Foundations of System Development

Martin Wirsing

in cooperation with
Axel Rauschmayer

WS 05/06

Goals

Get to know

- equational deduction and proof by rewriting
- confluence of rewriting systems
- canonical term algebras

Equational Deduction

How can we prove that a given equation e is a theorem of an equational theory T? We can do so by giving appropriate inference rules.

Given an equational theory $T = (\Sigma, E)$, with Σ as always assumed sensible, and a set E of possibly conditional Σ-equations, we say that an unconditional equation $(\forall X)\, t = t'$ is derivable from E, written $E \vdash (\forall X)\, t = t'$, iff it can be obtained by finite application of the following rules:

- **Reflexivity**. For each family of variables X with $t \in T_\Sigma(X)$,

$$\frac{}{(\forall X)\, t = t}$$

Equational Deduction

- **Symmetry**.

$$\frac{(\forall X)\, t = t'}{(\forall X)\, t' = t}$$

- **Transitivity**.

$$\frac{(\forall X)\, t = t' \qquad (\forall X)\, t' = t''}{(\forall X)\, t = t''}$$

- **Congruence**. For $f : s_1 \dots s_n \to s$, $f : s'_1 \dots s'_n \to s'$ in Σ with $s_1 \dots s_n s \equiv_\leq s'_1 \dots s'_n s'$, and for $t_i \in T_\Sigma(X)_{s_i}$, $t'_i \in T_\Sigma(X)_{s'_i}$, $1 \leq i \leq n$,

$$\frac{(\forall X)\, t_1 = t'_1 \quad \dots \quad (\forall X)\, t_n = t'_n}{(\forall X)\, f(t_1, \dots, t_n) = f(t'_1, \dots, t'_n)}$$

Equational Deduction

- **Modus ponens.** For families of variables X, Y, a substitution $\theta : X \to T_\Sigma(Y)$, and a conditional equation in E

$$(\forall X)\ t = t' \Leftarrow u_1 = v_1 \wedge \ldots \wedge u_n = v_n$$

$$\frac{(\forall Y)\ \overline{\theta}(u_1) = \overline{\theta}(v_1) \quad \ldots \quad (\forall Y)\ \overline{\theta}(u_n) = \overline{\theta}(v_n)}{(\forall Y)\ \overline{\theta}(t) = \overline{\theta}(t')}$$

We call $\vdash$ the provability relation, and define $T \vdash (\forall X)\ t = t'$ iff $E \vdash (\forall X)\ t = t'$.

Soundness and Completeness Theorem

Theorem

Let (Σ,E) be a (conditional equational specification, and let t,t' be two ground Σ-terms.

Then

T ⊢ t = t' if, and only, if T |= t = t' .

Equational Simplification

- From a computational point of view it is in general very inefficient to carry out **equational proofs** with the general inference rules that we have already discussed. It may be appropriate to use them for **theorem proving purposes**, but it would **not** be reasonable to use them for **equational programming purposes**.
- In Maude, the distinction between **theories**, with loose semantics, and **modules**, with initial algebra semantics, is not only a distinction of loose vs. initial, but it is also one of inefficient vs. efficient executability.
- That is, in modules we assume that equations can be efficiently executed by **equational simplification** from left to right; for theories we make no such assumptions.

Equational Simplification

Let $T = (\Sigma, E)$ be an equational theory. We say that the equations in E are admissible as equational simplification rules if each equation $(\forall X)\, t = t' \Leftarrow u_1 = v_1 \wedge \ldots \wedge u_n = v_n$ in E satisfies the following two properties:

- *fewer variables on the right side and condition*, that is, the families of variables $\mathrm{vars}(t')$, $\mathrm{vars}(u_i)$, and $\mathrm{vars}(v_i)$, $1 \leq i \leq n$, are all contained in $\mathrm{vars}(t) = X$, where, $\mathrm{vars}(t)$ denotes the family of variables actually appearing in t.
- *sort decreasingness*, that is, for any substitution $\theta : X \longrightarrow T_\Sigma(Y)$, and any $s \in S$, if $\overline{\theta}(t) \in T_\Sigma(Y)_s$, then $\overline{\theta}(t') \in T_\Sigma(Y)_s$.

