHALLENGE #12



## Automatic Programming

Do What I Mean (not 100\$ Line of code!, no programming bugs)

The holy grail of programming languages & systems

### 12. Devise a specification language or UI

1. That is easy for people to express designs (1,000x easier),
2. That computers can compile, and
3. That can describe all applications (is complete).

- System should “reason” about application
  - Ask about exception cases.
  - Ask about incomplete specification.
  - But not be onerous.
- This already exists in domain-specific areas.  
(i.e. 2 out of 3 already exists)
- An imitation game for a programming staff.

# JIM GRAY CHALLENGE #12



## Automatic Programming

Do What I Mean (not 100\$ Line of code!, no programming bugs)

The holy grail of programming languages & systems

12. Devise a specification language or UI
  1. That is easy for people to express designs (1,000x easier),
  2. That computers can compile, and
  3. That can describe all applications (is complete).
  - System should “reason” about application
    - Ask about exception cases.
    - Ask about incomplete specification
    - But not be erroneous



## now add auto-parallelization

- This already exists in domain-specific areas.  
(i.e. 2 out of 3 already exists)
- An imitation game for a programming staff.



# TIMING!

- CS ready to revisit programming
- data + parallel + distributed
- we have their attention!
- and things are already happening...



# DATALOG? OMG!

- ✿ been getting around in last 5 years!
  - ✿ despite god-awful syntax
  - ✿ a “domain-specific” language that’s playing the field
  
- ✿ our work on *declarative networking*, plus:
  - ✿ compilers ([bddbddb](#), stanford '05)
  - ✿ natural language processing ([dyna](#), hopkins, '05)
  - ✿ security protocols (AT&T, Stanford, '01)
  - ✿ information extraction (wisc/yahoo '07)
  - ✿ modular robotics ([meld](#), cmu, '07)
  - ✿ machine learning (berkeley/cmu '08)

