Version of July 8, 2009

Hyper Logic Programs in SILK For Business and Science: An Overview

Benjamin Grosof*

July 16, 2009
25-minute presentation at CULP 2009**

* Vulcan Inc., benjaming@vulcan.com
http://silk.semwebcentral.org
http://www.mit.edu/~bgrosof/

** Workshop on Commercial Users of Logic Programming, co-located with the 25th International Conference on Logic Programming (ICLP-2009), Pasadena, California
http://www.cs.kuleuven.be/~toms/CULP2009/

PROJECT HALO

VULCAN

Outline of Talk

- Intro to the SILK effort, and its parent Project Halo
- Hyper Logic Programs KR approach and expressive features
  - Higher-Order Defaults
  - Examples and Use Cases
  - Remedying FOL Semantics' Lack of Scalability
  - Comparison to other semantic rule systems and standards
    - RIF, BRMS, OWL, DBMS, etc.
- Conclusions and Directions
  - How You can be Involved

PROJECT HALO

VULCAN

SILK's ambitious Vision for longer-term Impact

- Key Knowledge Representation (KR) infrastructure sufficient to enable
  creation of global, widely-authored, very large knowledge bases (VLKBs)
  about science and business*
  that answer questions
  and proactively supply information,
  using powerful reasoning about rules and processes,
  that can be customized in their content and actions
  for individual organizations or people

- Newest part of Vulcan's Project Halo
  which addresses the problems of scale and brittleness in KBs,
  including the Knowledge Acquisition and UI aspects

* "Business" here is shorthand for human affairs, incl. government

PROJECT HALO

VULCAN

SILK Effort

- SILK = Semantic Inferencing on Large Knowledge
  - *What the next generation Web will be spun from*
- A KR Language and KR System with reasoner, UI, interchange
- Goal: Expressiveness + Semantics + Scalability + Web
- Begun in 2008
  - Part of Halo Advanced Research (HalAR), the new half of Project Halo
- Largest rule research program in the US (that we're aware of)
  - Primarily via contractors

PROJECT HALO

VULCAN

SILK Contributors current/past (partial list)

- Vulcan (Benjamin Grosof, Mark Greaves, Dave Gunning)
- Stony Brook University (Michael Kifer; students H. Wan, S. Liang, P. Fodor)
- SRI International (Vinay Chaudhri, David Martin, Ken Murray, Bill Jarrold)
- BBN Technologies (Mike Dean, Dave Kolas, Matt Rubin)
- Ontoprise GmbH (Daniel Hansch, Jurgen Angele)
- Automata (Paul Haley)
- Boeing (Peter Clark)
- Cycorp (Keith Goolsbey, Doug Lenat, Ben Rode)
- University of Texas (Bruce Porter, Ken Barker)
- University of Toronto (Sheila McIlraith; students S. Sohrabi, H. Ghaderi)
- University of Amsterdam (Bert Bredeweg)
- University of Freiburg (Georg Lausen)
- University of Michigan (Michael Wellman)
- Richard Fikes, consultant (Stanford University)
- Raphael Volz, consultant

Vulcan's Project Halo ; 1st system is AURA

- Vision of Digital Aristotle
  - Put the bulk of the world's scientific and similar knowledge on-line
  - Answer questions, act as personal tutor, with deep reasoning
- How to operationalize Digital Aristotle as a research effort?
- College-level science selected as initial domain focus
  - Medium wide, medium deep.
  - Good metrics available: textbook-type exam Q's. Initial domain task focus is:
    - Advanced Placement Exam (AP) in Physics, Chemistry, and Biology
      - Taken by USA high-school students to get credit for 1st-year college courses
- AURA expert system developed (see http://www.ai.sri.com/project/aura)
  - Novel combination of available techniques from AI
  - Controlled Natural Language, GUI, Frame-based KR, Problem-Solving
  - Students as users – formulate questions, formulate knowledge
  - Initial version 2004, then refined extensively and tested rigorously

