

REFPERSYS high-level goals and design ideas*

Basile STARYNKEVITCH[†] Abhishek CHAKRAVARTI[‡]

Nimesh NEEMA[§]

refpersys.org

October 2019 - May 2021

Abstract

REFPERSYS is a **REF**lexive and orthogonally **PERS**istent **SYSTEM** (as a GPLv3+ licensed free software¹) running on Linux; it is a hobby² but serious **research project** for many years, mostly aimed to experiment **open science** ideas close to **Artificial General Intelligence**³ dreams, and we don't expect useful or interesting results before several years of hard work.

audience : LINUX free software developers⁴ and computer scientists interested in an experimental open science approach to reflexive systems, orthogonal persistence, symbolic artificial intelligence, knowledge engines, etc....

Nota Bene: this report contains many [hyperlinks](#) to relevant sources so its [PDF](#) should rather be read on a computer screen, e.g. with [evince](#). Since it describes a circular design (with many [cycles](#) [[Hofstadter:1979:GEB](#)]), we recommend to read it twice (skipping footnotes and references on the first read).



This entire document is licensed under the Creative Commons Attribution-ShareAlike 4.0 International License. To view a copy of this license, visit creativecommons.org/licenses/by-sa/4.0/ or send a letter to Creative Commons, PO Box 1866, Mountain View, CA 94042, USA.

*This document has git commit `fb17387fbbb7e200`, was Lua- \LaTeX generated on *2021-May-17 18:55 MEST*, see gitlab.com/bstarynk/refpersys/ and its `doc/design-ideas` subdirectory. Its draft is downloadable, as a [PDF](#) file, from starynkevitch.net/Basile/refpersys-design.pdf...

[†]See starynkevitch.net/Basile/ and contact basile@starynkevitch.net, 92340 Bourg La Reine (near Paris), France.

[‡]chakravarti.avishek@gmail.com, FL 3C, 62B PGH Shah Road, Kolkata 700032, India.

[§]nimeshneema@gmail.com, 206 Sundaram Apartment, 38/2, Biyabani, Indore 452002, India.

¹Some code is available on gitlab.com/bstarynk/refpersys.

²[Basile Starynkevitch](#) (France) wants to find some research grant funding related to this. Please mention potential funding opportunities (call for research project proposals) by email to basile@starynkevitch.net.

³[Artificial General Intelligence](#)

⁴Those LINUX software developers are routinely *glancing inside, building* then using -from their published source code- quite large open source programs (such as [GCC](#), [SBCL](#), [CHICKEN-SCHEME](#), [HOP](#), [HAXE](#), [OCSIGEN](#), [EMACS](#), [SQLITE](#), [MARIADB](#), etc...) and perhaps even contributing to smaller free software projects like [NINJA](#), [libonion](#), etc... By the way, all these open source projects could be useful to or inspirational for REFPERSYS.

Contents

1	Social Necessity of AGI Systems with Long Term Development	2
2	REFPERSYS ambitions and goals	3
2.1	REFPERSYS core idea [1] ² s	3
2.2	REFPERSYS strange development cycle	7
2.3	REFPERSYS persistent heap	9
2.4	Agenda and multi-threading in REFPERSYS	12
2.5	Metaprogramming and introspection in REFPERSYS	14
3	The data and object models of REFPERSYS	17
3.1	how data should be processed in REFPERSYS	17
3.2	data at the low and high levels	18
3.2.1	values and quasi-values	18
3.2.2	implementation details	18
3.3	immutable values	20
3.3.1	immutable scalar values	21
3.3.2	immutable composite values	21
3.4	mutable objects	22
3.4.1	objects as frame-like data	22
3.4.2	concrete examples of objects	23
3.4.3	object payloads	25
3.5	File naming	25
3.6	Building <code>refpersys</code> executable	26
3.7	REFPERSYS workflow	26
3.8	the REFPERSYS object model	27
3.8.1	REFPERSYS inheritance graph	28
3.8.2	REFPERSYS message sending	29
4	Persistence in REFPERSYS	30
4.1	The textual data format of REFPERSYS	31
4.2	EBNF Grammar of Data Format	31
5	Metaprogramming in REFPERSYS	32
6	The primordial Read-Eval-Print-Loop of REFPERSYS	33
7	The Web interface of REFPERSYS	34

1 Social Necessity of AGI Systems with Long Term Development

Our complex, but fragile, world is facing dramatic and extremely challenging planet-wide issues, such global warming, demographic and political crises, economic and financial emergencies, and growing inequalities. In the light of such challenges, Artificial General Intelligence (AGI) systems are increasingly relevant. **@@TODO: explain how?**

As the slow, progressive Darwinian evolution of human intelligence shows, the limited intelligence of the *Homo Sapiens*⁵ species took more than a million years (about 30,000 generations) to continually evolve from an ape-like state.

Our observation of natural human intelligence (which has not yet been fully understood or modelled⁶) has led us to believe that there is no single, simple model of intelligence. Similarly, any AGI system must necessarily have a very complex and self-improving organisation.

We are aware than any progress towards AGI will be slow (many years, perhaps decades⁷) and progressive. Remember **Hofstadter's Law**: *"It always takes longer than you expect, even when you take into account Hofstadter's Law"* [**Hofstadter:1979:GEB**] and Brook's observations [**Brooks:1987:NSB**, **Brooks:1995:MM**] that *"if one woman can give birth in 9 months, 9 women cannot give birth to a baby in one month"*. For "giving birth" to REFPERSYS, a small team could need at least 9 years. However, intermediate results or side effects are not predictable but could be useful even during the REFPERSYS project.

We believe in free software (read also [this](#)), and we strongly believe that an AGI prototype should be some free software, exactly like most infrastructure software are (notably LINUX). See also the [SOFTWARE HERITAGE project](#) for interesting insights. REFPERSYS wants to be an AGI infrastructure, and there is work for many years (several years of work needed without any "artificial intelligence", just for the infrastructure).

An even partially successful AGI system might be useful to coordinate, run and manage other existing software (described through some knowledge given declara-

⁵In Latin, *Homo Sapiens* means "the human who knows what it knows" and, interestingly enough, relates to both [metaknowledge](#) and [Reflection](#).

⁶Half a billion euros of European taxpayers' money were spent on the [Human Brain Project](#), but did not lead to a complete, reproducible, artificial model of human intelligence; of course, it did fund interesting and successful research!

⁷An interesting parallel could be controlled nuclear fusion -which also bears some "bootstrapping" concepts- with [ITER](#); we expect REFPERSYS to cost several thousand times less at least; but even partial AGI success is as important for humanity as nuclear fusion produced electricity, and a future REFPERSYS might even help that ITER [megaproject](#) or other ones.

tively). Imagine how complex future [digital twins](#) of the entire planet Earth, designed to tackle with global warming, would need to be. For such dramatically complex usage, an AGI system (like REFPERSYS, if we succeed in making it) could be quite helpful to just drive and use such a “digital twin” simulation. Making it free software runnable on a free software operating system should benefit most of humanity (but keeping it proprietary won’t), and enable further or alternative experimentations. And “[there is no planet B](#)”⁸. So investing a few persons willing to working for nearly a decade is not too much for such a perspective.

2 REFPERSYS ambitions and goals

2.1 REFPERSYS core idea [1] ?s

The title of this subsection is *not* a typo⁹. We indeed mean both *ideas* (that is, software design and architectural concepts, guiding our daily implementation efforts) and *ideals* (that is, long term research objectives and ambitions).

The REFPERSYS¹⁰ system shares several -but not all- goals and design ideas (but no code) with [bismon](#) [[Starynkevitch:2019:bismon-draft](#)] but of course *not* [bismon](#)’s application¹¹ to [static source code analysis](#). Like [bismon](#), REFPERSYS is a **reflexive** (it uses [reflection](#)), **introspective** and **orthogonally persistent** system, but not for [static program analysis](#). Please read [Bismon](#)’s draft report [[Starynkevitch:2019:bismon-draft](#)] for a more precise definition of these concepts. **REFPERSYS is a long term¹² risky research project with an open science mindset and reproducible experiment ethics** [[zuboff:2015:big-other](#), [oneil:2016:weapons](#)], and a **free software licensed under GPLv3+, and targetted only for LINUX x86-64 computers.** A Linux system¹³ with at least 16 Gibytes of RAM, 4 *x86-64* cores, and 220 Gibytes of disk is required. The grand ambition of REFPERSYS is to become later an infrastructure for some strong AGI

⁸As reminded E.Macron, president of France, to the US Congress.

⁹It is a geeky pun on words with [shell globbing](#) and [regexpr](#) like syntax.

¹⁰For a **Reflexive Persistent System**

¹¹I Basile am not allowed and not funded to directly work on AGI -which still is my major personal scientific interest- but I do get funded on applied research projects like [DECODER](#) and try to push some AGI ideas into them.

¹²I don’t expect any significant AGI research results before \approx 2026.

¹³My own `ours.starynkevitch.net` computer, running *Debian/Unstable*, has 64 Gibytes of RAM, 24 cores (AMD 2970WX) and terabytes of disk space, including a terabyte of SSD.

system à la CAIA¹⁴ by Jacques Pitrat¹⁵ [Pitrat:1996:FGCS, Pitrat:2009:AST, Pitrat:2009:ArtifBeings], but before even approaching that goal a big lot of work is required, and REFPERSYS should be valuable by itself for other less ambitious and more pragmatcal purposes, perhaps some specialized collaborative web server (GPLv3+) to ease communication between human REFPERSYS developers, that is a mix of a wiki, a chat, and a tool for sharing document with drawings or graphics.

The development of REFPERSYS is (like the one of `bismon`, or of CAIA) a slow, incremental and gradual bootstrapping process with a meta-programming [dormoy:1992:meta, hernandez-phillips:2019:debugging-bootstrap] approach : features added to REFPERSYS in January 2020 are used to implement new features worked on a later REFPERSYS in March 2020.

As every practical software, REFPERSYS targets some defined machines: common Linux distribution running on some computer¹⁶. So the target machine of REFPERSYS is a quite complete and modern Linux system (such as a recent DEBIAN or UBUNTU desktop), with many useful packages, and administered by some human person¹⁷. The REFPERSYS system is published in “source” form, as a set of `git` versioned¹⁸ textual files (e.g. hopefully generated C files¹⁹, perhaps some `Makefile` or better yet an `OMAKE` build -most and more and more²⁰ of them being generated- or shell files or data files). Some of these files are generated, and the bootstrapping goal is to have *every* `git`-registered textual file been generated by REFPERSYS, with

¹⁴With explicit permission from J.Pitrat, CAIA source code -entirely generated by itself, about half a million lines of C code- is available on my (Basile’s) web page as `caia-su-24feb2016.tar.bz2`, and you could build it with `gcc -O -g [A-Z]*.c -rdynamic -ldl` then run `./a.out`. However, since I Basile sadly failed to convince J.Pitrat that open source [Lerner-Tirole:2000:economics-open-source, Weber:2004:SuccessOpenSource] software are -in our XXIth century- also an important way to transmit research ideas, there are no complete instructions to use it. Hence CAIA has an undocumented user interface as user-friendly as the one of `ed` but convenient enough to J.Pitrat alone! If you are capable of reading some comments in French and guessing the semantics of declarative “expert system” like rules (CAIA has more than a dozen of thousands of them), run it, then type `L EDITE` and start reverse-engineering that brilliant CAIA system.

