

Towards Certified Compilation of Financial Contracts

Danil Annenkov and Martin Elsmann

University of Copenhagen
Dept. of Computer Science (DIKU)
{daan,mael}@di.ku.dk

Abstract

We present an extension to a certified financial contract management system that allows for templated financial contracts and for integration with financial models through verified compilation into so-called payoff-expressions, which readily allow for determining the value of a contract in a given evaluation context, such as contexts created for simulations. The templating mechanism is useful both at the contract specification level, for writing generic reusable contracts, and for reuse of code that, without the templating mechanism, needs to be recompiled for different evaluation contexts. We report on the effect of using the certified system in the context of a GPGPU-based Monte Carlo simulation engine for pricing various over-the-counter (OTC) financial contracts. The full contract-management system, including the payoff-language compilation, is verified in the Coq proof assistant and certified Haskell code is extracted from Coq and integrated with an efficient OpenCL pricing engine.

1 Background and Motivation

New technologies are emerging that have potential for seriously disrupting the financial sector. In particular, blockchain technologies, such as Bitcoins [7] and the Ethereum Smart Contract peer-to-peer platform [9], have entered the realm of the global financial market and it becomes essential to ask to which degree users can trust that the underlying implementations are really behaving according to the specified properties. Unfortunately, the answers are not clear and errors may result in irreversible high-impact events.

The work presented here builds on a series of previous work on specifying financial contracts [1, 2, 4, 6, 8] and in particular on a certified financial contract management engine and its associated contract DSL [3]. This framework allows for expressing a wide variety of financial contracts (a fundamental notion in financial software) and for reasoning about their functional properties (e.g., horizon and causality). As in the previous work, the contract DSL that we consider is equipped with a denotational semantics, which is independent of stochastic aspects and depends only on an *external environment* $\text{ExtEnv} : \mathbb{N} \times \text{Label} \rightarrow \mathbb{R} + \mathbb{B}$, which maps observables (e.g., the price of a stock on a particular day) to values. In the work presented here, we present a certified compilation scheme that compiles a contract into a *payoff function*, which aggregates all cashflows in the contract, after discounting them according to some model. The result represents a single “snapshot” value of the contract. The payoff language, which is inspired by traditional payoff languages and is well suited for integration with Monte Carlo simulation techniques for pricing, is essentially a small subset of a C-like expression language enriched with notation for looking up observables in the external environment. We show that compilation from the contract DSL to the payoff language preserves the cashflow semantics.

The contract DSL described in [3], deals with concrete contracts, such as a one year European call option on the AAPL (Apple) stock with strike price \$100. The lack of genericity means that each time a new contract is created (even a very similar one), the contract management engine needs to compile the contract into the payoff language and further into a target language for embedding into the pricing engine. To avoid this recompilation problem, we introduce the notion of a *financial instrument*, which allows for templating of contracts and which can be

turned into a concrete contract by instantiating template variables with particular values. For example, a European call option instrument has template parameters such as maturity (the end date of the contract), strike, and the underlying asset that the option is based on. Compiling such a template once allows the engine to reuse compiled code, giving various parameter values as input to the pricing engine.

Moreover, an inherent property of contracts is that they evolve over time. This property is precisely captured by a contract reduction semantics. Each day, a contract becomes a new “smaller” contract, thus, for pricing purposes, contracts need to be recompiled daily, resulting in a dramatic compile time overhead. To avoid recompilation in this case, the generated payoff code is parameterized over the *current time* so that evaluating the payoff code at time t gives us the same result (upto discounting) as first advancing the contract to time t , then compiling it to the payoff code, and then evaluating the result.

The contract analysis and transformation code forms a core code base, which financial software crucially depends on. A certified programming approach using the Coq proof assistant allows us to prove the above desirable properties and to extract certified executable code.