PROJECT HALO

VULCAN

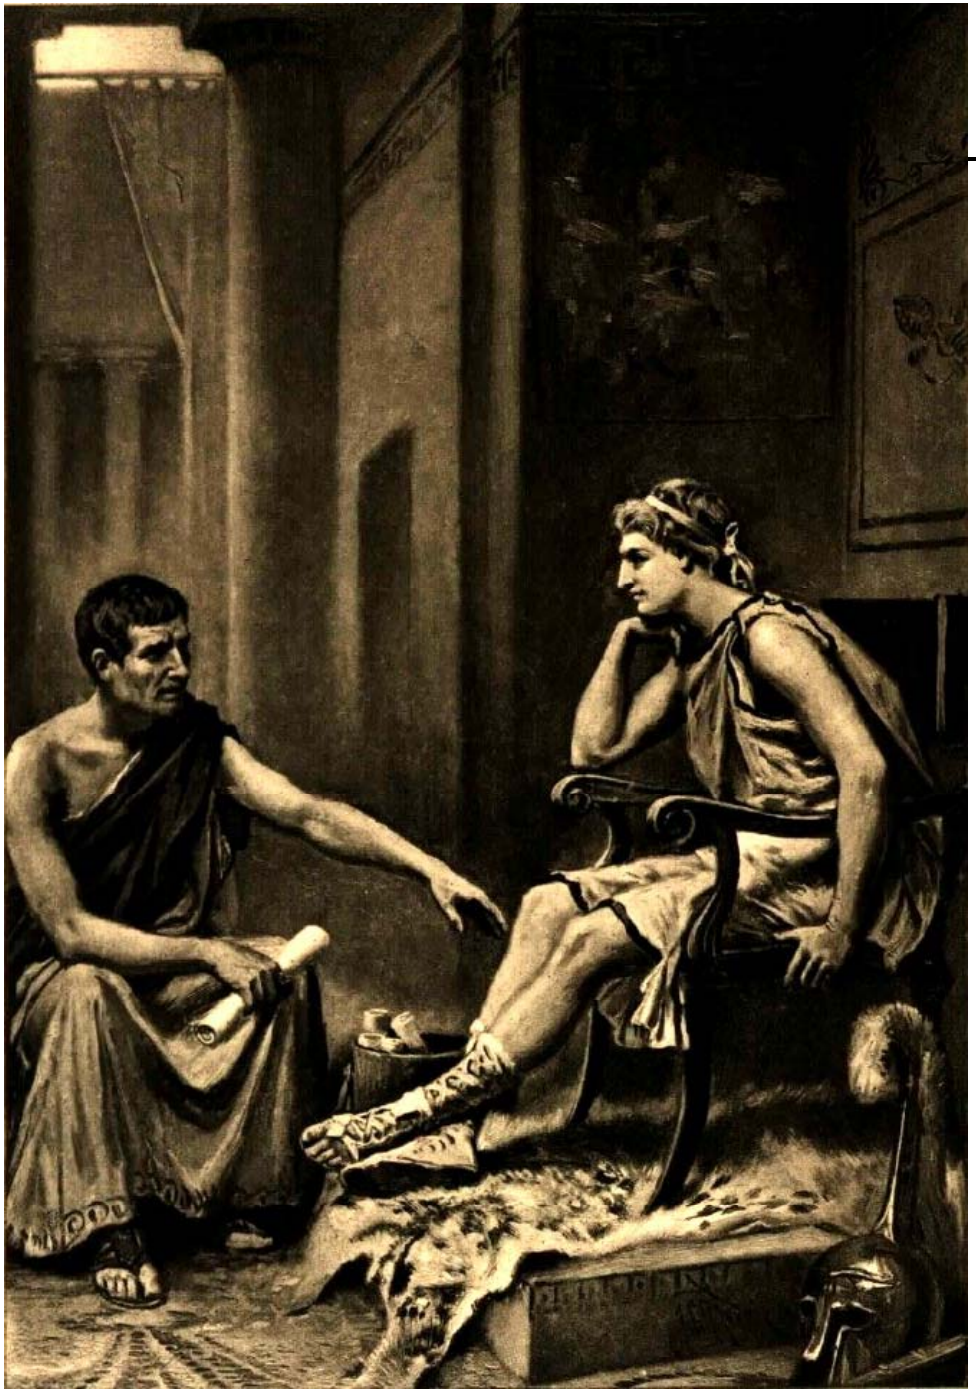

Aristotle Tutoring Alexander

VULCAN

Halo Enters Semantic Web Era; 2nd system is SMW+

- How to enable effective Knowledge Acquisition (KA)?
  + By Subject Matter Experts (SMEs), not programmers or knowledge engineers
  + Collaboratively – incorporate large #s of SMEs in KB construction & maintenance
  + Leveraging the Web
- Halo Extension to Semantic MediaWiki (SMW+) developed.
  - Open source extension of the MediaWiki software Wikipedia runs on
  - Supports RDF and OWL subset, interleaved tightly with hypertext
  - Rapid maturation of initial functionality
  - Standing queries. Data import/export. Plug-ins.
  - Upcoming release: simple semantic rules (Horn) and access control
  - Strong community uptake, early commercial adoption already
  - For more, see http://wiki.ontoprise.de
- But need better KR too, in part for sake of KA.
  - The underlying KR is the target for KA: "The KR is the deep UI"
  - Web knowledge interchange (with merging) for scalability of collaborative KA

PROJECT HALO

VULCAN

Goals for SILK KR Effort – Halo's 3rd system

- Expressiveness + Semantics + Scalability
  - Push the Frontier: high risk, high return
- Address requirements for AURA on AP task (& for SMW+)
  - Expressive power (e.g., defaults and processes)
  - Understandability via semantics and expressiveness
    - Raise abstraction level closer to the user's natural language and cognition
- Address requirements for long-term Digital Aristotle vision
  - Wider set of domains and tasks, via KR expressiveness and better KA
  - Knowledge interchange via semantics and expressiveness
  - Performance scalability of reasoning (incl. truth maintenance)

PROJECT HALO

VULCAN

Expressiveness "Brittleness" Areas Targeted

- Defaults/Exceptions/Defeasible *(incl. nonmonotonic reasoning, theory revision, argumentation, truth maintenance)*