¹⁵Jacques Pitrat has passed away on October 14th, 2019. See quickly also his old web page on `jacques.pitrat.pagesperso-orange.fr` and his interesting blog on `bootstrappingartificialintelligence.fr/WordPress3...`

¹⁶For several years, that computer is a desktop or powerful laptop running some DEBIAN. Later that could be some “virtual machine” e.g. some `DOCKER` container.

¹⁷For obvious cybersecurity reasons, automatic administration of that Linux distribution is out of scope. Also, since Basile Starynkevitch is still working (in October 2019) in a cybersecurity lab (of about 25 permanent staff) at `CEA/LIST`, cybersecurity concerns would be a conflict of interest.

¹⁸We crucially depend upon `git` specifically (e.g. `GitLab`), and porting REFPERSYS to some other versioning system -or to some other operating system than LINUX- would be a quite difficult task.

¹⁹However, notice that bootstrapped language implementations like `Scheme 48` or `OCaml` are keeping some `bytecode` form under version control, and `CHICKEN SCHEME` is, like `bismon`, `git`-keeping generated C files.

²⁰Of course, in a `chicken and egg` fashion, the initial version of REFPERSYS has to contain mostly hand-written files!

a **bootstrapped** approach²¹ similar to those of **self-hosting compilers**.

Within REFPERSYS, we call²² “source file” any Linux file which is `git`-versioned. We hope that more and more of these source files will be generated by the `refpersys` **ELF executable** program. **A significant milestone is the entire bootstrapping of REFPERSYS**, when all files (in textual form, to stay `git`-friendly, like **text based protocols** are more friendly for developers) can be regenerated by the `refpersys` executable, exactly in the same state as they were previously²³ : as a whole, our REFPERSYS system should become a **Quine program**, and CAIA is already one. So the **build automation** tool which compiles REFPERSYS should use file contents, not modification times to trigger compilation commands, since a full regeneration of such a bootstrapped REFPERSYS system will touch all files, without changing the content of any of them. Hence and very concretely, for building REFPERSYS the `OMake` build automation tool is preferable to GNU `make`.

For pragmatical reasons, **REFPERSYS needs a good garbage collector** (or GC [**appel:1991:garbage**, **wilson:1992:uniprocessorgc**, **baker:1995:cons**, **jones:2016:gchandbook**]), since fully compile-time GC [**mazur:2004:compile**] are too difficult to implement. Since multi-core x86-64 machines are very common, it should take advantage of them, so **REFPERSYS should follow a multi-threaded approach** above POSIX [**barney:2010:pthreads**] or **C++11 threads**. Our GC should be a **precise** garbage collector [**Rafkind:2009:PreciseGC**] and we may want to favor, like what was done in GCC MELT [**Starynkevitch:2007:Multistage**, **Starynkevitch-DSL2011**, **Starynkevitch-GCCMELTweb**], fast allocation of small memory zones which get quickly disposed of when becoming dead using a copying generational **Cheney-like GC algorithm** [**wilson:1992:uniprocessorgc**]. But mixing precise, sometimes generational GC techniques with multi-threading is a difficult programming task. But precise-GC friendly programming is simpler in generated C or C++ code than with hand-written code (because of explicit management of local GC roots and write bar-

²¹Observe that Linux source distributions like linuxfromscratch.org, or to a lesser extent **GenToo**, are also, when considered as a single system, fully bootstrapped.

²²Notice that, on purpose, our terminology is different of usual habits in the open source realm: almost all software projects (see also softwareheritage.org) are made of *computer files* typed by human developers in some **source-code editor** or some **IDE** such as `Emacs`, `vim` or `Code::Blocks`, according to the old **Unix philosophy**. Notice that large open source projects like the **LIBREOFFICE** suite, the **GCC** compiler collection or the **Firefox** browser tend to accept **plugins** instead of favoring old fashioned **command pipelines**, but multi-threaded applications may follow the **pipeline design pattern**. In contrast, we are impatient to reach the state where all REFPERSYS source files have been `git`-versioned but are all generated by a previous run of our `refpersys` executable. The REFPERSYS developer is interacting, through a web interface, with some running `refpersys` process, which is also some specialized web server (using HTTP).

²³Pedantically, some **fixpoint** of some very coarse-grained **operational semantics** related to **abstract interpretation** and **big step semantics**, each big step being the entire regeneration of the system, inspired by Futurama projections and **partial evaluation**.

riers, à la [QISH](#) or [OCAML](#): garbage collection invariants are boring and brittle to maintain in hand-written code).

[Reification](#) is an important concept in REFPERSYS, including (later) at the [knowledge representation](#) level with [semantic networks](#) and [frames](#). REFPERSYS [call stacks](#) are made of call frames known to our garbage collector (like [OCAML](#)'s ones). They could later be copied into data structures representing some [delimited continuations](#) [[Reynolds:1993:continuations](#), [Queinnec:2004:ContinWeb](#)], perhaps even representing and describing control [[fouet-starynkevitch:describing-control:1987](#), [Starynkevitch-1990-EUM](#), [Pitrat:2009:ArtifBeings](#)]. This should also enable [introspection](#), by permitting primitives inspecting the current call stack, perhaps using Ian Taylor's [libbacktrace](#). Also, such an introspection might perhaps be implemented [[mitchell:2001:alp](#)] with two nearly twin `refpersys` processes, one of them driving a `gdb` process²⁴.

REFPERSYS should (like CAIA and its predecessor MALICE did [[Pitrat:2009:AST](#), [Pitrat:1996:FGCS](#), [Pitrat:2009:ArtifBeings](#)]) have some expert system shell [[kumar:2015:importance-expo](#), [nigro:2008:meta](#)] and meta-rules to “dynamically compile” some subset of expert system rules and knowledge bases to procedural code (e.g. with a metaprogramming approach of generating *C* code, or `libgccjit` compiled code, then `dlopen(3)`-ing that code and running it at runtime. The `manydl.c` program show that this can practically be done many dozen of thousands of times on Linux desktops).

REFPERSYS will extensively use [metaprogramming](#) techniques, so it **should generate code** (like CAIA do) in a [transpiler](#) approach (in C, C++, -compiled into [plugins](#) and later [dynamically loaded](#) with `dlopen(3)` - maybe also JavaScript and HTML5 if we decide to have a web user interface). REFPERSYS could also later use [just-in-time compilation](#) libraries such as `libgccjit`. The domain-specific language of REFPERSYS²⁵ (a declarative one, with “expert system rules”) should gradually increase its expressiveness and become more and more declarative and closer to mathematical formalisms.

Most Linux distributions contain lots of useful libraries or software components for REFPERSYS long-term goals, notably machine learning open source libraries like [TENSORFLOW](#) [[charniak:2019:deep-learning](#)] or [GUDHI](#) [[chazal:2016:high](#)]. We might at some point also need messaging libraries like [OMQ](#), graphical user interfaces libraries à la [QT](#) or more probably web servicing libraries like [libonion](#) or [WT](#). To decrease efforts, we don't want to rewrite such libraries inside REFPERSYS (considered as a very high level, [declarative, domain-specific language](#)). Hence, we

²⁴Imagine some `popen` or some `g_spawn_async` or some `Poco::Process` of some `gdb refpersys 1234` process debugging the other one of pid 1234.

²⁵That domain-specific language has to be defined and implemented in a bootstrapped manner.

will need in REFPERSYS to generate some glue code, like [SWIG](#) does, from some **declarative description** (probably some frames or knowledge bases) of the [API](#) of these available libraries.

REFPERSYS should at first be **orthogonally persistent**. Like [BISMON](#) [[Starynkevitch:2019:bismon-draft](#)] it will load its state (its entire garbage-collected [heap](#)) from files at startup, and will dump its state²⁶ into files at shutdown. These state files are textual, in [JSON](#) format, and `git`-versioned, and should be portable to other 64 bits Linux computers. A [manifest file](#) describing the collection of files keeping the state is probably needed.

2.2 REFPERSYS strange development cycle

Ordinary software projects tend to follow a spiral development model [[boehm:1988:spiral](#)] as shown in figure 1. But REFPERSYS' development follows a [strange loop](#)

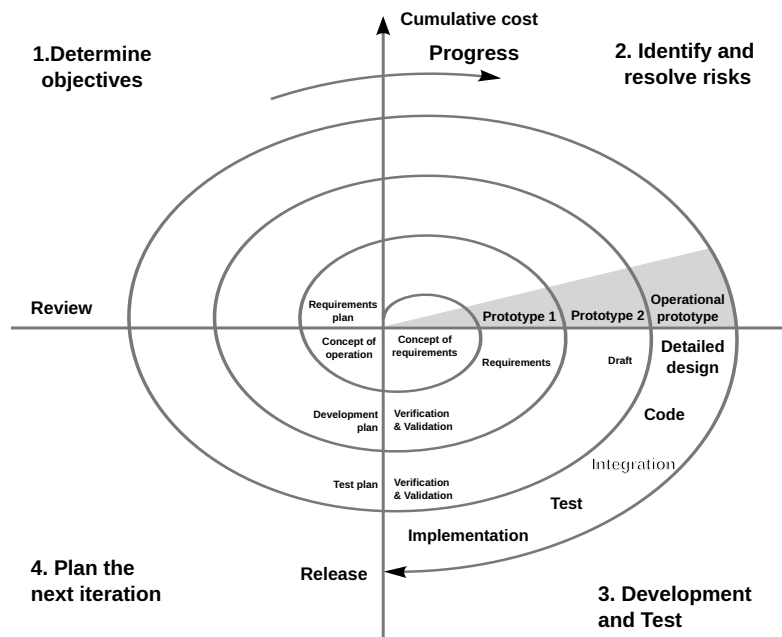



Figure 1: the **traditional spiral development** model (from Wikipedia [spiral model](#))

[[hofstadter:2007:strange-loop](#)], since it is bootstrapped in an [evolutionary prototyping](#) manner. It is more like a spiral staircase like in figure 2. The initial (floor) is just a persistent system, and we gradually add new code implementing more features (first entirely hand-written, later more and more parts of it replaced by REFPERSYS generated code). Of course the fun is in replacing existing hand-written

²⁶In a manner inspired by SBCL `save-lisp-and-die` primitive, or [POLYML](#) `export` primitive, or [marshalling facilities](#) of OCAML or PYTHON `pickle` module.

code (or low-level DSL) by more expressive and generated one. So we will continuously rewrite past formalizations as a more clever and expressive ones, taking more and more advantage of REFPERSYS whole-system introspective abilities. All of [EURISKO](#) [[Lenat:1983:Eurisko](#)], [CYC](#) [[Lenat:1991:ev-cycl](#)] and [SELF](#)²⁷ [[chambers:1991:efficient](#)] (or even [IO](#) or [SMALLTALK](#)) systems and their incremental development process are inspirational.