Term Positions and Subterm Occurrences

- Each Σ-term can be viewed as a **tree** in the obvious way.
- Each **position** in the tree can be denoted by a string of natural numbers, indicating the path that we must follow to go down in the tree and reach the position.
- At each level, the corresponding number in the string indicates the argument position on which we must go down, to finally reach the desired position.

  For example, the term

  f(h(d), q(b, a), g(a, k(c)))

  has the subterm k(c) at position 3.2.
- Given a Σ-term t and a position π we denote by **t/π** the **subterm occurring at that position**; thus,

  f(h(d), q(b, a), g(a, k(c)))/3.2 = k(c).

Term Replacement

Given a Σ-term t, a position π in it, and another Σ-term u we denote by

$$t[\pi \leftarrow u]$$

the replacement of t/π by u in t at position π. For example, consider the term $t = s(x) + s(s(y))$, its 2.1 position, and the term $u = x + z$. Then $t[2.1 \leftarrow u] = s(x) + s(x + z)$.

Note that, in general, $t[\pi \leftarrow u]$ need not be a well-formed Σ-term. However, it is a well-formed Σ-term if, for any $s \in S$,

$$t/\pi \in T_\Sigma(\mathsf{vars}(u) \cup \mathsf{vars}(t/\pi))_s \quad \Rightarrow \quad u \in T_\Sigma(\mathsf{vars}(u) \cup \mathsf{vars}(t/\pi))_s.$$

The Equational Rewriting Relation

Let $T = (\Sigma, E)$ be a theory whose equations E are admissible as equational simplification rules. Then we can use them from left to right to hopefully bring terms to a simpler form which can be interpreted as their evaluation.

This process is called equational simplification, or equational rewriting. For any S-indexed family X of variables, this defines two binary relations on $T_\Sigma(X)$, $\longrightarrow_E$, and its reflexive and transitive closure $\overset{*}{\longrightarrow}_E$. $\longrightarrow_E$ is defined recursively as follows:

For $t, t' \in T_\Sigma(X)$, we have $t \longrightarrow_E t'$ iff, either:

The Equational Rewriting Relation

- there is an equation $(\forall \text{vars}(u))\ u = v$ in E, a position π in t, and a substitution $\theta : \text{vars}(u) \longrightarrow T_\Sigma(X)$ such that, $t/\pi = \overline{\theta}(u)$, and $t' = t[\pi \leftarrow \overline{\theta}(v)]$, or

- there is a conditional equation $(\forall \text{vars}(u))\ u = v \Leftarrow u_1 = v_1 \wedge \ldots \wedge u_n = v_n$ in E, a position π in t, a substitution $\theta : \text{vars}(u) \longrightarrow T_\Sigma(X)$, and terms $w_i \in T_\Sigma(X)$, $1 \leq i \leq n$, such that:

  1. (*satisfaction of the condition*) $\overline{\theta}(u_i) \overset{*}{\longrightarrow}_E w_i$ and $\overline{\theta}(v_i) \overset{*}{\longrightarrow}_E w_i$, $1 \leq i \leq n$, and
  2. (*matching and replacement*) $t/\pi = \overline{\theta}(u)$, and $t' = t[\pi \leftarrow \overline{\theta}(v)]$.

Remarks on the Equational Rewriting Relation

- Note that, because of our assumption about sort-decreasingness of the equational simplification rules in E, whenever we have an equation $u = v$, with or without a condition, in E, and a term $t \in T_\Sigma(X)$ such that $t/\pi = \overline{\theta}(u)$, then $t[\pi \leftarrow \overline{\theta}(v)]$ is a well-formed Σ-term.

