

Dr. Dobb's interviews CEO Jans Aasman – Catching Up with Lisp

Dr. Dobb's: Jans, when I hear “Franz” I can't help but add “Lisp”. Is Lisp still at the core of what Franz is up to these days?

Aasman: I know you journalists want clarity, so my clear answer is a resounding “yes” and “no”.

Yes #1, we are still investing in our Lisp (Allegro CL). We recently released a new version, 8.2, and are planning v9.0 for later this year. We run on more platforms than ever, and we have our updated IDE running on the Mac. Also, real soon we'll announce our SMP (Symmetric Multi Processing) version of Allegro CL to automatically take advantage of all the multicore processors in a system.

Yes #2. AllegroGraph is the most ambitious internal projects we have ever worked on at Franz. It is a robust commercial graph database that is, in our humble opinion, the most scalable graph database for RDF data available, and it is all written in Lisp. So we are eating our own dog food. Actually, by doing this we have made many low level advancements in Allegro CL to improve data manipulation capabilities for “Big Data”. It is important to note that we own the entire technology stack which allows us to make optimizations all the way down to the hardware level.

No. What we see is that AllegroGraph is now more than half of our business, so in that sense you could say AllegroGraph is a new core.

Dr. Dobb's: Did Lisp play a part of the recently announced Franz-I0 Informatics partnership to deliver “fit for purpose” applications created by a Semantic application framework?

Aasman: No, not really. We used some Lisp in the project, just to do data manipulation but for the most part the work was done by AllegroGraph and the IO Informatics tools.

Dr. Dobb's: Why is Lisp so ideally suited to "semantics" such as with the Semantic web?

Aasman: There are several ways to answers this question, let me try a few.

- The main argument would be that Lisp is still the premier language for Artificial Intelligence and the Semantic Web is a sane and more modest revival of AI.
- The first applications written in Lisp were semantic networks. People would write knowledge processing applications and expert systems that were always built on top of a semantic network. Some of these semantic networks were more like frame-based system, where a frame is a thing that has a type (say 'person') and a bunch of arbitrary attributes (say 'first-name' 'last-name' 'dob', etc). These frame systems had, from day one, the capability to add any arbitrary attribute. The whole semantic web is just like a frame system. You have things that have a type and then a bunch of (arbitrary) attributes and then some logic around it to make it smart. Lisp was good at frames then and now with modern hardware it is even better at it now.
- One other reason that our Lisp is ideally suited for semantics is that we have a built in Prolog in our Lisp. Prolog is a declarative language build around logic and it is on many levels more powerful and practical than all the other logics invented by the semantic web community. We can use Prolog as a perfect rule-based system and query language for our triple store. So in our Lisp the rule and query language was built in from day one.

As an aside, many of the founders of the Semantic Web are Lisp

people.

- The first triple stores were written in Lisp (Ora Lassila from Nokia Labs).
- Jim Hendler and Dean Allemang wrote the book: the working Ontologist. Jim Hendler was also the driving force behind DAML that later became OWL
- Deborah McGuinness who was on all the committees that lead ultimately to ontologies.
- The first scalable reasoners for the Semantic Web were all written in Lisp (think of Fact by professor Ian Horrocks, think of Racer by Prof Haarslev and Prof Moller)
- etc, etc.
- Also see the paper The Use of Lisp in Semantic Web Applications.

Dr. Dobb's: What's new with Lisp?

Aasman: Do you mean with Lisp in general or the Lisp that Franz provides? Let me take the second interpretation: From our (biased) Franz perspective here are the exciting things:

- Symmetric Multi Processing
- The IDE and graphics system now run on all major platform, including the Mac
- You can debug Lisp programs visually in a stepper
- When you download our free version of Allegro CL you get (also for free) a professional object-oriented database (AllegroCache) and semantic web database (AllegroGraph).

Dr. Dobb's: What is a "dynamic object caching database"

Aasman: Dynamic object caching means that the language makes sure that the objects are in memory when you need them (and more importantly), get rid of them when you don't need them anymore. We also sell an object-oriented database (AllegroCache) that is totally integrated in the language itself. Not some wrapper around a relational database but a

completely native object-oriented database that is completely embedded in the language. So you can program as if objects are in memory but they are actually on disk.

Dr. Dobb's: Thanks for your time.