Each new feature -or small incremental change or a few of them (small `git` commits) - of REFPERSYS enables us to build and **generate** the next version of REFPERSYS, and a next feature is then added to that *improved* version, and so on repeatedly, etc....

Figure 2: the strange **REFPERSYS staircase development model** (from a [figure of Spiral stairs](#) by Lluïsa Iborra from the Noun Project)

The first significant milestone of REFPERSYS should be the ability to re-generate all its textual source files (and maybe even `git add` then `git commit` them). That would require first implementing some simple template based machinery²⁸, with the ability, like [QUINE](#) programs do, to regenerate all REFPERSYS source code (e.g.

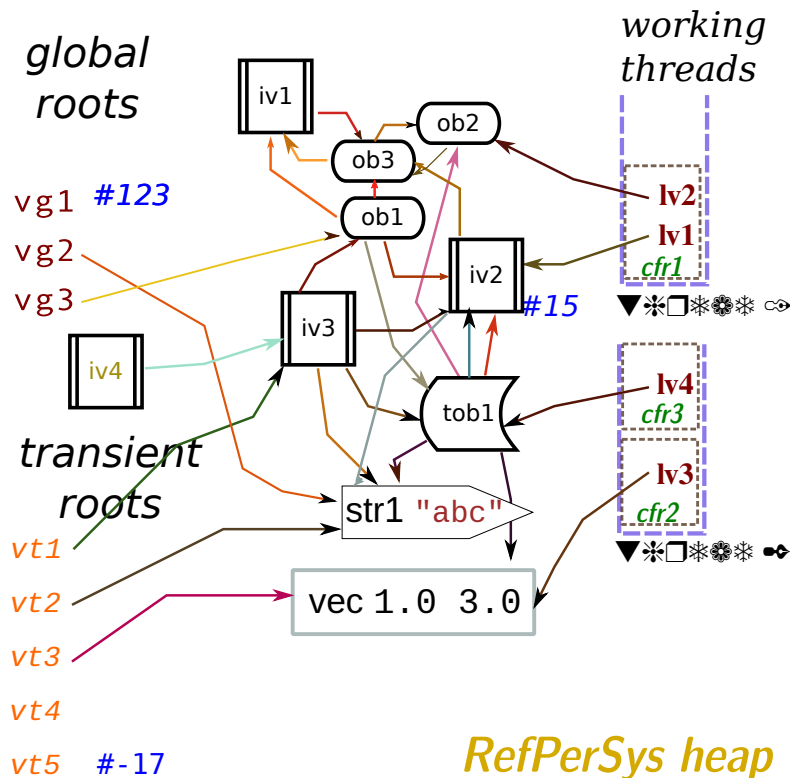
²⁷ SELF was even able (in hours of CPU time) to redefine its integers -even for arithmetic used inside its compiler- as [bignums](#).

²⁸ Perhaps inspired by simple designs like [DJANGO templates](#), but driven by frame-based REFPERSYS objects.

in C++, a `Makefile`, etc...). Actually REFPERSYS needs to conceptually have **self-modifying code** [Tschudin:2005:HarnessingSC], practically implemented by systematically doing most function calls through indirect function pointers (which gets updated with `dlsym(3)`).

2.3 REFPERSYS persistent heap

When REFPERSYS is running in some multi-threaded LINUX `process`, the REFPERSYS persistent heap is (like Bismon's one [Starynkevitch:2019:bismon-draft]) semantically like the memory heap of most `dynamic programming languages` (such as `PYTHON`, `GUILE`, `GO`, `SBCL`, etc ...). The figure 3 should give an intuition about that heap, when it is inside the `virtual address space` of some `refpersys` process. We strongly want to avoid any `GIL`, but multi-threaded precise efficient garbage collector implementations are quite difficult to code. However, notice that the persistence (dump as textual `git`-versioned disk files) of a heap uses algorithms similar to those of copying garbage collectors [wilson:1992:uniprocessorgc, jones:2016:gchandbook].



#123	a tagged 63 bits integer
<i>vg1</i>	a global persisted variable
<i>vt2</i>	a static transient variable
ob1	a mutable persistent object
iv2	an immutable constant composite value
iv4	an immutable but dead constant composite value (should be GC-ed)
tob1	a transient mutable object
<i>str1</i>	a constant UTF-8 string value "abc"
<i>vec</i>	a constant vector of floats [1.0; 3.0]
<i>lv2</i>	a local variable inside its call frame
<i>cfr1</i>	a call frame (simplified)
thread1	a working thread and its call stack (simplified)

In real life, the heap may be quite large (gigabytes) and contain hundreds of global roots or transient roots, millions of objects (sometimes transient, often persistent) and many millions immutable values (some of them composite and containing values, other scalar and containing non-pointer data like strings or vectors of float do), and dozen of working threads, each having thousands of call frames with dozens of local variables each.

Figure 3: the **REFPERSYS persistent heap** (simplified)

That figure 3 shows a few global and transient roots (both being processed by the garbage collector), and several threads each having its call stack (made of call frames) with local variables in it. In that figure, if μ and μ' are two memory zones or locations (like for an object such as `ob1`, or for an immutable value `iv2`), there is an arrow $\mu \rightarrow \mu'$ if some field ϕ of μ refers to μ' , that is (in C like notation) if $\boxed{\mu \rightarrow \phi = \mu'}$. Different arrow colors could mean different fields $\phi, \phi' \dots$ etc ... The heap is actually a large directed graph and may contain cycles (e.g. `ob1` \rightarrow `iv1` \rightarrow `ob3` \rightarrow `ob2` \rightarrow `ob3`). Most values are immutable values (some of them being composite, such as `iv1`). Some immutable values are scalar (e.g. strings). Notice that `iv4` is a dead value, unreachable from others; it should be later garbage collected. Only objects have a content which may change. Since REFPERSYS is multi-threaded, the access inside every object should be thread-safe and usually is protected by a *mutex* (or *read write lock*) which is part of that object²⁹.

Conceptually, REFPERSYS *tracing precise garbage collector* should traverse the graph of references to REFPERSYS values, starting from global or transient roots and local variables inside call frames of working threads. Each REFPERSYS **value** (immutable or object) is represented by a machine word (aligned, 64 bits) which usually contains a pointer, but sometimes some *tagged integer*. Immutable values are often “small” (typically, less than a few dozens of words of memory, sometimes a lot more) but objects are necessarily heavier since they contain some kind of lock. *closures* are immutable values, containing an object representing and giving their function code (as a C function pointer inside that object), and additional closed values. In practice our garbage collector processes not only values (either immutable values or objects), but also **quasi-values** : these are a single memory zone which is allocated using the garbage collector allocation protocol, traversed by the GC when something points to it, appears inside other values (in particular, as payload of objects), but by convention should not be passed as a genuine value. So the figure 3 is a simplification.

Some values (or objects) are dead; in the figure 3, the immutable value `iv4` is not reachable from roots or local variables on the call stack of working threads. So it is dead and should eventually be reclaimed by the garbage collector.

Values -either immutable values or changeable objects- in REFPERSYS can be either **persistent** (dumped in textual state files³⁰, then reloaded at restart of `refpersys` process) or **transient** (that is, not dumped and not appearing in state files).

²⁹Or by atomic pointers, probably the REFPERSYS class of an object is, inside it, given by some C++ field with an `std::atomic` pointer type, for efficiency reasons.

³⁰In the current implementation, REFPERSYS state files should appear under `persistore/` sub-directory, and the manifest file is `rps_manifest.json` at the top directory.

The **persistence** machinery - the dump - is conceptually simple and could run in several threads: start from global roots and traverse the memory graph but ignore transient objects and transient roots and memoize previously seen persistent objects. Of course, objects should not be persisted twice, and are referred by the **object id** or **objid** in the state files produced by the dump. That *objid* is alphanumeric, randomly generated and so hopefully globally unique -like `_2om48kc3k5R02d3ktW` for example- in our current implementation; exactly like **UUIDs** should be. Notice the conceptual similarity between REFPERSYS dump algorithm and its tracing garbage collector: both are traversing the graph of references inside the heap.

The global roots are objects. Use the C++ functions `rps_each_root_object` to iterate on them, `rps_add_root_object` to add one, `rps_remove_root_object` to remove one, `rps_is_root_object` to test if an object is a global root, `rps_set_root_objects` to get the set of all of them, and `rps_nb_root_objects` to get their number. Of course, some global roots can be transient objects, but all of them are roots for the garbage collector.

The initial loading machinery (recreating a suitable heap - and rebuilding a graph of references inspired by figure 3, without any transient stuff) from its previous dumped state) is first creating empty all objects, then later filling each of them. However, for efficiency, we may want to load the heap in parallel, using several loader threads. This could be easy if, after having created all objects as empty, and loaded plugins (i.e. `dlopen`-ing many `*.so` files), REFPERSYS processes each state file in a potentially different loading thread.

2.4 Agenda and multi-threading in REFPERSYS

Once REFPERSYS persistence is implemented and provides some meta-programming facilities, we can define and use some agenda machinery. The insight is that REFPERSYS is running several **[barney:2010:pthreads, butenhof:1997:programming] worker threads**³¹ known to its garbage collector (which might also need its own managing and synchronizing thread, which will mostly stay idle.). Our **agenda** is the central mechanism of REFPERSYS feeding these worker threads with some work to do, using **tasklets** representing a small amount of work to be done.

Each worker thread is indefinitely looping like this:

1. it runs occasionally some housekeeping processing, notably garbage collection work. This is where garbage collection gets synchronized. Occasionally,

³¹Concretely, this means `pthreads(7)`, perhaps wrapped as `C++11` threads, `QT5` threads, `GLIB` threads, etc

some new tasklets could be “auto-magically” inserted in the agenda at this point³², e.g. to run some code when some input data is available on some `file descriptor` for a `pipe(7)` or a `tcp(7)` socket, or to run some code every tenth of second, or to handle graceful termination when getting a `SIGTERM`³³ `signal(7)`.