2 The Contract Language

We assume a countably infinite set of program variables, ranged over by v . Moreover, we use n , i , f , and b to range over natural numbers, integers, floating point numbers, and booleans. We use p to range over parties. The contract language that we consider follows the style of [3] and is extended with template variables:

$$c ::= \text{zero} \mid \text{transfer}(p,p) \mid \text{scale}(e,c) \mid \text{translate}(t,c) \mid \text{checkWithin}(e,t,c,c) \mid \text{both}(c,c)$$

$$e ::= \text{op}(e, e, \dots, e) \mid \text{obs}(l,i) \mid f \mid b \quad t ::= n \mid v \quad \text{op} ::= \text{add} \mid \text{sub} \mid \text{mult} \mid \text{lt} \mid \text{neg} \mid \text{cond} \mid \dots$$

Expressions (e) may contain *observables*, which are interpreted in an external environment. A contract may be empty (**zero**), a transfer of one unit (for simplicity) (**transfer**), a scaled contract (**scale**), a translation of a contract into the future (**translate**), the composition of two contracts (**both**), or a generalized conditional **checkWithin**($cond, t, c_1, c_2$), which checks the condition $cond$ repeatedly during the period given by t and evaluates to c_1 if $cond = \text{true}$ or to c_2 if $cond$ never evaluates to **true** during the period t .

The main difference between the original version of the contract language and the version presented here is the introduction of *template expressions* (t), which, for instance, allows us to write contract templates with the contract maturity as a parameter. This feature requires refined reasoning about the temporal properties of contracts, such as causality. Certain constructs in the original contract language, such as **translate**(n, c) and **checkWithin**($cond, n, c_1, c_2$), are designed such that basic properties of the contract language, including the property of causality, are straightforward to reason about. In particular, the displacement numbers n in the above constructs are constant positive numbers. For templating, we refine the constructs to support template expressions in place of positive constants. One of the consequences of adding template variables is that the semantics of contracts now depends also on mappings of template variables in a *template environment* $\text{TEnv} : \text{Var} \rightarrow \mathbb{N}$, which is also the case for many temporal properties of contracts. For example, the type system for ensuring causality of contracts [3] and the concept of horizon are now parameterized by template environments.

3 The Payoff Intermediate Language

The main motivation behind the payoff language is to bridge the gap between the contract DSL and a traditional expression language, which is usually used to implement pricing engines. The

payoff language should be relatively straightforward to compile to various target languages such as Haskell, Futhark [5], or OpenCL.

$$\begin{aligned} il &::= \text{now} \mid \text{model}(l, t) \mid \text{if}(il, il, il) \mid \text{loopif}(il, il, il, t) \mid \text{payoff}(t, p, p) \mid \text{unop}(il) \mid \text{binop}(il, il) \\ \text{unop} &::= \text{neg} \mid \text{not} \quad \text{binop} ::= \text{add} \mid \text{mult} \mid \text{sub} \mid \text{lt} \mid \text{and} \mid \text{or} \mid \dots \quad t ::= n \mid i \mid v \mid \text{tplus}(t, t) \end{aligned}$$

The payoff language is an expression language ($il \in \text{IExpr}$) with binary and unary operations, extended with conditionals and generalized conditionals `loopif`, behaving similarly to `checkWithin`. Template expressions ($t \in \text{TExprZ}$) in this language are extensions of the template expressions of the contract language with integer literals and addition. Terms in the payoff language can be evaluated given a proper external environment, a proper template environment, and a *discount function* $d : \mathbb{N} \rightarrow \mathbb{R}$. The result of the evaluation is a single real value in contrast to the contract language for which the semantics is given in terms of traces.