# P2: DECLARATIVE INTERNET OVERLAYS

- ✿ chord routing, with:
  - ✿ multiple successors
  - ✿ stabilization
  - ✿ optimized finger maintenance
  - ✿ failure detection
- ✿ 48 “overlog” rules

```

/* The base tuples */
materialize(node, infinity, 1, keys(1)).
materialize(finger, 180, 160, keys(2)).
materialize(bestSucc, infinity, 1, keys(1)).
materialize(succDist, 10, 100, keys(2)).
materialize(succ, 10, 100, keys(2)).
materialize(pred, infinity, 100, keys(1)).
materialize(succCount, infinity, 1, keys(1)).
materialize(join, 10, 5, keys(1)).
materialize(landmark, infinity, 1, keys(1)).
materialize(fFix, infinity, 160, keys(2)).
materialize(nextFingerFix, infinity, 1, keys(1)).
materialize(pingNode, 10, infinity, keys(2)).
materialize(pendingPing, 10, infinity, keys(2)).

/** Lookups */
watch(lookupResults).
watch(lookup).

l1 lookupResults@R(R,K,S,SI,E) :- node@NI(NI,N),
    lookup@NI(K,R,E), bestSucc@NI(S,SI),
    K in (N,S].
l2 bestLookupDist@NI(NI,K,R,E,min<D>) :- node@NI(NI,N),
    lookup@NI(K,R,E), finger@NI(I,B,BI),
    D:=K - B - 1, B in (N,K).
l3 lookup@BI(min<BI>,K,R,E) :- node@NI(NI,N),
    bestLookupDist@NI(K,R,E,D),
    finger@NI(I,B,BI), D == K - B - 1,
    B in (N,K).

/** Neighbor Selection */
n1 succEvent@NI(S,SI) :- succ@NI(S,SI).
n2 succDist@NI(S,D) :- node@NI(NI,N),
    succEvent@NI(S,SI), D:=S - N - 1.
n3 bestSuccDist@NI(min<D>) :- succDist@NI(S,D).
n4 bestSucc@NI(S,SI) :- succ@NI(S,SI),
    bestSuccDist@NI(D), node@NI(NI,N),
    D == S - N - 1.
n5 finger@NI(0,S,SI) :- bestSucc@NI(S,SI).

/** Successor eviction */
s1 succCount@NI(count<*>) :- succ@NI(S,SI).
s2 evictSucc@NI(C) :- succCount@NI(C), C > 2.
s3 maxSuccDist@NI(max<D>) :- succ@NI(S,SI),
    node@NI(NI,N), evictSucc@NI(NI), D:=S - N - 1.
s4 delete succ@NI(S,SI) :- node@NI(NI,N),
    succ@NI(S,SI), maxSuccDist@NI(D),
    D == S - N - 1.

/** Finger fixing */
f1 fFix@NI(E,I) :- periodic@NI(E,10),
    nextFingerFix@NI(I).
f2 fFixEvent@NI(E,I) :- fFix@NI(E,I).
f3 lookup@NI(K,NI,E) :- fFixEvent@NI(E,I),
    node@NI(NI,N), K:=I << I + N.
f4 eagerFinger@NI(I,B,BI) :- fFix@NI(E,I),
    lookupResults@NI(K,B,BI,E).
f5 finger@NI(I,B,BI) :- eagerFinger@NI(I,B,BI).
f6 eagerFinger@NI(I,B,BI) :- node@NI(NI,N),
    eagerFinger@NI(I,B,BI),
    I:=I + 1, K:=I << I + N,
    K in (N,B), BI != NI.
f7 delete fFix@NI(E,I1) :- eagerFinger@NI(I,B,BI),
    fFix@NI(E,I1), I > 0, I1 == I - 1.
f8 nextFingerFix@NI(0) :- eagerFinger@NI(I,B,BI),
    ((I == 159) || (BI == NI)).
f9 nextFingerFix@NI(I) :- node@NI(NI,N),
    eagerFinger@NI(I,B,BI),

I:=I + 1, K:=I << I + N, K in (B,N),
NI != BI.

/** Churn Handling */
c1 joinEvent@NI(NI,E) :- join@NI(NI,E).
c2 joinReq@LI(LI,N,NI,E) :- joinEvent@NI(NI,E),
    node@NI(NI,N), landmark@NI(NI,LI), LI != "-".
c3 succ@NI(NI,N,NI) :- landmark@NI(NI,LI),
    joinEvent@NI(NI,E), node@NI(NI,N), LI == "-".
c4 lookup@LI(LI,N,NI,E) :- joinReq@LI(LI,N,NI,E).
c5 succ@NI(S,SI) :- join@NI(S,SI),
    lookupResults@NI(K,S,SI,E).

/** Stabilization */
sb1 stabilize@NI(NI,E) :- periodic@NI(NI,E,15).
sb2 stabilizeRequest@SI(SI,NI) :- stabilize@NI(NI,E),
    bestSucc@NI(S,SI).
sb3 sendPredecessor@PI1(PI1,P,PI) :- stabilizeRequest@NI(NI,PI1),
    pred@NI(P,PI), PI != "-".
sb4 succ@NI(P,PI) :- node@NI(NI,N), sendPredecessor@NI(NI,P,PI),
    bestSucc@NI(S,SI), P in (N,S).
sb5 sendSuccessors@SI(SI,NI) :- stabilize@NI(NI,E),
    succ@NI(S,SI).
sb6 returnSuccessor@PI(PI,S,SI) :- sendSuccessors@NI(NI,PI),
    succ@NI(S,SI).
sb7 succ@NI(S,SI) :- returnSuccessor@NI(S,SI).
sb7 notifyPredecessor@SI(SI,N,NI) :- stabilize@NI(NI,E),
    node@NI(NI,N), succ@NI(S,SI).
sb8 pred@NI(P,PI) :- node@NI(NI,N), notifyPredecessor@NI(NI,P,PI),
    pred@NI(P1,PI1), ((P1 == "-") || (P in (P1,N))). 

/** Connectivity Monitoring */
cm0 pingEvent@NI(NI,E) :- periodic@NI(NI,E,5).
cm1 pendingPing@NI(NI,PI,E) :- pingEvent@NI(NI,E),
    pingNode@NI(PI).
cm2 pingReq@PI(PI,NI,E) :- pendingPing@NI(PI,E).
cm3 delete pendingPing@NI(PI,E) :- pingResp@NI(PI,E).
cm4 pingResp@RI(RI,NI,E) :- pingReq@NI(RI,NI,E).
cm5 pingNode@NI(NI,SI) :- succ@NI(NI,S,SI), SI != NI.
cm6 pingNode@NI(NI,PI) :- pred@NI(NI,P,PI), PI != NI, PI != "-".
cm7 succ@NI(S,SI) :- succ@NI(NI,S,SI), pingResp@NI(SI,E).
cm8 pred@NI(P,PI) :- pred@NI(NI,P,PI), pingResp@NI(P,PI).
cm9 pred@NI("-","-",") :- pingEvent@NI(NI,E),
    pendingPing@NI(PI,E), pred@NI(P,PI).