2. it waits, if so needed (probably using PTHREADS condition variables), for the agenda to become non-empty
3. it chooses a tasklet τ to run inside the agenda. That tasklet is taken, so removed from the agenda.
4. it runs that tasklet τ for a small amount of time (a few dozen of milliseconds, typically), called a **step**³⁴. Of course during that step the agenda can (and usually will) change, and perhaps the same tasklet τ would be added again into the agenda, with maybe several other tasklets. Or on the contrary, running τ could remove one or several other tasklets $\tau_1, \tau_2 \dots$ from the agenda, and add other ones τ'_1, τ'_2, \dots there.
5. that loop is repeated (unless REFPERSYS is stopped).

The number of worker threads is fixed and small. Typically one worker thread per processor core (so 3 on a small laptop, 20 or 30 on a big desktop). Of course the agenda mechanism requires synchronization through locks or mutexes and PTHREAD condition variables [**barney:2010:pthreads**].

In addition of the worker thread, some additional slave threads could be needed, in particular to handle some event loop (and serve HTTP requests). Of course the running steps should appropriately lock objects, to avoid aftermath and synchronize properly their mutation.

The concrete organization of the REFPERSYS agenda has to be precisely defined. It could be, as **BISMON** has, a small data structure made of several first-in first-out queues, e.g. a queue of high priority tasklets, another of medium priority tasklets, one of low priority tasklets, etc..., with the agenda mechanism choosing in the non-empty queue of highest priority its tasklet staying in front.

³²Or such tasklets could be very carefully added into the agenda from non-worker threads organized in a producer-consumer fashion -such as those started by `libonion-`, respecting our GC invariants. This is a delicate issue !

³³Read also `signal-safety(7)` and consider using `signalfd(2)` or pipe-to-self tricks inspired by **QT approach to UNIX signal handling**. Notice that `timerfd_create(2)` might also be useful for tasklets to be added periodically in some `event loop` around `poll(2)`.

³⁴Calling blocking system calls such as `poll(2)` or `read(2)` from a pipe or socket should be forbidden here, because a step should run quickly, in milliseconds.

2.5 Metaprogramming and introspection in REFPERSYS

Metaprogramming is defined in [Wikipedia](#) as “a programming technique in which computer programs have the ability to treat other programs as their data. It means that a program can be designed to read, generate, analyze or transform other programs, and even modify itself while running”. That design idea is central to many Artificial Intelligence systems and AI inspired languages³⁵ and is also common in software engineering³⁶ [[Lenat:1983:Eurisko](#), [Lenat:1983:theory](#), [Lenat:1991:ev-cycl](#), [Pitrat:1996:FGCS](#), [Pitrat:2009:AST](#), [Pitrat:2009:ArtifBeings](#), [Pitrat:blog](#), [Queinnec:1996:LSP](#), [Queinnec:2004:ContinWeb](#), [Starynkevitch-1990-EUM](#), [Starynkevitch-DSL2011](#), [Starynkevitch-GCCMELTweb](#), [Starynkevitch:2007:Multistage](#), [Starynkevitch:2019:bismon-draft](#), [Tschudin:2005:HarnessingSC](#), [abelson:1996:sicp](#), [briot:1987:uniform](#), [chambers:1991:efficient](#), [cointe:1987:metaclasses](#), [dormoy:1992:meta](#), [fouet-starynkevitch:describing-control:1987](#), [greiner:1980:representation](#), [hernandez-phillips:2019:debugging-bootstrap](#), [hofstadter:2007:strange-loop](#), [kay:1996:early-smalltalk](#), [kelsey:1998:r5rs](#), [kumar:2015:importance-expert-systems](#), [matthews:2005:operational](#), [mazur:2004:compile](#), [nigro:2008:meta](#), [queinnec:2003:lisp](#), [Starynkevitch:2009:grow](#), [serrano:1995:bigloo](#)]. Generating some “source” code at build time is usual practice, advocated also by the [NINJA build system](#), and theorized (around 1930, before even computers existed) in the [CHURCH-TURING thesis](#). Related concepts include the famous (but [undecidable](#)) **halting problem** (whose proof involves a metaprogramming approach [[Hofstadter:1979:GEB](#)]), [hygienic macros](#), and [Rice’s theorem](#).

Practically speaking [[abelson:1996:sicp](#)], metaprogramming is easier achieved by explicitly representing (maybe incomplete) code with [abstract syntax trees](#) (or AST), maybe with some holes for [metavariables](#) for their later [explicit substitution](#), in the spirit of [DJANGO templates](#) or of [COMMON LISP macros](#) or [SCHEME macros](#). A practical way to implement such a template machinery for generating C or C++ code is given by [GCC MELT code chunks](#) [[Starynkevitch-DSL2011](#), [Starynkevitch-GCCMELTweb](#), [Starynkevitch:2009:grow](#), [Starynkevitch:2007:Multistage](#)], where a piece of C (or C++) code with holes (or metavariables) `$hellochunk` and `$msg` is given through the “macro-string” `#{/*$hellochunk#_here*/ printf("hello %s\n", $`
`...`

Later, such a macro-string or code chunk can be expanded by filling the holes, that is expanding the metavariables (e.g. `$msg`) appropriately. Such an expansion might be recursive, since some hole filling (or metavariable replacement) could in turn trigger

³⁵See also [SCHEME 48](#), [SBCL](#), [RUST](#), even [C++ templates](#), [CHICKEN SCHEME](#), [METAOCAML](#), the [ECLIPSE Constraint Programming System](#), [RASCAL](#), [NEMERLE](#), [COCCINELLE](#), [OCSIGEN](#), [GNU PROLOG](#), [CLIPS](#), [GPP](#), [SWIG](#), [ANTLR](#), [IBURG](#), [Gnu BISON](#), etc ...

³⁶A typical example is the [GCC compiler](#), or [AUTOCONF](#), and [transpiler approaches](#)

expansions of other macro-strings. In practice, REFPERSYS will use similar code chunks and macro-expansion to generate its C (or C++) code, and some initial ad-hoc [integrated development environment](#) (or IDE) will have to be coded, handling passively some persistent store. The expansion will be done through some scripting language (or *domain specific language*, a.k.a. DSL) which has to be implemented inside our IDE.

Metaprogramming involves code generation (using [source-to-source](#) ahead-of-time and/or [just-in-time](#)³⁷ compilation techniques [[Aho:2006:dragon-book](#)], and in REFPERSYS is useful for many tasks, such as generating the garbage collection support routines for scanning or forwarding, and the loading and dumping routines needed for persistence (in the spirit of [RPCGEN](#), [SWIG](#) and other [serialization](#) frameworks).

In REFPERSYS, metaprogramming is often and practically achieved (like in [[Starynkevitch-DSL2011](#), [Starynkevitch:2019:bismon-draft](#), [Pitrat:1996:FGCS](#), [Pitrat:2009:ArtifBeings](#)] and our `manydl.c` example program), by generating some C or C++ code in a temporary file³⁸ like `/tmp/rpsgen123.c`, compiling that file [[drepper:2011:write-shared-lib](#)] into a generated plugin `/tmp/rpsgen123.so` by running a process such as `gcc -fPIC -Wall -O -g -shared /tmp/rpsgen123.c -lsomething -o /tmp/rpsgen123.so` and waiting for its successful completion, then `dlopen(3)`-ing that newly generated `/tmp/rpsgen123.so`, in a manner compatible with our garbage collection and agenda invariants. We might later care about carefully `dlclose(3)`-ing that generated plugin, but in practice we accept some limited virtual memory plugin leak, and we could just dump appropriately our persistent state by mentioning in some generated [Manifest file](#) those plugins which should be saved (as generated C code) with the state.

[Reflection](#) is “the ability of a process to examine, introspect, and modify its own structure and behavior” and also, for [self-reflection](#), the capacity “to exercise introspection and to attempt to learn more about their fundamental nature and essence”. (Wikipedia). It is advocated (in [[Pitrat:2009:ArtifBeings](#)]) that a similar approach is (painfully) achievable in AI systems, and it would need both clever [backtracking](#) and [backtracing](#) techniques. Libraries such as Ian Taylor’s `libbacktrace` (which wants most of the code to be compiled with [DWARF debugging](#) information³⁹) are helpful.

³⁷Several JIT compilation libraries exist, notably `libgccjit` provided inside recent [GCC](#) compilers.

³⁸There are practical reasons to generate these temporary files outside of `/tmp/`, which gets cleaned at reboot.

³⁹In practice we should compile our or other C or C++ code with both `-O2 -g` passed while [invoking](#) [GCC](#) or `g++`, and this is indeed possible and practically works well enough.

Our precise garbage collector (see §3 below and [rafkind:2009:precise-gc], or QISH) wants local variables holding garbage collected pointers to be known to the GC. In practice, the REFPERSYS call frame is some explicit local `struct` named `_` in generated C code⁴⁰. Such explicited local frames can often be optimized by GCC or g++ (invoked with `-O2`).

As suggested by Pitrat (see [Pitrat:1996:FGCS, Pitrat:2009:AST, Pitrat:2009:ArtifBeings]), call stack reflection and backtrace is the elementary brick of more sophisticated *introspection* techniques. At some point, our REFPERSYS system should inspect its call stack and may take decisions after that. A typical approach would be to run such introspection once in a while (e.g. every 0.1 second on the average⁴¹, in the *inference engine* of some *expert system* or *knowledge base* component of REFPERSYS).

Since we aim to be able to re-generate most (and hopefully all) of REFPERSYS code (in C or in C++), having simple **coding conventions** does matter: every REFPERSYS-defined C or C++ identifier should start with `rps_` in lower, upper, or mixed case (e.g. also `RPS_` or `Rps_`). Every C or C++ function, even `static inline` ones appearing in header files, has its name starting with `rps_` and is *globally* unique to the entire `refpersys` program. The C (or C++) code should be automatically indented⁴² using **Gnu INDENT** or **ASTYLE**. Every named `struct` (in C) should have its tag matching `rps_*st`. Every typedef-ed data type should have its name matching `rps_*t`. Every named `enum` should have its tag matching `rps_*en` and the various enumerated values like `RPS_*`. Even in cases the C (or the C++) language allows several name spaces⁴³, we don't use that facility. Hence we refuse to code the common `typedef struct rpsfoo_t rpsfoo_t`; but prefer instead (inspired by **GTK**) coding `typedef struct rps_foo_st rps_foo_t`. Of course, names of *local variables* (that is *automatic variables* with their *lexical scope* limited to some small C or C++ *block*) could be as short as a single letter such as `i`. In general, our C or C++ code is written with the hope of being easily able to regenerate it.