4 Compiling Contracts to Payoffs

The contract language consist of two levels, namely constructors to build contracts (c) and expressions used in some of these constructors (`scale`, `checkWithin`, etc.). We compile both levels into a single payoff language. The compilation functions $\tau_e \llbracket - \rrbracket : \text{Expr} \times \text{TExprZ} \rightarrow \text{IExpr}$ and $\tau_c \llbracket - \rrbracket : \text{Contr} \times \text{TExprZ} \rightarrow \text{IExpr}$ are recursively defined on the syntax of expressions and contracts, respectively, taking the starting time $t_0 \in \text{TExprZ}$ as a parameter.

$$\begin{aligned} \tau_e \llbracket \text{cond}(b, e_0, e_1) \rrbracket_{t_0} &= \text{if}(\tau_e \llbracket b \rrbracket_{t_0}, \tau_e \llbracket e_0 \rrbracket_{t_0}, \tau_e \llbracket e_1 \rrbracket_{t_0}) & \tau_c \llbracket \text{zero} \rrbracket_{t_0} &= 0 \\ \tau_e \llbracket \text{obs}(l, i) \rrbracket_{t_0} &= \text{model}(l, \text{tplus}(t_0, i)) & \tau_c \llbracket \text{translate}(t, c) \rrbracket_{t_0} &= \tau_c \llbracket c \rrbracket_{\text{tplus}(t_0, t)} \\ \tau_c \llbracket \text{transfer}(p_1, p_2) \rrbracket_{t_0} &= \text{payoff}(t_0, p_1, p_2) & \tau_c \llbracket \text{both}(c_0, c_1) \rrbracket_{t_0} &= \text{add}(\tau_c \llbracket c_0 \rrbracket_{t_0}, \tau_c \llbracket c_1 \rrbracket_{t_0}) \\ \tau_c \llbracket \text{scale}(e, c) \rrbracket_{t_0} &= \text{mult}(\tau_e \llbracket e \rrbracket_{t_0}, \tau_c \llbracket c \rrbracket_{t_0}) \\ \tau_c \llbracket \text{checkWithin}(e, t, c_1, c_2) \rrbracket_{t_0} &= \text{loopif}(\tau_e \llbracket e \rrbracket_{t_0}, \tau_c \llbracket c_0 \rrbracket_{t_0}, \tau_c \llbracket c_1 \rrbracket_{t_0}, t) \end{aligned}$$

Let $\mathcal{E} \llbracket e \rrbracket : \text{ExtEnv} \times \text{TEnv} \rightarrow \mathbb{R} + \mathbb{B}$, $\mathcal{C} \llbracket c \rrbracket : \text{ExtEnv} \times \text{TEnv} \rightarrow \mathbb{N} \rightarrow \text{Party} \times \text{Party} \rightarrow \mathbb{R}$, and $\mathcal{IL} \llbracket il \rrbracket : \text{ExtEnv} \times \text{TEnv} \times (\mathbb{N} \rightarrow \mathbb{R}) \times \text{Party} \times \text{Party} \rightarrow \mathbb{R} + \mathbb{B}$ define the semantics of the contract expression sublanguage, the semantics of contracts, and the semantics of the payoff language, respectively. We also assume a function $\text{HOR} : \text{Contr} \times \text{TEnv} \rightarrow \mathbb{N}$ that returns a conservative upper bound on the length of a contract. The compilation function satisfies the following properties:

Theorem 1 (Soundness). *Assume parties p_1 and p_2 and discount function $d : \mathbb{N} \rightarrow \mathbb{R}$.*

- If $\tau_e \llbracket e \rrbracket_0 = il$ and $\mathcal{E} \llbracket c \rrbracket_{\rho, \delta} = v_1$ and $\mathcal{IL} \llbracket il \rrbracket_{\rho, \delta, d, p_1, p_2} = v_2$ then $v_1 = v_2$.
- If $\tau_c \llbracket c \rrbracket_0 = il$ and $\mathcal{C} \llbracket c \rrbracket_{\rho, \delta} = \text{trace}$, where $\text{trace} : \mathbb{N} \rightarrow \text{Party} \times \text{Party} \rightarrow \mathbb{R}$, and $\mathcal{IL} \llbracket il \rrbracket_{\rho, \delta, d, p_1, p_2} = v$ then $\sum_{t=0}^{\text{HOR}(c, \delta)} d(t) \times \text{trace}(t)(p_1, p_2) = v$.

