Linear Arithmetic

I start with a syntax that already contains $-$, $\leq$, $<$, $\geq$, $\not\approx$ and $\mathbb{Q}$. All these functions and relations are indeed expressible by first-order forumulas over 0, 1, $\approx$, and $>$.

For the semantics there are two approaches. Either providing axioms, i.e., closed formulas, for the above symbols and then considering all algebras satisfying the axioms, or fixing one particular algebra or a class of algebras.

6.2.1 Definition (LA Syntax)

The syntax of LA is

$$\Sigma_{\mathsf{LA}} = (\{\mathsf{LA}\}, \{0, 1, +, -\} \cup \mathbb{Q}, \{\leq, <, \not\approx, >, \geq\})$$

where $-$ is unitary and all other symbols have the usual arities.

Terms and formulas over Σ_{LA} are built in the classical free first-order way, see Section 3.1. All first-order notions, i.e., terms, atoms, equations, literals, clauses, etc. carry over to LA formulas. The atoms and terms built over the LA signature are written in their standard infix notation, i.e., I write $3 + 5$ instead of $+(3, 5)$. Note that the signature does not contain multiplication. A term $3x$ is just an abbreviation for a term $x + x + x$.

6.2.2 Definition (Linear Rational Arithmetic Standard Semantics)

The Σ_{LA} algebra $\mathcal{A}_{\mathsf{LRA}}$ is defined by $\mathsf{LA}^{\mathcal{A}_{\mathsf{LRA}}} = \mathbb{Q}$ and all other signature symbols are assigned the standard interpretations over the rationals.

Due to the expressive LA language there is no need for negative literals, because $(\neg <)^{\mathcal{A}_{\mathsf{LRA}}} = (\geq)^{\mathcal{A}_{\mathsf{LRA}}}$, $(\neg >)^{\mathcal{A}_{\mathsf{LRA}}} = (\leq)^{\mathcal{A}_{\mathsf{LRA}}}$, and $(\neg \approx)^{\mathcal{A}_{\mathsf{LRA}}} = (\not\approx)^{\mathcal{A}_{\mathsf{LRA}}}$.

Note the difference between the above standard semantics over Σ_{LA} and the free first-order semantics over Σ_{LA}, Definition 3.2.1. The equation $3 + 4 \approx 5$ has a model in the free first-order semantics, hence it is satisfiable, whereas in the standard model of linear rational arithmetic, Definition 6.2.2, the equation $3 + 4 \approx 5$ is false.

In addition, with respect to the standard LRA semantics the definitions of validity, satisfiability coincide with truth and the definition of unsatisfiability coincides with falsehood. This is the result of a single algebra semantics.

Fourier-Motzkin Quantifier Elimination

It is decidable whether a first-order formula over Σ_{LA} is true or false in the standard LRAsemantics. This was first discovered in 1826 by J. Fourier and re-discovered by T. Motzkin in 1936 and is called FM for short. Note that validity of a Σ_{LA} formula with respect to the standard semantics is undecidable

Similar to Congruence Closure, Section 6.1, the starting point of the procedure is a conjunction of atoms without atoms of the form $\not\approx$. These will eventually be replaced by a disjunction, i.e., an atom $t \not\approx s$ is replaced by $t < s \vee t > s$.

Every atom over the variables $x, y_1, \ldots, y_n$ can be converted into an equivalent atom $x \circ t[\vec{y}]$ or $0 \circ t[\vec{y}]$, where $\circ \in \{<, >, \leq, \geq, \approx, \not\approx\}$ and $t[\vec{y}]$ has the form $\sum_i q_i \cdot y_i + q_0$ where $q_i \in \mathbb{Q}$.

In other words, a variable x can be either isolated on one side of the atom or eliminated completely. This is the starting point of the FM calculus deciding a conjunction of LA atoms without $\not\approx$ modulo the isolation of variables and the reduction of ground formulas to $\top$, $\bot$.