⁴⁰Like **Bismon** does, see its `LOCAL_BM` macro. See also the `CAMLparami` and `CAMLlocalj` C macros of **OCAML**, and the `Py_VISIT` and `Py_DECREF` and other macros of **PYTHON**, the *foreign function interface* of **SBCL**, etc...

⁴¹Timing considerations are essential, practically speaking, in REFPERSYS. See `time(7)` man page.

⁴²With the social convention that REFPERSYS contributors are running `omake indent` or `make indent` before every `git commit`!

⁴³In C, having both a type and a label named `foo` is permitted, but we refuse such non-sense.

3 The data and object models of REFPERSYS

The data is what is processed by REFPERSYS, and is made of values (and, internally for the GC, also of quasi-values, which are pointers to GC-managed memory zones). The object model is defining our classes, our single inheritance mechanism, our message sending protocol (see §3.8.2).

3.1 how data should be processed in REFPERSYS

REFPERSYS aiming to be first a good old fashioned AI system (GOFAI), better known as *symbolic artificial intelligence* system, it is targeting mostly *symbolic computation*, in particular using a *semantic network* or other forms of mathematical finite but large *graph* representations, in particular *abstract syntax trees*⁴⁴ of generated programs, of internal rules or expressions, by some internal metaprogramming machinery. So REFPERSYS *objects* should have a finite but changing set of *attributes* or *properties* and be organized, as in most *object-oriented languages*. Hence, *documents*, *hypertext*, high-level *source code*, *ontologies*, *knowledge bases*, *expert systems*, implementation of some *inference engine* guided by *metarules*, etc ... should all be easily and conveniently representable⁴⁵ and processable, as some evolving sub-graph of REFPERSYS values.

Since REFPERSYS objects are the only mutable values, they keep not only their synchronization data, but also attributes or properties, components, and some extra *payload*⁴⁶. See also §3.4 below.

The REFPERSYS worker threads, organized in a small *thread pool*⁴⁷ are somehow organized in some *agenda*⁴⁸ mechanism. Informally, the agenda is a clever organization (perhaps a few mostly *FIFO queue* of elementary tasklets, or something more complex). Each such tasklet runs for a short time⁴⁹ and may, while running, update that agenda by adding further runnable tasklets to it, or by removing some of them. The agenda itself should be somehow reified and partly persistent, and tasklets are REFPERSYS objects. Of course some tasklets (e.g. those directly related to the user interface, e.g. *AJAX* or *QT* callbacks) are transient.

⁴⁴Practically speaking, abstract syntax trees are in fact at least finite *directed oriented graphs* and could even have cycles if you relate a *symbol* to its properties.

⁴⁵So *artifacts* like *XML* documents, *HTML5* or *XHTML* hypertexts, *JSON* data, *YAML* representations should all be easily representable and inspirational for REFPERSYS data and its processing.

⁴⁶From the GC point of view, payloads are quasi-values ...

⁴⁷Threads are heavy resources, each of them needing a call stack and, practically speaking, a processor core to run. We surely want to have at most a dozen of worker threads.

⁴⁸In Latin, “agenda” means “things which have to be done or completed”.

⁴⁹In practice, several dozens of milliseconds, to play nice with human interaction and be friendly with our garbage collector

3.2 data at the low and high levels

REFPERSYS mostly handle values⁵⁰, which can be either “light” immutable values or “heavy” mutable objects. Our data model is inspired by the [OBJVLISP](#) model (or [CLOS](#), see also the [Common Lisp HyperSpec](#)) common in most Lisp implementations [[queinnec:2003:lisp](#), [cointe:1987:metaclasses](#), [briot:1987:uniform](#)] and inspired by [SMALLTALK](#) [[kay:1996:early-smalltalk](#)]. Also, a value can be transient or persistent. Each REFPERSYS value fits in one 64 bits machine word⁵¹, so is nicely represented as an aligned pointer (ending with a 0 bit) or a tagged integer (63 bits, with the least significant bit being set to 1). Values are usually pointers to complex structures, so, per the [x86-64 calling conventions](#), are word aligned (address multiple of 8 bytes). Let’s call *genuine values* those that are not `null` and not tagged pointers (so either immutable values or objects in figure 3). These **genuine values** (and also quasi-values) are practically implemented as a [tagged union](#)⁵² and each of them start with a field (probably 16 bits) identifying their concrete type.

3.2.1 values and quasi-values

The REFPERSYS garbage collector manages both values and quasi-values (that is, a single non-empty sequence of memory words, used for some garbage collected data, e.g. inside objects). But only persistent values are dumped and reloaded in the persistent store. The values which are not dumped -so not reloaded on the next run- are called *transient* values.

For pragmatical reasons, our values⁵³ should be both ordered and hashed, since many [data structures](#) [[cormen:2009:introduction](#)], specified as some [abstract data type](#), either uses some ordering (e.g. in [red-black trees](#)) or some hash-code (e.g. various kinds of [hash tables](#)). Because of the weird and counter-intuitive semantics of [floating point](#) numbers, the [NaN](#) should be handled specifically (it is unordered), if we [box IEEE doubles](#).

3.2.2 implementation details

REFPERSYS takes advantage of some practical features⁵⁴ of C on Linux x86-64:

⁵⁰In particular, only [CLOSURES](#) are [applied](#), to arguments which are values (read more about [λ-calculus](#)); or messages are sent, to values with perhaps additional value arguments. Internally, our GC also handle quasi-values.

⁵¹Remember: REFPERSYS targets only Linux x86-64 systems!

⁵²An old example of tagged unions in C is the [X11 event structure](#), but [GUILLE](#) and [OCAML](#) use similar implementation tricks.

⁵³Of course, quasi-values need not to be ordered and hashed!

⁵⁴We don’t really care if these features are not exactly standard C11 [[c11-standard:2011](#)], because

- Practically, machine data pointers should be at least 64 bits (8 bytes) aligned⁵⁵ for large enough memory zones (i.e. most practical `struct`-s), and preferably 128 bits, that is 16 bytes, aligned. See also the `alignof` macro of `<stdalign.h>` and the [aligned type attribute](#).
- Limited [type-punning](#) abilities. Assume we have two `struct`-ures definitions, so `struct s1` and `struct s2`. Assume that both `s1` and `s2` *start* with the same *common* fields `unsigned num; then void*ptr; followed by char str[24];`. Assume that a pointer `p` points to a valid memory zone, whose alignment (respectively size) are at least all of `alignof(struct s1), sizeof(struct s1), alignof(struct s2), sizeof(struct s2)`: so we have `alignof(sizeof(*p)) >= alignof(struct s1) && sizeof(*p) >= sizeof(struct s1)` and `alignof(sizeof(*p)) >= alignof(struct s2) && sizeof(*p) >= sizeof(struct s2)`. Then: `((struct s1*)p)->num` and `((struct s2*)p)->num` both refer to the same memory location and number there; `((struct s1*)p)->ptr` and `((struct s2*)p)->ptr` both refer to the same memory location and pointer there; and of course `((struct s1*)p)->str` and `((struct s2*)p)->str` is the *same* string. See also `may_alias`, `warn_if_not_aligned`, `aligned`, `transparent_union` GCC [type attributes](#), the `-fms-extensions` option to GCC, and its [unnamed fields](#) ability.
- [Tail call](#) optimization, practically provided in *some* common cases by recent GCC or CLANG/LLVM compilers (requiring probably `-O2` compiler flag).
- Common [extensions to the C language](#), notably [statement exprs](#) (very useful), [label as values](#) (or “computed goto”-s), `typeof`, [zero-length arrays](#) and [flexible array members](#), [return addresses](#) and [other built-ins](#), may be used in REFPERSYS code.

Practically speaking, every REFPERSYS value or quasi-value (see our `Rps_QuasiZone` class) which sits in memory⁵⁶ is represented in some `class` inherited from `Rps_ZoneValue`. For instance, our string values have their memory zone type declared as `Rps_String`, but we use the `Rps_StringValue` class to construct them. In `Rps_String` the field `_sbuf` is a [flexible array member](#), and by convention contains `_bytsize + 1` bytes (terminated with a 0 byte), is validly UTF-8 encoded, aligned to 4 bytes and nul-byte terminated. Hash codes cannot be 0 and are lazily computed (so the

we strongly believe they are present on practical Linux x86-64 computers.

⁵⁵The x86-64 or AMD64 [instruction set architecture](#) allows in principle unaligned memory accesses, but these are very slow and unfriendly to [cache coherence](#) hardware implementations.

⁵⁶This excludes tagged integers, and that memory zone is at least word aligned to 8 bytes.

`rps_strhash` field is computed once when it was 0). The `Rps_Type::String` is some enumerator inside a global enum. Strings are ordered naturally, using `strcmp` on their `rps_strdata` bytes.

The `refpersys` executable is handling files either from the REFPERSYS home directory (obtained inside C++ code using a `rps_homedir()` call), given by `$REFPERSYS_HOME` or `$HOME` environment variables or thru the `-refpersys-home` program argument, or from the REFPERSYS load directory (by default the source directory, or given thru the `-load` program argument). User preferences should go into the REFPERSYS home directory, e.g. as the `.refpersys.json` file there.

@@TODO: should explain more implementation details in C++ terms?

3.3 immutable values

By definition, immutable values don't change. All their useful bits⁵⁷ stay unchanged as long as the value is alive. Some values are scalar (strings, vectors of floats, perhaps bitmaps⁵⁸ if we reify them). Other values are composite.

Since objects are fundamental, we want to keep finite collections of them. In particular, REFPERSYS will reify (represent as first-class *immutable* values) **tuples** of objects and finite **sets** of objects as values, and also `,` and they are the common composite values of REFPERSYS. A tuple is obviously represented by boxing a sequence of object references (i.e. pointers). A set would be represented by an ordered sequence of object pointers, with membership efficiently testable by a $O(\log n)$ time **binary search algorithm**). We expect most of tuples and sets to be small and fitting in an **L1 or L2 cache line**, so their processing should be efficient.

In REFPERSYS, **closures** -that is first-class procedural values, like in SCHEME⁵⁹ [**kelsey:1998:r5rs**, **matthews:2005:operational**, **Queinnec:1996:LSP**, **Queinnec:2004:ContinWeb**, **abelson:1996:sicp**], **HOP** or **BIGLOO** [**serrano:1995:bigloo**], **COMMON LISP**, **JAVASCRIPT** - are also immutable values. The closed values -binding free variables of the closure- inside such closures are arbitrary, but fixed, and won't change during the lifetime of that closure. The function code inside them is given by some fixed object reifying that code, and probably useful to generate the "source" code (e.g. as generated C or C++ code) of that function. The **mangled name**, later "`dlsym(3)`-ed", of that function in its **ELF *.so shared object** file [**drepper:2011:write-shared-lib**,

⁵⁷For housekeeping purposes, our garbage collector may reserve a few bits, e.g. for **tri-color** marking [**wilson:1992:uniprocessorgc**]

⁵⁸In practice, bitmaps or pixmaps would rather be the payload of some objects, see below.