  - A kinematics problem situation has standard earth gravity, and no air resistance. [physics AP]
  - A given organism has the anatomy/behavior that is typical/normal for its species, e.g., a bat has 2 wings and flies. [bio AP]
  - Price info for an airplane ticket on Alaska Air's website is accurate and up to date. [e-shopping]

  ❖ Practical reasoning almost always involves a potential for exceptions
- Hypotheticals
  - If Apollo astronaut Joe golfed a ball on the moon, then standard earth gravity would not apply. [negative hypothetical] [*conflict* between defaults, resolved by *priority* among them]
  - If I had swerved my car 5 seconds later than I did, I would have hit the debris in the left lane with my tire. [*counterfactual*]
- Actions and Causality
  - If a doorkey is incompletely inserted into the keyhole, turning the key will fail. [*precondition*]
  - During the mitotic stage of prometaphase, a cell's nuclear envelope fragments [biology AP]
  - After a customer submits an order on the website, Amazon will email a confirmation and ship the item. [Event-Condition-Action (*ECA*) rule] [policy]
- Processes (i.e., representing and reasoning about processes)
  - Mitosis has five stages; its successful completion results in two cells. [compose] [partial description]
  - If Amazon learns that it will take an unexpectedly long time to stock an ordered item, then it emails the customer and offers to cancel the order without penalty. [exception handling]
  - A Stillco sensor-based negative feedback thermal regulator is adequate to ensure the overnight vat fermentation of the apple mash will proceed within desired bounds of the alcohol concentration parameter. [science-based business process]

Ubiquitous in science, commonsense, business, etc. All are interrelated.