- Note also that $\mathsf{vars}(t[\pi \leftarrow \overline{\theta}(v)]) \subseteq \mathsf{vars}(t)$, and therefore, whenever $t \longrightarrow_E^* t'$ we have $\mathsf{vars}(t') \subseteq \mathsf{vars}(t)$.

Notation: We define the binary relation $t \downarrow_E t'$ on pairs of terms $t, t' \in T_\Sigma(X)$ by,

$$t \downarrow_E t' \quad \Leftrightarrow \quad (\exists w \in T_\Sigma(X))\ t \longrightarrow_E^* w \ \wedge\ t' \longrightarrow_E^* w.$$

Soundness of the Equational Rewriting Relation

Theorem: Equational rewriting is a sound inference system, in the sense that

$$t \longrightarrow_E^* t' \quad \Rightarrow \quad E \vdash (\forall \mathsf{vars}(t))\ t = t'.$$

Proof by induction on the number of the rewrite steps.

Confluence

Suppose we have a theory $T = (\Sigma, E)$ whose equations are admissible as equational simplification rules. Then we can do sound inference by equational simplification with E, but the equations E may still be quite unusable, because depending on the order and place of equation application we may get different and unrelatable results. That is, in general, equational simplification can be nondeterministic.

The minimum requirement to make equational simplification deterministic is confluence, also called the Church-Rosser property. We say that E is confluent iff whenever we have $t \longrightarrow_E^* u$ and $t \longrightarrow_E^* v$, then $u \downarrow_E v$. We call E ground confluent iff the above property is guaranteed only for terms without variables $t \in T_\Sigma$.

Completeness of Confluent Equations

We have seen that equational simplification is a sound inference system. But can we prove any equation just by simplification? The answer is, yes, if we are confluent.

Theorem: Let the equations in E be admissible as equational simplification rules and confluent. Then,

$$E \vdash (\forall \mathsf{vars}(t) \cup \mathsf{vars}(t'))t = t' \quad \Leftrightarrow \quad t \downarrow_E t'.$$

Proof: The ($\Leftarrow$) part follows easily from the soundness theorem for equational simplification, the Variable Expansion Lemma, Symmetry, and Transitivity. The proof of the ($\Rightarrow$) part is by induction on the depth of the proof term. Proof of transitivity needs the confluence property.

Termination

Equational theories $T = (\Sigma, E)$ whose equations are ground confluent can be viewed as equational programs, and support a style of functional programming in which expresions are evaluated by simplification. This is what OBJ "objects" and Maude functional modules are: equational programs evaluated by simplification.

In general, however, such programs can be nonterminating. Terminating equational programs are obtained when the rewriting relation is terminating.

We say that the equations E are terminating as simplification rules when there is no infinite chain of rewrites

$$t \longrightarrow_E t_1 \longrightarrow_E t_2 \ldots t_{n-1} \longrightarrow_E t_n \longrightarrow \ldots$$

Canonical Forms

If the equations in $T = (\Sigma, E)$ are confluent (resp. ground confluent) and terminating, then for each term $t \in T_\Sigma(X)$ (resp. $t \in T_\Sigma$) there is a unique term $can_E(t) \in T_\Sigma(X)$ (resp. $can_E(t) \in T_\Sigma$) called its canonical form such that:

- $t \overset{*}{\longrightarrow}_E can_E(t)$, and
- $can_E(t)$ cannot be further rewritten.

Indeed, by the termination assumption, we can always simplify t to such a term; and uniqueness is then forced by the confluence (resp. ground confluence) property. The term $can_E(t)$ should be understood as the value returned by the equational program $T = (\Sigma, E)$ for the input expression t.