⁵⁹Try for example GNU **GUILE** following **this tutorial**.

levine:1999:linkers-loaders] is somehow related⁶⁰ to the *objid* of that object. Closures are absolutely essential in REFPERSYS, since they are the only way to refer to executable machine code. Even method implementations are using closures, since the **virtual method table** in REFPERSYS classes (actually, their payload) is referring to closures (is is an association between selectors (like in **Objective-C**, but reified as objects) and closures implementing methods, à la OBJVLISP [**cointe:1987:metaclasses, abadi:1995:imperative**]).

We may consider also having in REFPERSYS immutable node⁶¹ instances: like mutable objects (see §3.4 below), each of them would have a class (with single-inheritance), attributes and components. But since they are immutable, they have no *objid*, no locking mechanism, and their class, attributes and components would be fixed and defined at their creation time.

In C++ code, values are `Rps_Value`, a “smart” value container, actually a single word. That class is specialized into helper subclasses, such as `Rps_StringValue`, `Rps_DoubleValue` etc.

3.3.1 immutable scalar values

Immutable scalar values include:

- **strings**, persisted as JSON strings; their internal representation is `Rps_String` using an UTF-8 encoding. In C++, use `Rps_String::make` to make one, or construct an `Rps_StringValue` with a C or C++ UTF-8 encoded string, a `std::string`, or a `QString`.
- boxed **doubles**⁶² (which cannot hold an IEEE-754 NaN, which is incomparable), persisted as JSON doubles; their internal boxed representation is `Rps_Double`. In C++, use `Rps_Double::make` to make one, or construct an `Rps_DoubleValue` with a C double.

3.3.2 immutable composite values

Immutable composite values include:

- **tuples** of object references. Their memory representation is a `Rps_TupleObzone`, and it has both `Rps_TupleOb::make` and `Rps_TupleOb::collect`

⁶⁰For a fictional example, an object of *objid* `_61NdgIkKhWD041a094` might be related to some ELF function of name close to `rps_61NdgIkKhWD041a094`, etc ...

⁶¹REFPERSYSnodes are generalizing BISMOM nodes [**Starynkevitch:2019:bismon-draft**], which are immutable, have an object connective and a sequence of sons which are arbitrary values. Perhaps “node” is a wrong word and we could name them “records” or “structures”.

⁶²See floating-point-gui.de for more about IEEE 754 double precision numbers on current computers. This is actually a difficult topic.

static functions. But use `Rps_TupleValue` to build them. **@@TODO: should explain more**

- **set** of object references; Their memory representation is a `Rps_SetOb` zone, and it has both `Rps_SetOb::make` and `Rps_SetOb::collect` static functions. But use `Rps_SetValue` to build them. **@@TODO: should explain more**

@@TODO: should explain a lot more

3.4 mutable objects

Practically speaking, mutable objects are heavy, since they should carry inside them locking devices⁶³ for multi-threading support. And each object carries an association between attributes (playing the role of arbitrary keys) and their corresponding value. In addition, an object can carry its payload⁶⁴ for stuff which does not fit into that model. For example, an object may carry as its payload a dictionary associating machine strings to values, or an hash-table of triplets, or an opened `FILE*` handle, or `file descriptor`⁶⁵, or some `TENSORFLOW` or `GHUDI` data for machine learning purposes, maybe something related to `libonion` or `ZeroMQ`, some `GMPLIB` big number, some `PPL` polyhedra, maybe some `QT` graphical widget, etc Of course every object has its class (which is itself an object, having a `metaclass`) and carries a function pointer⁶⁶, for REFPERSYS closures.

Notice that a generational GC approach moving data is possible for some immutable values, but not for objects and their optional payload, since REFPERSYS objects contain locks and payloads that are dealt with through external functions requiring fixed, unchangeable, pointers.

In C++ code, object references are `Rps_ObjectRef`, a “smart” object container, actually a single word. They are persisted by their *objid* in JSON format as a string.

3.4.1 objects as frame-like data

REFPERSYS objects are quite flexible, even more than `JAVASCRIPT`⁶⁷ ones; they take

⁶³At the implementation level, think of some `mutex` or preferably some `read-write lock`, so `pthread_mutex_init` or `pthread_rwlock_init` or `C++11 equivalents`.

⁶⁴Every payload belongs to a single object, its owner!

⁶⁵`Finalizers` are practically not enough to handle these, even if they are useful, in our GC, as a last resort measure!

⁶⁶That function pointer should be, for efficiency reasons (we don’t want to lock an object to get that pointer!) `atomic` in C parlance, and might be set using `dlsym(3)`.

⁶⁷Remember that in `JAVASCRIPT` `ob.f1` is defined to be the same as `ob['f1']`, the equivalent of

some inspiration from RLL [[greiner:1980:representation](#)], EURISKO [[Lenat:1983:Eurisko](#), [Lenat:1983:theory](#)], CYC and its CYCL [[Lenat:1991:ev-cycl](#)], and more recently the *Semantic Web* and its OWL. An object has attributes (usually a few ones, but perhaps many of them) and components (again, perhaps 0 or a few of them, but in rarer cases thousands of them), and some optional payload (quite often, it is missing). Notice the similarity with JAVASCRIPT objects (which calls *fields* with *keys* what REFPERSYS calls attributes, and array elements what are REFPERSYS components). However, JAVASCRIPT is (like IO or SELF [[chambers:1991:efficient](#)]) a [prototype-based](#) object programming language, while REFPERSYS remains, like JAVA or C#, or COMMON LISP (and of course C++ or GO), a more or less [class-based](#) object programming language with single [inheritance](#): in REFPERSYS (like in C# or JAVA) all objects are indirect instances of the same single top-level class (which is reified as an object, using some [metaclass](#) machinery).

Through their flexible attributes (each of them can be fetched, added, removed or changed during the lifetime of the containing object), REFPERSYS objects⁶⁸ can be used to represent *frames* (read about [frame languages](#)) or [semantic networks](#) and other forms of [graph databases](#) (held entirely in memory).

REFPERSYS objects keep (like BISON objects do [[Starynkevitch:2019:bismon-draft](#)]) their modification time (or *objmtime*) and that timestamp⁶⁹ may be useful to decide some further processing (à la make).

@@TODO: should explain a lot more

3.4.2 concrete examples of objects

In the below examples, ρ “*some foo*” would be the “reification” of *some foo*, that is a REFPERSYS value for that *some foo*. When we are certain that *some foo* is represented by a REFPERSYS mutable *object* (not some immutable value), we would write Ω “*some foo*” instead of ρ “*some foo*”

REFPERSYS contributors should be known to the REFPERSYS system⁷⁰. So a typical contributor would be “reified” by an object of objid `_3a9otsskmcJ04v9S7n` representing him, shown in figure 4.

fetching attribute or field `f1` from an object `ob`. And `ob[1]` is like component of index 1 in object `ob`.

⁶⁸REFPERSYS objects are also -conceptually- inspired by the [OBJECT framework](#) of GNOME.

⁶⁹The timestamp is implemented as a `double` floating point representing elapsed seconds since the [UNIX Epoch](#), obtained with `clock_gettime(2)` using `CLOCK_REALTIME`.

⁷⁰In 2019, we ignore subtle issues like [European GDPR](#) since it focuses on *transmission* of personal data, assuming that every contributor to REFPERSYS consciously decided to contribute to REFPERSYS on his own free will. We believe, similarly, that `git` users also have decided to use it with their freedom. We are not lawyers.

∈	Ω “contributor class”	the class
Ω “first name” :	string "Basile"	attribute
Ω “last name” :	string "Starynkevitch"	attribute
Ω “email” :	string "basile@starynkevitch.net"	attribute
Ω “year of birth” :	tagged integer #1959	attribute
Ω “friends” :	set { _6e08kcozzUh801dHVt7 _7zpWFJ2npj001zfVvq }	attribute

Ω “Basile STARYNKEVITCH” ≡ _3a9otsskmcJ04v9S7n

Figure 4: An example of object representing a contributor

The REFPERSYS system could flexibly add additional information, e.g. with an attribute Ω “*authored*” associated to some tuples of objects authored by that Ω “Basile STARYNKEVITCH”. Should his email change, that could be represented by overwriting attribute Ω “email” with a string such as "basile@refpersys.org". It is important that attributes are themselves objects: one could imagine that the object Ω “email” could contain an attribute Ω “how to display” associated with a closure, which would be applied to show that email cleverly (e.g. as some `HTML5 element`).

Another example of object might be some code chunk to safely print into \$fil an integer \$i, e.g. `#{ if ($fil != NULL) fprintf($fil, "%d", $i);}#` might be represented as in figure 5:

∈	Ω “code chunk class”	the class
Ω “metavariables” :	set { Ω “\$i”, Ω “\$fil” }	attribute
Ω “target language” :	Ω “C language”	attribute
[0]	string " if ("	component
[1]	object Ω “\$fil”	component
[2]	string " != NULL) fprintf("	component
[3]	object Ω “\$fil”	component
[4]	string ", \"%d\", "	component
[5]	object Ω “\$i”	component
[6]	string ");"	component

Figure 5: a simplified example of code chunk

REFPERSYS is coded **only** for LINUX/X86-64 with an English locale, using UTF-8.

3.4.3 object payloads

Objects may have some *optional* unique **payload**. The payload is owned by a single object (its owner). The payload data is generally mutable, and contains stuff which does not fit into the object model (see §3.8 below). The payload may be persistent, but some payloads are obviously transient (e.g. an opened file handle, or a Qt widget). At the implementation level, a payload is some garbage-collected quasi-value whose “type” starts with `Payl` in C++. By convention, the class of an object may require a particular payload: for example objects which are classes have to have a payload which is a class information.

Payload may be:

- a class information (inside classes, `Rps_Type::PaylClassInfo` and `Rps_PayloadClassInfo`) which gives the superclass object and the dictionary of methods (see §3.8.2 and figure 12).
- a mutable set of objects `Rps_Type::PaylSetOb` (see figure 13).
- a mutable vector of objects `Rps_Type::PaylVectOb` (see figure ??).
- a mutable vector of values `Rps_Type::PaylVectVal`
- a mutable association from objects to values `Rps_Type::PaylAssoc`
- a mutable binary relation between objects `Rps_Type::PaylRelation`
- a mutable string buffer `Rps_Type::PaylStrBuf`
- etc ...