PROJECT HALO

VULCAN

Outline of Talk

- Intro to the SILK effort, and its parent Project Halo
- Hyper Logic Programs KR approach and expressive features
  - Higher-Order Defaults
  - Examples and Use Cases
  - Comparison to other semantic rule systems and standards
    - RIF, BRMS, OWL, DBMS, etc.
- Conclusions and Directions
  - How You can be Involved

PROJECT HALO

VULCAN

SILK's New KR: Hyper Logic Programs

- Hyper Logic Programs KR combines new features
  - Defaults and Weakened Classical, cf. generalized Courteous LP
    - Higher-order defaults with priorities, cancellation, contraposition, multi-way conflicts
    - Sound interchange with classical logic (via hypermonotonic mapping)
  - External Actions, Events, & Queries, cf. generalized Production LP
    - Via procedural attachments. Including built-ins.
    - Enables interoperation with Production/ECA rules (via SweetRules technique)
    - Brings Actions (and events) to the semantic party

  with previous advanced features
  - Higher-order and Frames, cf. Hilog and F-Logic
  - Webized syntax and Knowledge Interchange, cf. RIF/RuleML and OWL/RDF
  - Closed-World, cf. well-founded unstratified NAF
  - Good Efficiency of reasoner performance. With persistency & truth maintenance.
  - Equality, Lloyd-Topor, Aggregation, Functions, Skolemization, Integrity Constraints

PROJECT HALO

VULCAN

SILK – Summary

- Status: prototype engine, language, and theory for expressive heart
  - V1 adds Higher-Order Defaults to Flora
  - Extensive requirements analysis, use cases, benchmarking; ReCyc translation
  - V2 in development: adds Java API, Actions/Events, Interchange with RIF and Classical, …
- Radically extends the KR power of W3C OWL, SPARQL, and RIF – and of SQL
  - Defaults and robust conflict handling – *cope with knowledge quality and context*
  - Higher-order and flexible meta-reasoning – *elevate meta-data to meta-knowledge*
  - Actions and events, cf. production rules and process models – *activate knowledge*
- Raises the KR abstraction level for business users (SMEs) and NL KA/UI
- Use cases in business policies, ontology mapping, e-commerce, biomed, …
- Redefining the KR playing field for semantic web, business rules, and rule-based process management
  - Defaults and Higher-Order – yet retain computational web scalability
  - Escape from FOL's Extreme Brittleness – yet retain grade-AAA model-theoretic semantics

PROJECT HALO

VULCAN

More Rationale about Hyper LP KR

- Give up reasoning by cases
  - Source of exponential worst-case complexity in classical, disjunctive LP, stable LP
  - Can hope to reintroduce in restricted or altered form, or develop work-arounds, later
  - But there are many apps not requiring it, e.g., DBMS, BRMS

- Can realistically hope to be web-scalable performance-wise, unlike highly expressive classical
  - Polynomial computational complexity, under non-onerous restrictions
    - Same complexity as Horn rules!! (Must be careful of recursion through functions.)
  - Many optimizations available
  - Established track record of high scalability for relational databases

PROJECT HALO

VULCAN

New Theory & Algorithms for Higher-Order Defaults

- Combines Courteous + Hilog, and generalizes
- New approach to defaults: "argumentation theories"
  - Meta-rules specify when rules are defeated
  - [Wan, Grosof, Kifer, *et al.* ICLP-2009]
- Extends straightforwardly to combine with other key features
  - E.g., Frame syntax, external Actions
- Significantly improves on previous Courteous approach in other ways
  - Eliminates a complex transformation
  - Much simpler to implement
    - 20-30 background rules instead of 1000's of lines of code
  - Much faster when updating the premises
  - More flexible control of edge-case behaviors
  - Much simpler to analyze theoretically

PROJECT HALO

VULCAN

SILK Current Status – More

- New approach to representing causal change in processes
  - Uses defaults
- Use cases, incl. survey
  - Science AP
  - Business domains
- ReCyc: Rough prototype translator from Cyc to SILK
  - 3 Million axioms from ResearchCyc (translates 99% of the KB)
- Benchmarking of relevant rule systems
  - OpenRuleBench [Liang *et al.* WWW-2009]
- SILK V2 is in development. Near term steps include:
  - Add expressive features, e.g., Weakened Classical, external Actions
  - Webize more fully, e.g., knowledge interchange, UI