```

```

/* The base tuples */
materialize(node, infinity, 1, keys(1)).
materialize(finger, 180, 160, keys(2)).
materialize(bestSucc, infinity, 1, keys(1)).
materialize(succDist, 10, 100, keys(2)).
materialize(succ, 10, 100, keys(2)).
materialize(pred, infinity, 100, keys(1)).
materialize(succCount, infinity, 1, keys(1)).
materialize(join, 10, 5, keys(1)).
materialize(landmark, infinity, 1, keys(1)).
materialize(fFix, infinity, 160, keys(2)).
materialize(nextFingerFix, infinity, 1, keys(1)).
materialize(pingNode, 10, infinity, keys(2)).
materialize(pendingPing, 10, infinity, keys(2)).

/** Lookups */
watch(lookupResults).
watch(lookup).

l1 lookupResults@R(R,K,S,SI,E) :- node@NI(NI,N),
    lookup@NI(K,R,E), bestSucc@NI(S,SI),
    K in (N,S].
l2 bestLookupDist@NI(NI,K,R,E,min<D>) :- node@NI(NI,N),
    lookup@NI(K,R,E), finger@NI(I,B,BI),
    D:=K - B - 1, B in (N,K).
l3 lookup@BI(min<BI>,K,R,E) :- node@NI(NI,N),
    bestLookupDist@NI(K,R,E,D),
    finger@NI(I,B,BI), D == K - B - 1,
    B in (N,K).

/** Neighbor Selection */
n1 succEvent@NI(S,SI) :- succ@NI(S,SI).
n2 succDist@NI(S,D) :- node@NI(NI,N),
    succEvent@NI(S,SI), D:=S - N - 1.
n3 bestSuccDist@NI(min<D>) :- succDist@NI(S,D).
n4 bestSucc@NI(S,SI) :- succ@NI(S,SI),
    bestSuccDist@NI(D), node@NI(NI,N),
    D == S - N - 1.
n5 finger@NI(0,S,SI) :- bestSucc@NI(S,SI).

/** Successor eviction */
s1 succCount@NI(count<*>) :- succ@NI(S,SI).
s2 evictSucc@NI(C) :- succCount@NI(C), C > 2.
s3 maxSuccDist@NI(max<D>) :- succ@NI(S,SI),
    node@NI(NI,N), evictSucc@NI(NI), D:=S - N - 1.
s4 delete succ@NI(S,SI) :- node@NI(NI,N),
    succ@NI(S,SI), maxSuccDist@NI(D),
    D == S - N - 1.

/** Finger fixing */
f1 fFix@NI(E,I) :- periodic@NI(E,10),
    nextFingerFix@NI(I).
f2 fFixEvent@NI(E,I) :- fFix@NI(E,I).
f3 lookup@NI(K,NI,E) :- fFixEvent@NI(E,I),
    node@NI(NI,N), K:=I << I + N.
f4 eagerFinger@NI(I,B,BI) :- fFix@NI(E,I),
    lookupResults@NI(K,B,BI,E).
f5 finger@NI(I,B,BI) :- eagerFinger@NI(I,B,BI).
f6 eagerFinger@NI(I,B,BI) :- node@NI(NI,N),
    eagerFinger@NI(I,B,BI),
    I:=I + 1, K:=I << I + N,
    K in (N,B), BI != NI.
f7 delete fFix@NI(E,I1) :- eagerFinger@NI(I,B,BI),
    fFix@NI(E,I1), I > 0, I1 == I - 1.
f8 nextFingerFix@NI(0) :- eagerFinger@NI(I,B,BI),
    ((I == 159) || (BI == NI)).
f9 nextFingerFix@NI(I) :- node@NI(NI,N),
    eagerFinger@NI(I,B,BI),

I:=I + 1, K:=I << I + N, K in (B,N),
NI != BI.

/** Churn Handling */
c1 joinEvent@NI(NI,E) :- join@NI(NI,E).
c2 joinReq@LI(LI,N,NI,E) :- joinEvent@NI(NI,E),
    node@NI(NI,N), landmark@NI(NI,LI), LI != "-".
c3 succ@NI(NI,N,NI) :- landmark@NI(NI,LI),
    joinEvent@NI(NI,E), node@NI(NI,N), LI == "-".
c4 lookup@LI(LI,N,NI,E) :- joinReq@LI(LI,N,NI,E).
c5 succ@NI(S,SI) :- join@NI(S,SI),
    lookupResults@NI(K,S,SI,E).

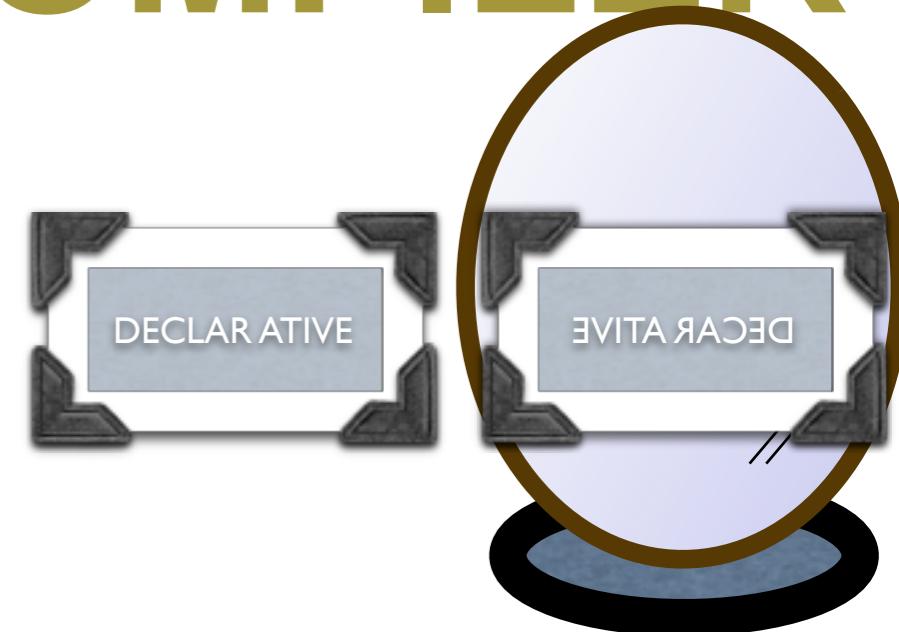
/** Stabilization */
sb1 stabilize@NI(NI,E) :- periodic@NI(NI,E,15).
sb2 stabilizeRequest@SI(SI,NI) :- stabilize@NI(NI,E),
    bestSucc@NI(S,SI).
sb3 sendPredecessor@PI1(PI1,P,PI) :- stabilizeRequest@NI(NI,PI1),
    pred@NI(P,PI), PI != "-".
sb4 succ@NI(P,PI) :- node@NI(NI,N), sendPredecessor@NI(NI,P,PI),
    bestSucc@NI(S,SI), P in (N,S).
sb5 sendSuccessors@SI(SI,NI) :- stabilize@NI(NI,E),
    succ@NI(S,SI).
sb6 returnSuccessor@PI(PI,S,SI) :- sendSuccessors@NI(NI,PI),
    succ@NI(S,SI).
sb7 succ@NI(S,SI) :- returnSuccessor@NI(S,SI).
sb7 notifyPredecessor@SI(SI,N,NI) :- stabilize@NI(NI,E),
    node@NI(NI,N), succ@NI(S,SI).
sb8 pred@NI(P,PI) :- node@NI(NI,N), notifyPredecessor@NI(NI,P,PI),
    pred@NI(P1,PI1), ((P1 == "-") || (P in (P1,N))). 

/** Connectivity Monitoring */
cm0 pingEvent@NI(NI,E) :- periodic@NI(NI,E,5).
cm1 pendingPing@NI(NI,PI,E) :- pingEvent@NI(NI,E),
    pingNode@NI(PI).
cm2 pingReq@PI(PI,NI,E) :- pendingPing@NI(PI,E).
cm3 delete pendingPing@NI(PI,E) :- pingResp@NI(PI,E).
cm4 pingResp@RI(RI,NI,E) :- pingReq@NI(RI,NI,E).
cm5 pingNode@NI(NI,SI) :- succ@NI(NI,S,SI), SI != NI.
cm6 pingNode@NI(NI,PI) :- pred@NI(NI,P,PI), PI != NI, PI != "-".
cm7 succ@NI(S,SI) :- succ@NI(NI,S,SI), pingResp@NI(SI,E).
cm8 pred@NI(P,PI) :- pred@NI(NI,P,PI), pingResp@NI(P,PI).
cm9 pred@NI("-","-",") :- pingEvent@NI(NI,E),
    pendingPing@NI(PI,E), pred@NI(P,PI).

