

가 가 . 가 ,
(
)
,
C++ 가 . Clojure
,
,
,
(F#
) 가 C#
(Microsoft .Net
(C#
) 가
,
,
,
C#
C#
FP FP 3)

Who Do You Love?

가 가 Lisp, Prolog Smalltalk 가 . 2000
가 Common Lisp[Steele Jr et al. 1990] 가
,
,
가 “ ”
, Common Lisp
가 Clojure
가 ,
가
가 “ ”
“ ? ”
” 4)
2005 ISV가 C/C++ Java .Net ' '
(: Ruby on Rails)가 JVM CLR
, Common Lisp
,
, C++ 가
(C++) . Common Lisp
SQL
. 2005 Clojure JVM Common Lisp 가
JVM Common Lisp DotLisp[Hickey 2003],
CLR Lisp, CL JVM jFli[Hickey
2004], API Foil[Hickey and Thorsen 2005] . CL JVM IPC
Clojure가 5)

- Clojure

1)

The objective for Clojure can be summarized most succinctly as: I wanted a language as acceptable as Java or C#, but supporting a much simpler programming model, to use for the kinds of information system development I had been doing professionally. I started working on Clojure in 2005, during a sabbatical I funded out of retirement savings. The purpose of the sabbatical was to give myself the opportunity to work on whatever I found interesting, without regard to outcome, commercial viability or the opinions of others. One might say these are prerequisites for working on Lisps or functional languages. I budgeted for two years of self-directed work, and Clojure was one of two projects I pursued. After about a year I decided the other project (a cochlear modeling and machine listening problem) was more of a research endeavor that might require two to five more years, so I dedicated myself at that point to getting Clojure to a useful state. I announced and released the first version of Clojure, as an open source project, in the fall of 2007. I did 100% of the implementation work during this time, and more than 90% through the 1.2 release in 2010. Subsequent to release Clojure benefited greatly from the feedback, suggestions and effort of its community. I am accepted by the community as “benevolent dictator for life” (BDFL) and continue to make all decisions relating to its evolution. Clojure is full of the great ideas of others, but I alone take responsibility for its faults.

2)

I had been doing commercial software development since 1987, almost always as the development lead and primary architect, first in C++, then Java and C#, as was common then and, in testament to institutional inertia, is still now thirty years later. I worked on scheduling systems, broadcast automation, yield management, audio recognition, exit poll tabulation and election projection et al. I would broadly characterize my work, and the work most commonly done by professional programmers, as information systems programming. Most developers are primarily engaged in making systems that acquire, extract, transform, maintain, analyze, transmit and render information—facts about the world. Most often, this information documents some human activity, be that of customers, suppliers, advertisers, travelers, voters, members, students, patients etc and must deal with all the irregularity thereof. This is in stark contrast to artificial systems, e.g., programming language compilers, which make up their own rules, in fully enumerated spaces, can eliminate irregularity and can reject anything which does not conform. Information system programmers have the thankless task of attempting to superimpose some- what regular models over information and real-world activity that refuses to comply. For instance, in music scheduling, trying to decide: whether artists have songs or songs have artists. Or optimizing a scheduler to spread out the plays of songs by each artist, except on ‘twofer Tuesdays’ when songs by each artist must be played in adjacent pairs. In information systems programming, twofer Tuesdays are everywhere.

3)

By the mid 1990’s I was a C++ expert, taught advanced C++ as an adjunct at NYU, and was a proponent and advocate of the benefits of static typing (but neglected the tradeoffs, sorry students!). I was happily discovering type parameterization tricks [Hickey 1996], running a const-correct shop etc. However, over time, in my experience, the suitability-to-task of these statically typed class models for information systems programming was quite low, and the benefits of the type checking minimal, especially in addressing the number one actual problem faced by programmers: the overwhelming complexity inherent in imperative, stateful programming. As programs grew large, they required increasingly Herculean efforts to change while maintaining all of the presumptions around state and relationships, never mind dealing with race conditions as concurrency was increasingly in play. And we faced encroaching, and eventually crippling, coupling, and huge codebases, due directly to specificity (best-practice encapsulation, abstraction, and parameterization notwithstanding). C++ builds of over an hour were common. In the years immediately preceding work on Clojure I worked on

the system to be used for the national exit poll in the U.S. This system involved automating large statistical models. Early on we decided that an imperative approach to doing the stats was a misfit, being not mathematical and also due to the expected concurrency throughout the system. After exploring F# for the task (and finding it insufficiently expressive), we decided to code the stats in C# like the rest of the system (this was a Microsoft .Net shop). I designed some custom immutable data structures whose use would be co-aligned with the figures in the statistical specifications (so statisticians who did not know C# could look at the code and understand it), and a functional library for manipulating the structures and doing the calculations. This was a great success, yielding much simpler code that we did not, and still do not, worry about. However, the resulting C# code was, in the eyes of both new and experienced C# developers, bizarre and non-idiomatic. Thus functional programming was a win, but FP in a non-FP language was not something I could see being widely applied.

4)

I had always been a language geek, playing with Lisp, Prolog and Smalltalk in my spare time. When I became an independent consultant in 2000, I spent more time with Common Lisp [Steele Jr et al. 1990] and wrote a couple of real systems in it. It was a revelation. Huge layers of unnecessary complexity simply vanished. I had the flexibility to use exactly as much language as was needed for the problem. The percentage of code directly related to the domain increased. Development was much faster, the resulting program was more general and easier to change. It was impossible to avoid the sinking feeling that I had been “doing it wrong” by using C++/Java/C# my whole career. I needed another choice more like Common Lisp or I wouldn’t be able to continue as a professional software developer. While language features matter, the primary hurdle to language adoption by professionals is acceptability to developers and stakeholders. Thus, Clojure did have to be at minimum practical for developers I would work with, and companies I might work for, or else I couldn’t use it to make a living. I took this as an agenda item for its design but not as motivation to seek their input in advance, because at the time the advice from professional developers about Lisp was “it’s dead” and about functional programming was “what’s that?”, and about writing a new programming language was “you’re crazy”.

5)

The year 2005 may have represented the nadir of language diversity in commercial software development, with ISVs primarily using C/C++ and many businesses considering themselves either Java or .Net ‘shops’. This was before the ascent of dynamic languages for web development (e.g., Ruby on Rails), and many companies were reluctant to deploy and operate anything that didn’t run on the JVM or CLR. I did commercial work in Common Lisp twice - a scheduling system and a yield management system. The first time the program had to be rewritten in C++ in order to be acceptable for deployment, an arduous task (though I was a seasoned C++ programmer) that took longer than the initial development, was much more code, and did not run significantly faster. In the second case again Common Lisp was not acceptable for deployment, so I designed the program to generate SQL stored procedures for delivery to the client and runtime execution. Prior to embarking on Clojure in 2005, I had made several attempts to create a bridge between the JVM and Common Lisp, in order to increase the latter’s practicality and acceptability. These were DotLisp [Hickey 2003], an interpreted Lisp with host interop for the CLR, jFli [Hickey 2004], a library that embedded a JVM in CL, and Foil [Hickey and Thorsen 2005], a library that exposed a similar API but used IPC between the CL runtime and the JVM. None of these yielded production-level solutions, but they definitely informed Clojure’s ultimate host syntax and fed the idea of Clojure being hosted.

From:

<http://moro.kr/> - **Various Ways**

Permanent link:

<http://moro.kr/open/a-history-of-clojure>

Last update: **2022/08/30 08:11**