The calculus operates on a set of atoms N. The normal forms are conjunctions of atoms $s \circ t$ where s, t do not contain any variables. These can be obviously eventually reduced to $\top$ or $\bot$. The FM calculus consists of two rules:

Substitute $N \uplus \{x \approx t\} \Rightarrow_{\mathsf{FM}} N\{x \mapsto t\}$

provided x does not occur in t

Eliminate $N \uplus \bigcup_i \{x \circ_i^1 t_i\} \uplus \bigcup_j \{x \circ_j^2 s_j\} \Rightarrow_{\mathsf{FM}}$
$N \cup \bigcup_{i,j} \{t_i \circ_{i.j} s_j\}$

provided x does not occur in N nor in the t_i, s_j, $\circ_i^1 \in \{<, \leq\}$, $\circ_j^2 \in \{>, \geq\}$, and $\circ_{i,j} = >$ if $\circ_i^1 = <$ or $\circ_j^2 = >$, and $\circ_{i,j} = \geq$ otherwise

If all variables in N are implicitely existentially quantified, i.e., N stands for $\exists \vec{x}.N$, then the above two rules constitute a sound and complete decision procedure for conjunctions of LA atoms without $\not\approx$.

6.2.3 Lemma (FM Termination on a Conjunction of Atoms)

FM terminates on a conjunction of atoms.

6.2.4 Lemma (FM Soundness and Completeness on a Conjunction of Atoms)

$N \Rightarrow^*_{\mathsf{FM}} \top$ iff $\mathcal{A}_{\mathsf{LRA}} \models \exists \vec{x}.N$.
$N \Rightarrow^*_{\mathsf{FM}} \bot$ iff $\mathcal{A}_{\mathsf{LRA}} \not\models \exists \vec{x}.N$.

The FM calculus on conjunctions of atoms can be extended to arbitrary closed LRA first-order formulas ϕ. I always assume that different quantifier occurrences in ϕ bind different variables. This can always be obtained by renaming one variable.

The first step is to eliminate $\top$, $\bot$ from ϕ and to transform ϕ in negation normal form, see Section 3.9. The resulting formula only contains the operators $\forall$, $\exists$, $\wedge$, $\vee$, $\neg$, where all negation symbols occur in front of atoms.

The following rule can be used to remove the negation symbols as well:

ElimNeg $\chi[\neg\, s \circ^1 t]_p \;\Rightarrow_{\mathsf{FM}}\; \chi[s \circ^2 t]_p$

where the pairs $(\circ_1, \circ_2)$ are given by pairs $(<, \geq), (\leq, >), (\approx, \not\approx)$ and their symmetric variants

The above two FM rules on conjunctions cannot cope with atoms $s \not\approx t$, so they are eliminated as well:

Elim$\not\approx$ $\chi[s \not\approx t]_p \;\Rightarrow_{\mathsf{FM}}\; \chi[s < t \vee s > t]_p$

The next step is to compute a *Prenex Normal Form*, a formula $\{\exists, \forall\} x_1 \ldots \{\exists, \forall\} x_n.\phi$ where ϕ does not contain any quantifiers. This can be done by simply applying the mini-scoping rules, see Section 3.9, in the opposite direction:

Prenex1 $\chi[(\forall x.\psi_1) \circ \psi_2]_p \Rightarrow_{\mathsf{FM}} \chi[\forall x.(\psi_1 \circ \psi_2)]_p$

provided $\circ \in \{\wedge, \vee\}$, $x \notin \mathsf{fvars}(\psi_2)$

Prenex2 $\chi[(\exists x.\psi_1) \circ \psi_2]_p \Rightarrow_{\mathsf{FM}} \chi[\exists x.(\psi_1 \circ \psi_2)]_p$

provided $\circ \in \{\wedge, \vee\}$, $x \notin \mathsf{fvars}(\psi_2)$

Prenex3 $\chi[(\forall x.\psi_1) \wedge (\forall y.\psi_2)]_p \Rightarrow_{\mathsf{FM}} \chi[\forall x.(\psi_1 \wedge \psi_2\{y \mapsto x\})]_p$

Prenex4 $\chi[(\exists x.\psi_1) \vee (\exists y.\psi_2)]_p \Rightarrow_{\mathsf{FM}} \chi[\exists x.(\psi_1 \vee \psi_2\{y \mapsto x\})]_p$

where Prenex3 and Prenex4 are preferred over Prenex1 and Prenex2.