```

# EVITA RACED: OVERLOG METACOMPILER

- ✿ go to the mirror, boy
  - ✿ would you build a system with that?
- ✿ overlog optimizer in ~150 overlog rules
  - ✿ system r + histograms + magic + stratification checks
  - ✿ exercise: system r => cascades?
    - ✿ 24 hours, 27 rules
  - ✿ porting to Pig@yahoo



Tyson Condie, et al.  
VLDB '08



# OUR LEVERAGE?

- ✿ we actually can auto-parallelize datalog
  - ✿ “just” dataflow, but well beyond embarrassing
- ✿ good for decentralized (datacenter) kernels
  - ✿ paxos & commit, DHTs, multicast, ...
- ✿ obviously good for queries, etc.
  - ✿ and we’re tackling fancy machine learning



# GETTING THERE

- ✿ reality check: everyone hates datalog
  - ✿ so toss is out!
- ✿ design a new language, for hackers
  - ✿ attractive syntax, tools
  - ✿ get the complexity in a box
  - ✿ integrate with the rest of the universe
  - ✿ go fast
  - ✿ exploit/sell what we know
- ✿ remember:
  - ✿ the language train only leaves once a decade!
  - ✿ our turn to drive

# CHALLENGE 12

you may say i'm a dreamer  
... but i'm not the only one