PROJECT HALO

VULCAN

Ecology Ex. of Causal Process Reasoning in SILK

```
/*    Toxic discharge into a river causes fish die-off.    */
/* Init. facts, and an "exclusion" constraint that fish count has a unique value */
  occupies(trout,Squamish).
  fishCount(s0,Squamish,trout,400).
  !- fishCount(?s,?r,?f,?C1) and fishCount(?s,?r,?f,?C2) | ?C1 != ?C2.
/* Action/event description that specifies causal change, i.e., effect on next state */
  @tdf1  fishCount(?s+1,?r,?f,0) :- occurs(?s,toxicDischarge,?r) and occupies(?f,?r).
/* Persistence ("frame") axiom */
  @pef1  fishCount(?s+1,?r,?f,?p) :- fishCount(?s,?r,?f,?p).
/* Action effect axiom has higher priority than persistence axiom */
  @pr1  overrides(tdf1,pef1).
/* An action instance occurs */
  @UhOh  occurs(s0+1,toxicDischarge,Squamish).
```

As desired: |= fishCount(s0+1,Squamish,trout,400) and
fishCount(s0+2,Squamish,trout,0).

Notes: @ prefixes a rule label. ? prefixes a variable. :- means if. !- prefixes an exclusion, and means "it's a conflict if". In an exclusion, | means given that.

E-Commerce Ex. of Causal Process Reas. in SILK

```
/*    E-commerce delivery logistics. */
/* Initial fact, and prevention constraint that location is unique */
  loc(s0,PlasmaTV46,LasVegasWH).
  !- loc(?s,?item,?posn1) and loc(?s,?item,?posn2) | ?posn1 != ?posn2.
/* Action/event description that specifies causal change, i.e., effect on next state */
  @mov1  loc(?s+1,?item,?addr) and neg loc(?s+1,?item,?warehouse) :-
         shipment(?s,?item,?warehouse,?addr) and loc(?s,?item,?warehouse).
/* Persistence ("frame") axioms about location */
  @pel1  loc(?s+1,?item,?posn) :- loc(?s,?item,?posn).
  @pel2  neg loc(?s+1,?item,?posn) :- neg loc(?s,?item,?posn).
/* Action effect axiom has higher priority than the persistence axioms */
  overrides(mov1,pel1).   overrides(mov1,pel2).
/* An action instance occurs */
  @deliv57  shipment(s0+1,PlasmaTV46, WH_LasVegasNV, 9_Fog_St_SeattleWA).
```

As desired: |= loc(s0+2,PlasmaTV46, 9_Fog_St_SeattleWA) and
neg loc (s0+2,PlasmaTV46, WH_LasVegasNV).

Notes: @ prefixes a rule label. ? prefixes a variable. :- means if. !- prefixes an exclusion, and means "it's a conflict if". In an exclusion, | means given that.

Trust Mgmt. Ex. of Higher-Order Defaults in SILK

illustrating also basic Knowledge-level Communication, and Frame syntax

In Frame syntax: subject[property -> object] *stands for* property(subject,object).

```
/*    Trust policy administration by multiple agents, about user permissions */
/* Admin. Bob controls printing privileges including revocation (neg). */
  Bob[controls -> print].   Bob[controls -> neg print].    /* neg print means it's disallowed.*/
  Cara[controls -> ?priv].  /* Cara is the most senior admin., so controls all privileges. */
/* If an administrator controls a privilege and states at a time (t) that a user has a privilege,
   then the user is granted that privilege. Observe that ?priv is a higher-order variable. */
  @grant(?t) ?priv(?user) :- ?admin[states(?t) -> ?priv(?user)] and ?admin[controls(?priv)].
/* More recent statements have higher priority, in case of conflict.  */
  overrides(grant(?t2), grant(?t1)) :- ?t2 > ?t1.
/* Admin.'s Bob and Cara make conflicting statements over time about Al's printing */
  Cara[states(2007) -> print(Al)].   Cara[states(2007) -> webPage(Al)].
   Bob[states(2008) -> neg print(Al)].
```

As desired: |= neg print(Al). webPage(Al).