Finally, for the resulting formula $\{\exists, \forall\} x_1 \ldots \{\exists, \forall\} x_n.\phi$ in prenex normal form the FM algorithm computes a DNF of ϕ by exhaustively applying the rule PushConj, Section 2.5.2.

The result is a formula $\{\exists, \forall\} x_1 \ldots \{\exists, \forall\} x_n.\phi$ where ϕ is a DNF of atoms without containing an atom of the form $s \not\approx t$.

Then FM on formulas considers the quantifiers iteratively in an innermost way. For the formula $\{\exists, \forall\} x_1 \ldots \{\exists, \forall\} x_n.\phi$ always the innermost quantifier $\{\exists, \forall\} x_n$ is considered.

If it is an existential quantifier, $\exists x_n$, then the FM rules Substitute, Eliminate are applied to the variable x_n for each conjunct C_i of $\phi = C_1 \vee \ldots \vee C_n$. The result is a formula $\{\exists, \forall\} x_1 \ldots \{\exists, \forall\} x_{n-1}.(C'_1 \vee \ldots \vee C'_n)$ which is again in prenex DNF. Furthermore, by Lemma 6.2.4 it is equivalent to $\{\exists, \forall\} x_1 \ldots \{\exists, \forall\} x_n.\phi$.

If the innermost quantifier is a universal quantifier $\forall x_n$, then the formula is replaced by $\{\exists, \forall\} x_1 \ldots \{\exists, \forall\} x_{n-1} \neg \exists x_n . \neg \phi$ and the above steps for negation normal form and DNF are repeated for $\neg \phi$ resulting in an equivalent formula $\{\exists, \forall\} x_1 \ldots \{\exists, \forall\} x_{n-1} \neg \exists x_n . \phi'$ where ϕ' is in DNF and does not contain negation symbols nor atoms $s \not\approx t$.

Then the FM rules Substitute, Eliminate are applied to the variable x_n for each conjunct C_i of $\phi' = C_1 \vee \ldots \vee C_n$. The result is an equivalent formula $\{\exists, \forall\} x_1 \ldots \{\exists, \forall\} x_{n-1} . \neg (C'_1 \vee \ldots \vee C'_n)$. Finally, the above steps for negation normal form and DNF are repeated for $\neg (C'_1 \vee \ldots \vee C'_n)$ resulting in an equivalent formula $\{\exists, \forall\} x_1 \ldots \{\exists, \forall\} x_{n-1} . \phi''$ where ϕ' is in DNF and does not contain negation symbols nor atoms $s \not\approx t$. This completes for FM decision procedure for LRA formulas.

Every LRA formula can by reduced to $\top$ or $\bot$ via the FM decision procedure. Therefore LRA is called a *complete* theory, i.e., every closed formula over the signature of LRA is either true or false.

LA formulas over the rationals and over the reals are indistinguishable by first-order formulas over the signature of LRA. These properties do not hold for extended signatures, e.g., then additional free symbols are introduced. Furthermore, FM is no decision procedures over the integers, even if the LA syntax is restricted to integer constants.

FM Complexity

The complexity of the FM calculus depends mostly on the quantifier alternations in $\{\exists, \forall\} x_1 \ldots \{\exists, \forall\} x_n.\phi$.

In case an existential quantifier $\exists$ is eliminated, the formula size grows worst-case quadratically, therefore $O(n^2)$ runtime. For m quantifiers $\exists \ldots \exists$: a naive implementation needs worst-case $O(n^{2^m})$ runtime. It is not known whether an optimized implementation with simply exponential runtime is possible.

If there are m quantifier alternations $\exists\forall\exists\forall \ldots \exists\forall$, a CNF to DNF conversion is required after each step. Each conversion has a worst-case exponential run time, see Section 2.5. Therefore, the overall procedure has a worst-case non-elementary runtime.

Simplex

The Simplex algorithm is the prime algorithm for solving optimization problems of systems of linear inequations over the rationals. For automated reasoning optimization at the level of conjunctions of inequations is not in focus. Rather, solvability of a set of linear inequations as a subproblem of some theory combination is the typical application. In this context the simplex algorithm is useful as well, due to its incremental nature. If an inequation $t \circ c$, $\circ \in \{\leq, \geq, <, >\}$, $t = \sum a_i x_i$, $a_i, c \in \mathbb{Q}$, is added to a set N of inequations where the simplex algorithm has already found a solution for N, the algorithm needs not to start from scratch. Instead it continues with the solution found for N. In practice, it turns out that then typically only few steps are needed to derive a solution for $N \cup \{t \circ d\}$ if it exists.

