Master Internship Proposal - 2025/2026

Non-aggregative Models of Resource Composition

D.Galmiche, D. Larchey-Wendling, D. Méry
LORIA, CNRS, Nancy
email: galmiche@loria.fr, larchey@loria.fr, dmery@loria.fr

1 Context

Nowadays computer systems are highly complex. Most of them are intrinsically distributed, dynamically and reconfigurable (Big Data, Internet of Things, Cloud Computing) and concurrent at both the physical level (CPU cores) and the logical level (threads). This level of complexity raises important technical problems. For instance, how to ensure that a system consisting of a huge multitude of heterogeneous components, built by different developers, provides the expected services? How to ensure that the components do not block waiting one for another (deadlock), or that one component has to wait indefinitely (denial of service)? How to check that the update of a component's software does not leak confidential information to third parties?
Understanding complex systems is only possible due to their *modularity* and the components of a modular system can be viewed, at a higher level of abstraction, as *resources*, that can be either static or dynamic (a piece of data or a process), physical or logical (CPUs or threads), having simple atomic (a memory cell) or more elaborated structure (such as a linked list or a pipeline). The main operation on a resource domain is *composition*, formally understood as the product of a partial monoid. The resource domains having received most attention so far have *aggregative* composition operations, i.e. composition operations that behave like unions of disjoint sets, as in Separation and Bunched Logics [6, 11]. Some of the algebraic (equational) properties of aggregative compositions, i.e. compositions which obey the principle that *the whole is the sum of its parts*, have already been considered in existing works [3, 9]. However, because the aggregative nature of a composition operator cannot be fully understood in terms of its equational properties alone, we need to introduce more structural aspects like *preordering* the resources or measuring their *size*. For instance, an expected property of an aggregative operator is that it should keep its operands as discernable parts of its result. It should also make the whole composite *bigger* than any of its parts. However, many natural resource compositions do not follow those principles.

For an intuitive example, consider a *merging* composition about dual resources that annihilate each other (like particles and anti-particles in physics). Merging a resource (particle) a with its dual resource (anti-particle) b, we get the empty resource e (intangible amount of energy). Obviously, the size of the composition $a * b$ is not bigger than the size of a or the size of b, and only looking at e, one cannot tell if it is the result of $a * b$, or of the merging of another dual pair c and d.

2 Subject

A first goal of the internship is to investigate non-aggregative models of resource composition and study what kind of resource logics arise from these new resource models. More precisely, and as a first case study, we are interested in investigating what kind of logic (in terms of BI-like sharing and separation connectives) arises from non-aggregative operators like merging and fusion.

However, not all interesting aspects of non-aggregative models can be formalized solely in terms of separation connectives. For instance, time and space distributions are in fact better handled via modalities [1, 2, 4]. Current works mixing modal and resource logics rely on models in which the modalities are interpreted over an accessibility relation that is completely independent of the preordering relation over which the resources are interpreted. Such an orthogonal view of combination helps reusing existing techniques from both fields, but makes it more difficult for the resulting logics to deal with properties involving the interplay of modal and resource aspects.
For example, in our recent temporal extensions of BI called LTBI [7] the resources are distributed over points in time at which they are considered to be available. The formula $A * (A -\!\!* B)$ thus means that if A is satisfied at point t with resources r_1 and $A -\!\!* B$ is satisfied at point t with resources r_2, then B should be satisfied at point t with resources $r_1 * r_2$.
The interpretation of the resource sensitive connectives does not depend on time. Changing the conclusion to "then B should be satisfied at point $t + |r_1|$" would allow for a more interacting example where the satisfiability (availability in time) of B would depend on the size of r_1 (meaning that the bigger A is, the longer it takes to produce B).

A second goal of the internship is to study extensions and variants of LTBI in which interactions between modal and separating connectives are enhanced with measures on the size of resources and investigate the expressive power of such variants.

References

[1] Gabrielle Anderson and David Pym. A calculus and logic of bunched resources and processes. *Theoretical Computer Science*, 614(C):63–96, February 2016.

[2] Stephen Brookes and Peter W. O'Hearn. Concurrent separation logic. *ACM SIGLOG News*, 3(3):47–65, August 2016.

[3] James Brotherston and Jules Villard. Parametric Completeness for Separation Theories. In *Principles of Programming Languages, POPL*, pages 453–464. ACM, 2014.

[4] Jean-René Courtault, Didier Galmiche, and David Pym. A logic of separating modalities. *Theoretical Computer Science*, 637:30–58, July 2016.

[5] Didier Galmiche and Daniel Méry. Semantic labelled tableaux for propositional BI without bottom. *Journal of Logic and Computation*, 13(5):707–753, 2003.

[6] Didier Galmiche, Daniel Méry, and David Pym. The semantics of BI and Resource Tableaux. *Mathematical Structures in Computer Science*, 15(6):1033–1088, 2005.

[7] Didier Galmiche and Daniel Méry. Labelled Tableaux for Linear Time Bunched Implication Logic In *In 8th International Conference on Formal Structures for Computation and Deduction, FSCD 2023*, Roma, Italy, July 2023.

[8] Jonas Braband Jensen and Lars Birkedal. Fictional separation logic. In Helmut Seidl, editor, *Programming Languages and Systems*, pages 377–396, Berlin, Heidelberg, 2012. Springer Berlin Heidelberg.

[9] Dominique Larchey-Wendling and Didier Galmiche. Looking at separation algebras with boolean bi-eyes. In Josep Díaz, Ivan Lanese, and Davide Sangiorgi, editors, *TCS 2014*, volume 8705 of *LNCS*, pages 326–340. Springer, 2014.

[10] Peter W. O'Hearn and David Pym. The Logic of Bunched Implications. *Bulletin of Symbolic Logic*, 5(2):215–244, 1999.

[11] Peter W. O'Hearn, John C. Reynolds, and Hongseok Yang. Local reasoning about programs that alter data structures. In *Proceedings of the 15th International Workshop on Computer Science Logic*, CSL '01, pages 1–19, 2001.