/* Currently, Al is permitted a webpage but not to print. */

Notes: @ prefixes a rule label. ? prefixes a variable. :- means if. !- prefixes an exclusion, and means "it's a conflict if". In an exclusion, | means given that.

Use Cases for SILK beyond commercial state of art

- There are many!
- Existing use cases from SILK's research-y or standards-design roots
  - E.g., from RIF, RuleML, SWSL documents and prototypes
  - E-commerce, financial, health, trust, SOA, policies, regulations, mobile, biomed, defense, etc.
  - Many of these are not yet implementable in current well-supported, well-performing commercially deployed systems
    - E.g., they use defaults
    - E.g., they use feature combinations that are not easily available

PROJECT HALO

VULCAN

Outline of Talk

- Intro to the SILK effort, and its parent Project Halo
- Hyper Logic Programs KR approach and expressive features
  - Higher-Order Defaults
  - Examples and Use Cases
  - Remedying FOL Semantics' Lack of Scalability
  - Comparison to other semantic rule systems and standards
    - RIF, BRMS, OWL, DBMS, etc.
- Conclusions and Directions
  - How You can be Involved

PROJECT HALO

VULCAN

Remedying FOL Semantics' Lack of Scalability

- Hyper LP handles conflict robustly
  - Whereas FOL is a "Glass Bubble" – it's perfectly brittle semantically in face of contradictions from …
  - Quality problems/errors in the data and knowledge
  - Conflict when merging KBs

    E.g., OWL beyond the RL subset suffers this problem

A VLKB with a million or billion axioms formed by merging from multiple Web sources, is unlikely to have zero KB/KA conflicts from:

- Human knowledge entry/editing
- Implicit context, cross-source ontology interpretation
- Updating cross-source
- Source trustworthiness

- *Weakening provides a critical advantage for VLKB scalability*
  - *semantically, as well as computationally*

PROJECT HALO

VULCAN

FOL: A Glass Bubble

Extreme sensitivity to conflict limits its scalability in # of axioms and # of merges

Keijo Kopra from Finland as he competes in the Iittala Cup glass-blowing competition June 7, 2008. (Reuters)

Features Comparison – More Systems & Stds

Level	Groups of Features	*SILK1*	*Flora*	*RIF-BLD*	*Jena*	*Onto-broker*	*Jess*	*IBM C.R.*	*DLV*	*SQL*	*SPA-RQL*	*Common Logic*	*OWL2 RL*	*OWL2 DL*
Basic	Horn chain. etc.	Y	Y	Y	Y	Y	Y	Y	Y	R.	R.	Y	R.	R.
Advanced	*(Level summary)*	Most!	lots	some	some	some	some	some	some	some	some	some	some	some
	Equality	Y	Y	Y	R.	R.	R.	N	Y	R.	R.	Y	R.	Y
	Functions	Y	Y	Y	N	N	N	Y	Y	N	N	Y	N	N
	Frames etc.	Y	Y	R.	R.	R.	R.	R.	R.	R.	R.	R.	R.	R.
	Closed-World	Y	Y	N	N	Y	R.	R.	most	R.	R.	N	N	N
	Higher-Order	Y	Y	N	N	N	R.	N	N	R.	R.	Y	R. bit	R. bit
	Actions	Dev.	N	N	N	N	Y	Y	N	N	N	N	N	N
	Base Defaults	Y	N	N	N	N	N	Y	N	N	N	N	N	N
	Webized	Dev.	R.	Y	Y	R.	R.	R.	R.	N	Y	Y	Y	Y
Hyper	*(Level summary)*	1st!	N	N	N	N	N	N	N	N	N	N	N	N
	H-O. Defaults	Y	N	N	N	N	N	N	N	N	N	N	N	N
	Weak. Classi.	Dev.	N	N	N	N	N	N	N	N	N	N	N	N
Misc.		NA	NA	NA	NA	NA	NA	NA	NA	NA	NA	NA	NA	NA
	Other Expres.	Dev.	inherit.	-	-	-	events	-	disju.	R.	R.	classical	-	classic.
	Efficiency	good	good	NA	fair	good	fair	poor	good	NA	NA	NA	NA	NA

- Summarizes detailed analysis of 40 KR expressive features, 17 systems.
- Notes: Dev. = Developing, R. = Restricted; C.R.=Common Rules; disju.=disjunctive.