Firstly, the problem is rescritcted to non-strict inequations. Starting point is a set N (conjunction) of (non-strict) inequations of the form $(\sum_{x_j \in X} a_{i,j} x_j) \circ_i c_i$ where $\circ_i \in \{\geq, \leq\}$ for all i. Note that an equation $\sum a_i x_i = c$ can be encoded by two inequations $\{\sum a_i x_i \leq c, \sum a_i x_i \geq c\}$.

The variables occurring in N are assumed to be totally ordered by some ordering $\prec$. The ordering $\prec$ will eventualy guarantee termination of the simplex algorithm, see Definition 6.2.10 and Theorem 6.2.11 below. I assume the x_j to be all different, without loss of generality $x_j \prec x_{j+1}$, and I assume that all coefficients are normalized by the gcd of the $a_{i,j}$ for all j: if the gcd is different from 1 for one inequation, it is used for division of all coefficients of the inequation.

The goal is to decide whether there exists an assignment β from the x_j into $\mathbb{Q}$ such that

$$\mathsf{LRA}(\beta) \models \bigwedge_i [(\sum_{x_j \in X} a_{i,j} x_j) \circ_i c_i]$$

or equivalently, $\mathsf{LRA}(\beta) \models N$. So the x_j are free variables, i.e., placeholders for concrete values, i.e., existentially quantified.

The first step is to transform the set N of inequations into two disjoint sets E, B of equations and simple bounds, respectively. The set E contains equations of the form $y_i \approx \sum_{x_j \in X} a_{i,j} x_j$, where the y_i are fresh and the set B contains the respective simple bounds $y_i \circ_i c_i$. In case the original inequation from N was already a simple bound, i.e., of the form $x_j \circ_j c_j$ it is simply moved to B. If in N left hand sides of ineqations $(\sum_{x_j \in X} a_{i,j} x_j) \circ_i c_i$ are shared, it is sufficient to introduce one equation for the respective left hand side. The y_i are also part of the total ordering $\prec$ on all variables.

The two representations are equivalent:

$$\mathrm{LRA}(\beta) \models N$$

iff

$$\mathrm{LRA}(\beta[y_i \mapsto \beta(\textstyle\sum_{x_j \in X} a_{i,j} x_j)]) \models E$$

and

$$\mathrm{LRA}(\beta[y_i \mapsto \beta(\textstyle\sum_{x_j \in X} a_{i,j} x_j)]) \models B.$$

Given E and B a variable z is called *dependent* if it occurs on the left hand side of an equation in E, i.e., there is an equation $(z \approx \sum_{x_j \in X} a_{i,j} x_j) \in E$, and in case such a defining equation for z does not exist in E the variable z is called *independent*. Note that by construction the initial y_i are all dependent and do not occur on the right hand side of an equation.

Given a dependant variable x, an independent variable y, and a set of equations E, the *pivot* operation exchanges the roles of x, y in E where y occurs with non-zero coefficient in the defining equation of x. Let $(x \approx ay + t) \in E$ be the defining equation of x in E. When writing $(x \approx ay + t)$ for some equation, I always assume that $y \notin \mathsf{vars}(t)$. Let E' be E without the defining equation of x. Then

$$\mathsf{piv}(E, x, y) := \{y \approx \frac{1}{a}x + \frac{1}{-a}t\} \cup E'\{y \mapsto (\frac{1}{a}x + \frac{1}{-a}t)\}$$

Given an assignment β, an independent variable y, a rational value c, and a set of equations E then the *update* of β with respect to y, c, and E is

$$\mathsf{upd}(\beta, y, c, E) := \beta[y \mapsto c, \{x \mapsto \beta[y \mapsto c](t) \mid x \approx t \in E\}]$$