Some payloads could be erased or replaced by another kind of payload, but some payloads are not erasable; for example, it makes no sense to replace a class information payload by a string buffer one in the same class object. Hence payloads have a `is_erasable` member function in C++.

3.5 File naming

Our *C++17 hand-written files* are named like `*_rps.cc` for source files, and `*_rps.hh` for header files, with a special case for the common `refpersys.hh` super-header file including most others. If they use QT5 extensions requiring its `moc` they would be named `*qrps.cc` for Qt C++ source files and `*qrps.hh` Qt C++ header files.

Documentation goes under `doc/` (preferably in \LaTeX , probably the `LUALATEX` variant). It could need `inkscape` version 0.92 or better and `GRAPHVIZ` version 2.40 or better.

Temporary C++ generated files, including those generated by `moc` should be named with something starting with an underscore `_` if they don't need to be `git commit-ed`.

Permanent C++ generated files which have to be version controlled so `git add-ed` go under `refpersys/generated/` directory.

Every hand written C++ file should have a proper `GPLv3+` comment at start. The copyright owner is the REFPERSYS team. We mention `refpersys.org`, and we list every human member of it.

3.6 Building `refpersys` executable

Dependencies: We need the latest `JsonCpp` library, at least its version 1.7. We need `QT5`, at least version 5.12. Our build automation system is `omake` version 0.10.3 or better. We depend upon `GNU bash` installed as `/bin/bash`. We need also `pkg-config` version 0.29 or better, suitably configured to play nice with at least `QT5`. We could later need `GNU unistring` (for UTF-8 processing) and maybe `ANTLR4` (as a parser generator).

3.7 REFPERSYS workflow

As long as we are very few and part-time on that REFPERSYS project, we essentially use `git` as an improved *centralized* version control system à la `svn` (so the *distributed* nature of `git` is irrelevant for us in 2019. It could become important when the REFPERSYS matures and generates a lot of files). By social convention: we `git commit` often (e.g. every hour or two of work). Before that, we `omake indent` and we ensure that the code is buildable with `omake clean` followed by `omake -project` before any `git commit`. We format and indent manually written C++ code with `omake indent` or using our `indent-cxx-files.sh` shell script before any `git push`.

Our `git commit` messages given by `git log` are starting with a short sentence in English (ASCII characters only). If more than one sentence is needed, the following ones should start with a blank line.

Glossary

Artificial General Intelligence [Artificial general intelligence](#) (AGI) refers to the specific capacity of a machine to learn and understand any intellectual task that can be performed by a human being. It is the primary goal of some artificial intelligence research, and is sometimes referred to as “strong AI” or “full AI”. . 1

metaknowledge [Metaknowledge](#) is knowledge about knowledge, and is an inclusive term spanning several disciplines. Bibliography, the academic study of books, and epistemology, the philosophical study of knowledge, are examples of metaknowledge. Even the tagging of documents could be considered as metaknowledge. In AGI, metaknowledge refers to the knowledge about knowledge-based systems. A declarative system might be guided by metarules, that is “expert system” rules to compile or interpret other rules (maybe themselves). . 2

Reflection [Reflection](#) is (according to *Wikipedia*) “the ability of a process to examine, introspect, and modify its own structure and behavior”, and related to [Self-reflection](#), the capacity of humans (and hopefully of artificial cognitive systems, like REFPERSYS should become) “to exercise introspection and to attempt to learn more about their fundamental nature and essence”. . 2

In practice, such a code chunk representation could be more compact; we could assume that both the `if` keyword of C and the `fprintf` and `NULL` C identifiers occur frequently enough to be refactored and each reified into its own object. Of course, it might make sense to add an Ω “author” attribute in our code chunk, whose value would be our `_3a9otsskmcJ04v9S7n` object of figure 4 above.

Obviously, some kind of data don’t fit exactly into such simple objects. Some objects might represent a big hashtable of triplets, and such data has to be the payload of that object. Other objects might reify sorted dictionaries mapping strings to values, and their payload could be some [red-black trees](#) whose internal nodes would be GC-managed quasi-values. And an opened `FILE*` would be represented and reified as some object carrying a payload with some `FILE* rps_filehandle; field.`

@@TODO: TO BE WRITTEN

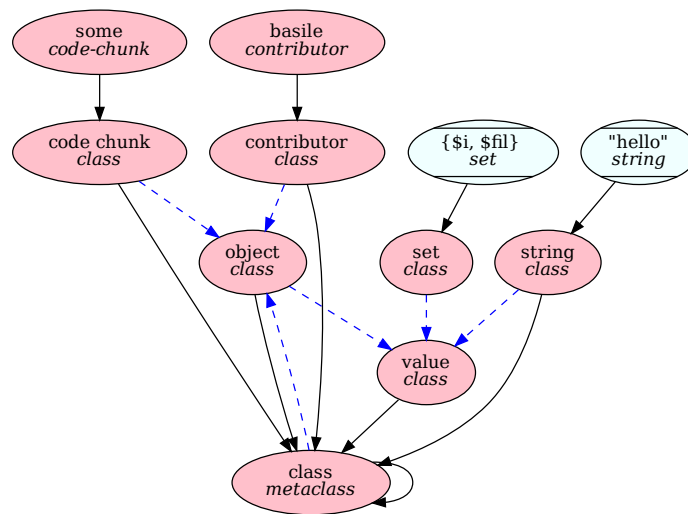
3.8 the REFPERSYS object model

Every REFPERSYS value belongs to some single class, reified as a particular REFPERSYS object. We first explain the inheritance graph (see §3.8.1), and then the

message sending protocol (see §3.8.2).

3.8.1 REFPERSYS inheritance graph

Every non-nil REFPERSYS value belongs to a single class, defining its behavior by the set of selectors it is understanding for message sending. The figure 6 shows (with simplification) the single-inheritance graph of REFPERSYS values. In practice we expect many hundreds of classes and at least many hundred thousands values in a mature persistent store.



- objects (pink background)
- immutable values (azure background, horizontal lines)
- $v \rightarrow \omega$ (straight black arrow) means : v is instance of class ω
- $\omega_1 \dashrightarrow \omega_2$ (dashed blue arrow) means : class ω_1 is subclass of ω_2

Figure 6: The simplified inheritance graph in REFPERSYS

That figure 6 shows several objects:

- the `OBJECT` class, superclass of every object;
- the `CLASS` metaclass, class of every class;
- the `VALUE` class, ultimate class of every value;
- the `CODE_CHUNK` class, for code chunks like in figure 5;

- the `CONTRIBUTOR` class, for reified contributors like in figure 4;
- the `STRING` class, of immutable string values;
- the `SET` class, of immutable set of objects;
- the `BASILE` contributor object of figure 4;
- some *code-chunk* object, like in figure 5;

and several immutable values:

- the "hello" string;
- set of two chunk metavariables { $\Omega(\$i), \Omega(\$fil)$ } which appears as value of attribute Ω “metavariables” in figure 5

A REFPERSYS class object should contain, in its payload some *class information*:

- the sequence of its direct then indirect super classes
- a flexible *dispatch table* or *virtual method table*⁷¹ associating selectors to closures handling messages with them. We call that association the direct *method dictionary* of that class. It should be implemented efficiently (perhaps using caching techniques local to each occurrence of sending).

Both are changeable. Any object can change its class at will at any time. A class can have new methods added or removed at any time. A class can change its superclass at will⁷².

3.8.2 REFPERSYS message sending

Our message sending protocol is inspired by those of SMALLTALK, COMMON LISP, OCAML, JAVA. Every message send has a receiver ρ (the target of the message sending), a selector - some object ω_{sel} (what do we send) and optional extra arguments $\alpha_1 \dots \alpha_n$ (so $n = 0$ when we don't have extra arguments). Conceptually, what is happening is some loop:

- let κ be initially the class of ρ , the receiver or target.

⁷¹Our methods are *always* virtual, like for SMALLTALK; Conceptually REFPERSYS don't have any non virtual methods.

⁷²However, this should be done with care, avoiding additional circularities in the inheritance graph.

- look into the method dictionary δ of class κ ; if the selector ω_{sel} is associated to method μ (some closure), apply that μ to $\rho, \alpha_1, \dots, \alpha_n$; the result of this application is the result of the message sending.
- if κ is the topmost VALUE class, the message sending has failed.
- otherwise, replace κ by its direct superclass κ' and repeat the method lookup (second step here).

In practice, we might try to use caching techniques (but later) à la SELF or JAVASCRIPT implementation to accelerate message sending. We should define what happens when no method is found (perhaps using some MESSAGE-NOT-UNDERSTOOD built-in selector à la SMALLTALK [kay:1996:early-smalltalk]), taking $\omega_{sel}, \alpha_1, \dots, \alpha_n$ as arguments, or triggering some exception machinery.

4 Persistence in REFPERSYS

The persistence of REFPERSYS is an essential feature. The `refpersys` program starts by loading its persistent state (from various textual files under `persistore/` directory⁷³). In the usual case, a `refpersys` process dumps its persistent state before exiting.

A **manifest file** named `rps_manifest.json` is describing the entire persisted state and referencing indirectly other files. The figure 7 is giving the syntax of that file.

```
manifest ← {
  "format": "RefPerSysFormat2019A"    mandatory format id
  "spaceset": [ idspace ... ]         oids of spaces
  "globalroots": [ idroot ... ]       oids of global roots
  "plugins": [ idplugin ... ]         oids of dlopen-ed plugins
}
```

Figure 7: syntax of the manifest file `rps_manifest.json`

If `_8J6vNYtP5E800eCr5q` is a space oid id_{space} , then the persistent space data is in JSON file `persistore/sp_8J6vNYtP5E800eCr5q-rps.json`.

⁷³Of course, some other directory can be given through explicit program arguments to the `refpersys` executable.

For plugins, if `_7GIB3ma21I200tfqDs` is a plugin oid id_{plugin} , its generated C++ source code should go into the file `generated/rps_7GIB3ma21I200tfqDs-mod.cc` and the corresponding `dlopen`-ed plugin in `plugins/rps_7GIB3ma21I200tfqDs-mod.so` ELF shared object file.

The loader will `rps_add_root_object` every root object of given id_{root} .

@@TODO: improve

4.1 The textual data format of REFPERSYS

Each space file of id_{space} starts with a prologue whose syntax is in figure

```
space-prologue ← {
    "format": "RefPerSysFormat2019A"    mandatory format id
    "spaceid":  $id_{space}$                 the id of the current space
    "nbojects": number-of-objects
}
```

Figure 8: JSON syntax of the prologue of space id_{space}

then each object content of some given oid is preceded by a comment like `//+ob oid` , for example an object of oid `_3fzIPzN1WfV01GGQSt` starts with a comment `//+ob_3fzIPzN1WfV01GGQSt` line. The following object content is described in figure 11 below.