24

Outline of Talk

- Intro to the SILK effort, and its parent Project Halo
- Hyper Logic Programs KR approach and expressive features
  - Higher-Order Defaults
  - Examples and Use Cases
  - Comparison to other semantic rule systems and standards
    - RIF, BRMS, OWL, DBMS, etc.
- Conclusions and Directions
  - How You can be Involved

PROJECT HALO

VULCAN

Future Directions for SILK

- Process – more complex
- Natural Language KA and UI
- Parallelism in reasoning
- Connectors to Semantic Web, legacy BRMS and DBMS
- Uncertainty, Disjunction

- And Use Cases, of course

- Halo is part of an increasingly-integrated strategy at Vulcan to invest in semantics and advanced knowledge tools
  - Other investments: Radar Networks, ZoomInfo, Evri, etc.
  - Semantic MediaWiki+ is an early spinout

PROJECT HALO

VULCAN

Impact Opportunities for SILK and HalAR

- Improve by orders of magnitude:
  - Scale of practical semantic default+actions reasoning
    - <~1000 rules ⇒ ?100,000+ rules
  - Collaboration costs of multifold KB merging when there's conflict (as is usual)
    - Can take human out of the loop at run time
  - Population of users capable of specifying semantic rules
    - "KR Power to the People!" Leverage Aura and SMW+ KA/UI front-ends.
- Synergize best of last 20 years of pure-research progress in LP KR
  - ⇒ Redefine KR playing field of semantic web, business rules, & process management
- Provide a key missing research piece for SOA / web services
  - Enable building shared business/govt KBs on processes & policies ⇒ virtuous circle
- Hope: be like advance of the Relational model in DBMS
  - Will Hyper LP be to the 2010s what Relational was to 1970s-80s?

Key KR infrastruct. for widely-authored VLKBs for science and business that answer questions, proactively supply information, and reason powerfully

PROJECT HALO

VULCAN

How You can be Involved

- General Contact: Benjamin Grosof benjaming@vulcan.com
  - Suggest design, use cases, experts, cooperations
- Visit the SILK webpage and sign up for the mailing list so you'll be alerted of announcements about SILK
  - URL: http://silk.semwebcentral.org
  - Mailing list: silk-announce@semwebcentral.org (very low volume)
- Provide comments on SILK language design
  - Initial public draft in ~ fall 2009
  - Plan to propose a RIF extension with defaults and actions
    - Corresponding to a large expressive subset of SILK
- Try out SILK software
  - Prototype, free for research use
  - V1 public release in ~ fall 2009; V2 in 2010; selected earlier users sooner
  - Also SMW+ upcoming release will have simple semantic LP rules of SILK-y flavor
    - In ~ fall 2009. Limited to Horn.

PROJECT HALO

VULCAN

Acknowledgements

- SILK contributors
  - (previously listed)

- Contributors to several key previous KR efforts
  - RuleML and SWSL (Semantic Web Services Language) standards designs
  - SweetRules and Flora-2 systems

- Especially:
  - Michael Kifer, Mark Greaves, Dave Gunning, Mike Dean, Hui Wan, Paul Haley, Vinay Chaudhri, David Martin, Keith Goolsbey, Harold Boley, Said Tabet, and Guizhen Yang

PROJECT HALO

VULCAN

SILK –
Transforming Knowledge

Thank You

Disclaimer: The preceding slides represent the views of the author only.
All brands, logos and products are trademarks or registered trademarks of their respective companies.