A Simplex problem state is a quintuple $(E; B; \beta; S; s)$ where E is a set of equations; B a set of simple bounds; β an assignment to all variables in E, B; S a set of derived bounds, and s the status of the problem with $s \in \{\top, \text{IV}, \text{DV}, \bot\}$. The state $s = \top$ indicates that $\text{LRA}(\beta) \models S$; the state $s = \text{IV}$ that potentially $\text{LRA}(\beta) \not\models x \circ c$ for some independent variable x, $x \circ c \in S$; the state $s = \text{DV}$ that $\text{LRA}(\beta) \models x \circ c$ for all independent variables x, $x \circ c \in S$, but potentially $\text{LRA}(\beta) \not\models x' \circ c'$ for some dependent variable x', $x' \circ c' \in S$; and the state $s = \bot$ that the problem is unsatisfiable.

The following states can be distinguished:

$(E; B; \beta_0; \emptyset; \top)$	is the start state for N and its transformation into E, B, and assignment $\beta_0(x) := 0$ for all $x \in \mathsf{vars}(E \cup B)$
$(E; \emptyset; \beta; S; \top)$	is a final state, where $\mathsf{LRA}(\beta) \models E \cup S$ and hence the problem is solvable
$(E; B; \beta; S; \bot)$	is a final state, where $E \cup B \cup S$ has no model

The important invariants of the simplex rules are:

(i) for every dependent variable there is exactly one equation in E defining the variable and

(ii) dependent variables do not occur on the right hand side of an equation,

(iii) $\mathrm{LRA}(\beta) \models E$

These invariants are maintained by a pivot (piv) or an update (upd) operation.

EstablishBound $(E; B \uplus \{x \circ c\}; \beta; S; \top) \Rightarrow_{\text{SIMP}} (E; B; \beta; S \cup \{x \circ c\}; \text{IV})$

AckBounds $(E; B; \beta; S; s) \Rightarrow_{\text{SIMP}} (E; B; \beta; S; \top)$

if $\text{LRA}(\beta) \models S$, $s \in \{\text{IV}, \text{DV}\}$

FixIndepVar $(E; B; \beta; S; \text{IV}) \Rightarrow_{\text{SIMP}} (E; B; \text{upd}(\beta, x, c, E); S; \text{IV})$

if $(x \circ c) \in S$, $\text{LRA}(\beta) \not\models x \circ c$, x independent

AckIndepBound $(E; B; \beta; S; \mathrm{IV}) \Rightarrow_{\mathrm{SIMP}} (E; B; \beta; S; \mathrm{DV})$

if $\mathrm{LRA}(\beta) \models x \circ c$, for all independent variables x with bounds $x \circ c$ in S

FixDepVar$\leq$ $(E; B; \beta; S; \mathrm{DV}) \Rightarrow_{\mathrm{SIMP}} (E'; B; \mathrm{upd}(\beta, x, c, E'); S; \mathrm{DV})$

if $(x \leq c) \in S$, x dependent, $\mathrm{LRA}(\beta) \not\models x \leq c$, there is an independent variable y and equation $(x \approx ay + t) \in E$ where ($a < 0$ and $\beta(y) < c'$ for all $(y \leq c') \in S$) or ($a > 0$ and $\beta(y) > c'$ for all $(y \geq c') \in S$) and $E' := \mathrm{piv}(E, x, y)$

FixDepVar$\geq$ $(E; B; \beta; S; \mathrm{DV}) \Rightarrow_{\mathrm{SIMP}} (E'; B; \mathrm{upd}(\beta, x, c, E'); S; \mathrm{DV})$

if $(x \geq c) \in S$, x dependent, $\mathrm{LRA}(\beta) \not\models x \geq c$, there is an independent variable y and equation $(x \approx ay + t) \in E$ where ($a > 0$ and $\beta(y) < c'$ for all $(y \leq c') \in S$) or ($a < 0$ and $\beta(y) > c'$ for all $(y \geq c') \in S$) and $E' := \mathrm{piv}(E, x, y)$

FailBounds $(E; B; \beta; S; \top) \Rightarrow_{\text{SIMP}} (E; B; \beta; S; \bot)$

if there are two contradicting bounds $x \leq c_1$ and $x \geq c_2$ in $B \cup S$ for some variable x