To avoid recompilation of a contract when time moves forward, we define a function `cutPayoff()`. This function is defined recursively on the syntax of intermediate language expressions. The only interesting case is the case for `payoff` expressions:

$$\text{cutPayoff}(\text{payoff}(t, p_1, p_2)) = \text{if}(\text{lt}(t, \text{now}), 0, \text{payoff}(t, p_1, p_2))$$

The function guards the `payoff` expression with a condition guarding whether this payoff should have effect. For the remaining cases, the function recurses on subexpressions and returns otherwise unmodified expressions.

Avoiding recompilation can significantly improve performance especially on GPGPU devices. On the other hand, additional conditionals are introduced, which results in a number of additional checks at runtime. Experiments conducted with “hand-compiled” OpenCL code, which

was semantically equivalent to the payoff language code, show that for the simple contracts, like European options, additional conditions, introduced by `cutPayoff()` do not significantly influence performance. The estimated overhead was around 2.5 percent, while compilation time is in the order of a magnitude bigger than the total execution time.

5 Conclusion

This work extends the certified contract management system of [3] with template expressions, which allows for drastic performance improvements and reusability in terms of the concept of instruments. Along with changes to the contract language, we developed a formalization of the payoff intermediate language in Coq. Our approach introduces the abstract syntax of the payoff language as an inductive data type with the semantics of the payoff language and the compilation from the contract DSL defined as partial functions (using Coq's `Option` data type).

A number of important properties (including soundness) of the translation from contracts to the payoff language have been proven in Coq. Moreover, Coq's code extraction mechanism is used to obtain a certified compiler implementation in the Haskell programming language. We are currently working on establishing the important property of the new time-parameterized payoff evaluation function, which states that given a contract c and its compilation into payoff code il , evaluating il at time t gives us the same result (upto discounting) as first advancing c to time t , then compiling it to payoff code, and then evaluating the result at time 0.

References

- [1] Jesper Andersen, Ebbe Elsborg, Fritz Henglein, Jakob Grue Simonsen, and Christian Stefansen. Compositional specification of commercial contracts. *International Journal on Software Tools for Technology Transfer*, 8(6):485–516, 2006.
- [2] B.R.T Arnold, A. Van Deursen, and M. Res. An algebraic specification of a language for describing financial products. In *ICSE-17 Workshop on Formal Methods Application in Software Engineering*, pages 6–13, 1995.
- [3] Patrick Bahr, Jost Berthold, and Martin Elsmann. Certified symbolic management of financial multi-party contracts. In *Proceedings of the 20th ACM SIGPLAN International Conference on Functional Programming*, ICFP'2015, pages 315–327, September 2015.
- [4] Simon Frankau, Diomidis Spinellis, Nick Nassuphis, and Christoph Burgard. Commercial uses: Going functional on exotic trades. *Journal of Functional Programming*, 19(1):27–45, 2009.
- [5] Troels Henriksen, Martin Elsmann, and Cosmin E Oancea. Size slicing: a hybrid approach to size inference in Futhark. In *Proceedings of the 3rd ACM SIGPLAN workshop on Functional high-performance computing*, pages 31–42. ACM, 2014.
- [6] Tom Hvitved, Felix Klaedtke, and Eugen Zalescu. A trace-based model for multiparty contracts. *The Journal of Logic and Algebraic Programming*, 81(2):72–98, 2012.
- [7] Satoshi Nakamoto. Bitcoin: A peer-to-peer electronic cash system, 2008.
- [8] Simon Peyton Jones, Jean-Marc Eber, and Julian Seward. Composing contracts: an adventure in financial engineering (functional pearl). In *Proceedings of the ACM SIGPLAN International Conference on Functional Programming*, ICFP'2000, September 2000.
- [9] Gavin Wood. Ethereum: A secure decentralised generalised transaction ledger, 2015. Homestead revision, Founder, Ethereum & Ethcore, gavin@ethcore.io.