The Canonical Term Algebra

If the equations in $T = (\Sigma, E)$ are ground confluent and terminating, then the S-indexed family of terms $Can_{\Sigma/E}$ with

$$|Can_{\Sigma/E,s} = \{can_E(t) \mid t \in T_{\Sigma,s}\}$$

can be made into a Σ-algebra as follows:

- for each constant $a : nil \longrightarrow s$ in Σ, we define $a_{Can_{\Sigma/E}} = can_E(a)$, and
- for each $f : s_1 \ldots s_n \rightarrow s$ in Σ, and $t_i \in Can_{\Sigma/E,s_i}$, $1 \leq i \leq n$, we define $f^{s_1 \ldots s_n, s}_{Can_{\Sigma/E}}(t_1, \ldots, t_n) = can_E(f(t_1, \ldots, t_n))$.

Example of Canonical Term Algebra

We shall see in what follows that the canonical term algebra is initial among those satisfying the equations E.

The canonical term algebra $Can_{\Sigma/E}$ is in some sense the most intuitive representation of the initial algebra from a computational point of view.

For example, the equations in the NATURAL module are ground confluent and terminating. Its canonical forms are the natural numbers in Peano notation. And its operations are the successor and addition functions.

Indeed, given two Peano natural numbers n, m the general definition of $f^{s_1 \ldots s_n, s}_{Can_{\Sigma/E}}$ specializes for $f = \_ + \_$ to the definition of addition, $n +_{Can_{\mathtt{NATURAL}}} m = can_{\mathtt{NATURAL}}(n + m)$, so that $\_ +_{Can_{\mathtt{NATURAL}}} \_$ is the addition function.

Example of Canonical Term Algebra

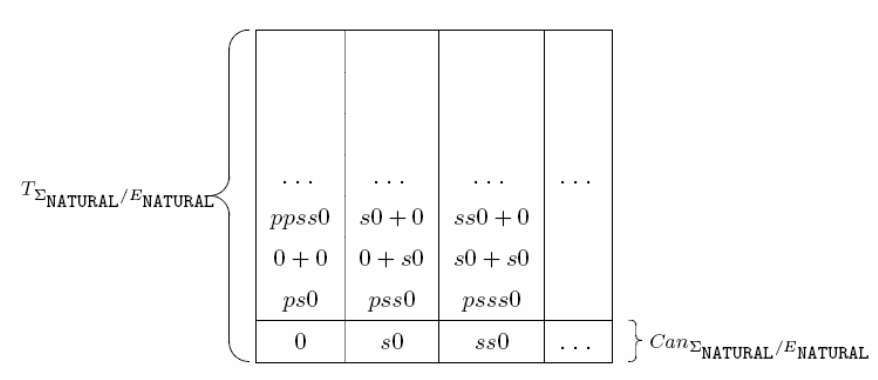

The Canonical Term Algebra is Initial

Theorem: Assume that the equations in $T = (\Sigma, E)$ are ground confluent and terminating. Then, $Can_{\Sigma/E}$ is isomorphic to $T_{\Sigma/E}$ and is therefore initial in $\mathbf{Alg}_{(\Sigma,E)}$.

Proof: Show that $T_{\Sigma/E}$ and $Can_{\Sigma/E}$ are isomorphic:

Define for each $s \in S$ a function $can_{E,s} : T_{\Sigma/E,s} \longrightarrow Can_{\Sigma/E,s}$ by, $can_{E,s}[t] = can_E(t)$. This is independent of the choice of t and therefore well-defined because, by definition of $\equiv_E$, we have $t \equiv_E t'$ iff $E \vdash (\forall \emptyset)\ t = t'$ iff (by E confluent) $t \downarrow_E t'$, iff $can_E(t) = can_E(t')$. Furthermore, $can_{E,s}$ is surjective by construction and injective because of the above chain of equivalences; therefore it is bijective.

Summary

- Execution of equational specifications in Maude is based on equational simplification.
- If the equational rules are confluent, then simplification is equivalent to equational deduction.
- If the equational rules are confluent and terminating, then the canonical term algebra is an initial model.