@@TODO: review and improve ! We use a JSON format to persist our state. Our immutable values could easily be represented in textual syntax, for example a set of three objects of objids `_0iaOILq4pj20097DNb`, `_1R4TeqlLvhs03o0mGN`, `_7m9EMmdyQKU00euKwB` might be represented as the following JSON object:

```
{ "vtyp" : set
  "elem" : [ "_0iaOILq4pj20097DNb",
             "_1R4TeqlLvhs03o0mGN",
             "_7m9EMmdyQKU00euKwB" ] }
```

4.2 EBNF Grammar of Data Format

The figure 9 gives the JSON syntax of scalar values or object references in persisted state files. The figure 10 is **@INCOMPLETE@** and gives the JSON syntax of composite values in persisted state files. The figure 11 gives the JSON syntax of object contents inside space files.

The syntax of object contents is given in figure 11. The *payload-kind* there is either some C identifier (if it starts with a letter: A ... Z or a ... z) or some objid (if it

value	←	int	tagged integers
		float	double precision floating point numbers
		string	string of Unicode characters enclosed in double quotes
		object	reference to mutable <i>object</i> with a globally unique objid
		set	set of ordered unique <i>objects</i>
		tuple	set of ordered (and possibly duplicate) <i>objects</i>
		closure	function closing over an environment of <i>values</i>
int	←	α ,	a 63-bit integer represented as an JSON number type
		$\alpha \in \mathbb{Z}$,	
		$-2^{62} \leq \alpha \leq 2^{62} - 1$	
float	←	<i>floating-point-number</i> ,	an IEEE 754 double, with a dot
string	←	" α ",	a UTF-8 string represented as an JSON string type
		{ "string": σ }	when string σ looks like an objid
object	←	$_ \alpha$,	a Base-62 number prefixed with an underscore

Figure 9: JSON syntax of scalar values and object references

starts with an underscore $_ \dots$). When it is some C identifier *ident*, an extern "C" function (of signature `rpsldpysig_t`, defined in file `refpersys.hh`) named `rpsldpy_`*ident* is found by `dlsym(3)` then invoked at load time. When it is some objid *objid*, the `rpsldpy`*objid* function is called. For example, class objects have "payload": "class" as a JSON field in their state file, so are loaded by calling `rpsldpy_class`. If (later) we would have "payload": "_2j66FFjmS7n03HNNBn", then `rpsldpy_2j66FFjmS7n03HNNBn` should be called.

The JSON representation of payloads vary. The figure 12 explains class related payload.

The figure 13 gives the format of mutable set of objects payload. The "setob": array might be empty.

The figure ?? gives the format of mutable vector of objects payload. The "vectob": array might be empty. Some components could be null.

5 Metaprogramming in REFPERSYS

Metaprogramming, that is generation of "program" text files (e.g. of C++ code, JavaScript code) is an important insight of REFPERSYS. Generation of such C++ files is inspired by [Starynkevitch-DSL2011] and the code chunks there.

A code chunk is a REFPERSYS object, of class `code_chunk` (of `oid_3rXxMck40kz03RxRLM`, which is conceptually a mix of strings and holes or metavariables expanded into fur-

```

set ← {
  "vtype": "set",
  "elem": [  $\omega_1$ ,  $\omega_2$ ,  $\omega_3$ , ... ]
}
where the  $\omega_i$  are objects represented by objid-s

tuple ← {
  "vtype": "tuple",
  "comp": [  $\omega_1$ ,  $\omega_2$ ,  $\omega_3$ , ... ]
}
where the  $\omega_i$  are objects represented by objid-s

closure ← {
  "vtype": "closure"
  "fn":  $\omega_{fun}$            the object giving a function
  "env": [  $v_1$ ,  $v_2$ , ... ]
}
where the  $v_i$  are JSON for closed values.

```

Figure 10: JSON syntax of composite values

ther code (by a machinery to be defined later, and hopefully to be generated with some code chunks).

@@TODO: complete, explain and improve !

6 The primordial Read-Eval-Print-Loop of REFPERSYS

In commits around b986354245ee38db24 (November 2020), a *GNU readline*-based **Read-Eval-Print-Loop** of REFPERSYS is developed, in the -still handwritten-C++ file `repl_rps.cc`.

That initial REPL language should be capable of calling most of the public C++ functions (in our C++ file `refpersys.hh`) creating objects and values and modifying them.

Auto-completion (in particular of existing names, and of *existing* oids) is practically essential. This is why *GNU readline* is needed.

Given the lack of popularity of Lisp syntax in 2020 or 2021, the syntax of the REPL language should be closer to JavaScript or to Python, even if the semantics of REFPERSYS is heavily Lisp-inspired. See markup file `doc/repl.md` for details.

The lexer C++ routine `rps_repl_lexer` in our C++ file `repl_rps.cc` is returning (for each lexical token) a *pair* of values. For example, an integer like `-23` or `0xffff` is lexed as a pair `int, -23` or `int, 65535` respectively, where `int` is a primordial REFPERSYS object (actually a REFPERSYS class, of `oid_2A2mrPpR3Qf03p6o5b`). An existing object ω (given by its name, or by its oid) is lexed as a pair `object, ω` . So `int` would be lexed as the pair `object, int`. A

```

object-content ← {
    "oid": oid,           the string oid of the current object
    "mtime": modtimeoid,  its modification time
    "class": class-oidoid, the oid of its class
    [ "payload": payload-kind ] optional payload kind
    [ "comps": [ value ... ], ] optional components
    [ "attrs": [ attr-entry ... ], ] optional attributes
    @@@ incomplete @@@
}

attr-entry ← {
    "at": object           the oid of the attribute key
    "va": value           the corresponding attribute value
}

```

Figure 11: JSON syntax of object contents inside space files

```

class-payload :
    "payload": "class",
    "class_super": object,  the oid of the superclass
    "class_methodict": [ method-entry ... ]  method dictionary

method-entry ← {
    "methosel": oid           the oid of the method selector
    "methclos": closure       the corresponding method closure
}

```

Figure 12: JSON syntax of class payload

new (unknown) name σ would be lexed as `symbol`, string σ , e.g. some unknown `foo` would be lexed as `symbol`, `"foo"`. Delimiters δ are kept in the global `repl_delim` string dictionary (of `oid_627ngdqrVfF020ugC5`) -associating strings like `"` to objects ω_δ of class `repl_delimiter`- and lexed as a pair `repl_delimiter, ω_δ` .

7 The Web interface of REFPERSYS

In commits around `6d44cba00aa9b0a51` (May 2021) a Web interface is worked upon (but buggy). You might run the `./refpersys -AWEB,REPL -W.` (or maybe just `./refpersys -W.`) shell command and browse the `localhost:9090` URL. The file `webroot/index.html.rps` is then served, and hopefully should be template-expanded. Inspired by PHP, processing instructions like

```
<?refpersys suffix='rpshtml' action='_2s15Gjb7sw004EcMqf' rps_json='{ "foo":1 }' ?>
```

```
set-objects-payload :
    "payload": "setob",
    "setob": [           oid ... ] sorted oids of elements
```

Figure 13: JSON syntax of mutable set of objects payload

```
vector-objects-payload :
    "payload": "vectob",
    "vectob": [           oid ... ] sorted oids of elements
```

Figure 14: JSON syntax of mutable vector of objects payload

should be expanded by closure applications⁷⁴. However, such processing instructions should be on one single line (for example in file `webroot/index.html.rps`).

The advantage of a Web interface is at first to be able to generate and use HTML5, and hopefully to generate web forms.

In May 2021 the `refpersys` program should not be used as a web server accessible to outside, there are cybersecurity concerns. It could in a month be run on some isolated laptop to hopefully make a tiny demo. With websockets, some JavaScript code could be generated and running in the web browser, in parallel of the agenda mechanism running inside `refpersys` process. This has to be implemented, hopefully with generated C++ code which would generate JavaScript code and send it to web browser, which would use websockets to communicate asynchronously with the `refpersys` process.

At the initial stage, we aim primarily to allow the user to enter C++ code for a specific plugin, and to display the details for a selected object.

In the case of the former (entering C++ code), we would require to take advantage of Javascript code editing plugin, such as that provided by codemirror.net. We would need to add the `codemirror.js` file to the `webroot/js/` directory, and include it from our `index.html.rps` file. We would then create a textarea in our index HTML file, and apply the codemirror API on it. We would also require an HTTP POST endpoint in the RefPerSys web server that can process the plugin code received by it.

Handling of HTTP requests is done with the help of `libonion`⁷⁵.

⁷⁴In this example, the applied closure has as connective the object of oid `_2s15Gjb7sw004EcMqf`. See our C++ function `rps_serve_onion_expanded_stream` in file `httpweb_rps.cc` for details.

⁷⁵an HTTP server open source library from coralbits.com/static/onion/

Glossary

Artificial General Intelligence [Artificial general intelligence \(AGI\)](#) refers to the specific capacity of a machine to learn and understand any intellectual task that can be performed by a human being. It is the primary goal of some artificial intelligence research, and is sometimes referred to as “strong AI” or “full AI”. . 1

metaknowledge [Metaknowledge](#) is knowledge about knowledge, and is an inclusive term spanning several disciplines. Bibliography, the academic study of books, and epistemology, the philosophical study of knowledge, are examples of metaknowledge. Even the tagging of documents could be considered as metaknowledge. In AGI, metaknowledge refers to the knowledge about knowledge-based systems. A declarative system might be guided by metarules, that is “expert system” rules to compile or interpret other rules (maybe themselves). . 2

Reflection [Reflection](#) is (according to *Wikipedia*) “the ability of a process to examine, introspect, and modify its own structure and behavior”, and related to [Self-reflection](#), the capacity of humans (and hopefully of artificial cognitive systems, like REFPERSYS should become) “to exercise introspection and to attempt to learn more about their fundamental nature and essence”. . 2

Index

A

abstract syntax tree, 14
agenda, 12
AST, 14

E

evolution
 Darwinian, 2

F

free software, 2

G

genuine
 value, 18

I

IDE, 15
infrastructure, 2
integrated development environment, 15

M

metaprogramming, 14

O

objid, 12

P

payload, 11, 17, 25

Q

quasi-value, 11, 17, 25

R

REFPERSYS, 1
`rps_manifest.json`, 30

S

SOFTWARE HERITAGE, 2
string, 21

T

tree
 abstract syntax, 14
tuple, 20

V

value, 11