FailDepVar$\leq$ $(E; B; \beta; S; \text{DV}) \Rightarrow_{\text{SIMP}} (E; B; \beta; S; \bot)$

if $(x \leq c) \in S$, x dependent, $\text{LRA}(\beta) \not\models x \leq c$ and there is no independent variable y and equation $(x \approx ay + t) \in E$ where ($a < 0$ and $\beta(y) < c'$ for all $(y \leq c') \in S$) or ($a > 0$ and $\beta(y) > c'$ for all $(y \geq c') \in S$)

FailDepVar$\geq$ $(E; B; \beta; S; \text{DV}) \Rightarrow_{\text{SIMP}} (E; B; \beta; S; \bot)$

if $(x \geq c) \in S$, x dependent, $\beta \not\models_{\text{LA}} x \geq c$ and there is no independent variable y and equation $(x \approx ay + t) \in E$ where (if $a > 0$ and $\beta(y) < c'$ for all $(y \leq c') \in S$) or (if $a < 0$ and $\beta(y) > c'$ for all $(y \geq c') \in S$)

6.2.7 Lemma (Simplex State Invariants)

The following invariants hold for any state $(E_i; B_i; \beta_i; S_i; s_i)$ derived by $\Rightarrow_{\text{SIMP}}$ on a start state $(E_0; B_0; \beta_0; \emptyset; \top)$:

(i) for every dependent variable there is exactly one equation in E defining the variable

(ii) dependent variables do not occur on the right hand side of an equation

(iii) $\text{LRA}(\beta) \models E_i$

(iv) for all independant variables x either $\beta_i(x) = 0$ or $\beta_i(x) = c$ for some bound $x \circ c \in S_i$

(v) for all assignemnts α it holds $\text{LRA}(\alpha) \models E_0$ iff $\text{LRA}(\alpha) \models E_i$

6.2.8 Lemma (Simplex Run Invariants)

For any run of $\Rightarrow_{\text{SIMP}}$ from start state $(E_0; B_0; \beta_0; \emptyset; \top) \Rightarrow_{\text{SIMP}} (E_1; B_1; \beta_1; S_1; s_1) \Rightarrow_{\text{SIMP}} \ldots$:

(i) the set $\{\beta_o, \beta_1, \ldots\}$ is finite

(ii) if the sets of dependent and independent variables for two equational systems E_i, E_j coincide, then $E_i = E_j$

(iii) the set $\{E_o, E_1, \ldots\}$ is finite

(iv) let S_i not contain contradictory bounds, then $(E_i; B_i; \beta_i; S_i; s_i) \Rightarrow_{\text{SIMP}}^{\text{FIV},*}$ is finite

6.2.9 Corollary (Infinite Runs Contain a Cycle)

Let $(E_0; B_0; \beta_0; \emptyset; \top) \Rightarrow_{\mathsf{SIMP}} (E_1; B_1; \beta_1; S_1; s_1) \Rightarrow_{\mathsf{SIMP}} \ldots$ be an infinite run. Then there are two states $(E_i; B_i; \beta_i; S_i; s_i)$, $(E_k; B_k; \beta_k; S_k; s_k)$ such that $i \neq k$ and $(E_i; B_i; \beta_i; S_i; s_i) = (E_k; B_k; \beta_k; S_k; s_k)$.

6.2.10 Definition (Reasonable Strategy)

A *reasonable* strategy prefers FailBounds over EstablishBounds and the FixDepVar rules select minimal variables x, y in the ordering $\prec$.

6.2.11 Theorem (Simplex Soundness, Completeness & Termination)

Given a reasonable strategy and initial set N of inequations and its separation into E and B :

(i) $\Rightarrow_{\text{SIMP}}$ terminates on $(E; B; \beta_0; \emptyset; \top)$,

(ii) if $(E; B; \beta_0; \emptyset; \top) \Rightarrow^*_{\text{SIMP}} (E'; B'; \beta; S; \bot)$ then N has no solution,

(iii) if $(E; B; \beta_0; \emptyset; \top) \Rightarrow^*_{\text{SIMP}} (E'; \emptyset; \beta; B; \top)$ and $(E; \emptyset; \beta; B; \top)$ is a normal form, then $\text{LRA}(\beta) \models N$,

(iv) all final states $(E'; B'; \beta; S; s)$ match either (ii) or (iii).